Quantifying the Benefits of Targeting for Pandemic Response

Sergio Camelo[a], Dragos F. Ciocan[b], Dan A. Iancu[b,c], Xavier S. Warnes[c], and Spyros I. Zoumpoulis[d]

[a]Institute for Computational and Mathematical Engineering, Stanford University
[b]Technology and Operations Management, INSEAD
[c]Operations, Information & Technology, Graduate School of Business, Stanford University
[d]Decision Sciences, INSEAD

ABSTRACT

To respond to pandemics such as COVID-19, policy makers have relied on interventions that target specific age groups or activities. Such targeting is potentially contentious, so rigorously quantifying its benefits and downsides is critical for designing effective and equitable pandemic control policies. We propose a flexible modeling framework and algorithms to compute optimally targeted interventions that coordinate across two dimensions of heterogeneity: age of different groups and the specific activities that individuals engage in during the course of a day. We showcase a complete implementation in a case study focused on the COVID-19 epidemic in the Île-de-France region of France, based on hospitalization, community mobility, social contacts and economic data. We find that optimized dual-targeted policies generate substantial complementarities that lead to Pareto improvements, reducing the number of deaths and the economic losses overall and reducing the time in confinement for each age group, compared to less targeted interventions. These policies have a simple and explainable structure. Since dual-targeted policies could lead to increased discrepancies in the confinements faced by distinct groups, we also quantify the impact of requirements that explicitly limit such disparities, and find that satisfactory trade-offs may be achievable through limited targeting.

Keywords: Pandemic management, Confinement, Targeted interventions, Optimization, COVID-19

Introduction

The COVID-19 pandemic has forced policy makers worldwide to rely on a range of large-scale population confinement measures in an effort to contain the disease spread. In determining these measures, a key recognition has been that substantial differences exist in the health and economic impact produced by different individuals engaged in distinct activities. Targeting confinements to account for such heterogeneity could be an important lever to mitigate a pandemic's impact, but could also lead to potentially contentious and discriminatory measures. This work is aimed at developing a rigorous framework to quantify the benefits and downsides of such targeted interventions, and applying it to the COVID-19 pandemic as a real-world case study.

One real-world contentious example of targeting has been to differentiate confinements based on *age groups*, e.g., sheltering older individuals who might face higher health risks if infected, or restricting younger groups who might create higher infection risks. Such measures, focusing on confinements or other interventions, and targeting age or other population characteristics, have been studied in the literature (1; 2; 3; 4; 5; 6; 7; 8; 9; 10) and implemented in several settings – e.g., with stricter confinements applied to older groups in Finland (11), Ireland (12), Israel (13) and Moscow (14), or curfews applied to children and youth in Bosnia and Herzegovina (15) and Turkey (16) – but some of the measures were eventually deemed unconstitutional and overturned (13; 15). A different example of targeting has stemmed from the recognition that different *activities* (more specifically, population interactions in locations of certain activities), such as work, schooling, transport, leisure, result in significantly different patterns of social contacts and new infections. This has been shown to be critical when modeling pandemic spread (17; 18; 19) and has been recognized in numerous implementations that differentially confine various activities (e.g., closures of schools, workplaces, recreation venues, etc.), and even some that differentiate based on *both* age groups and activities (e.g., dedicated hours when only the senior population was allowed to shop at supermarkets (20)).

As these examples suggest, targeted confinements have merits but also pose potentially significant downsides. On the one hand, targeting can generate improvements in both health and economic outcomes, giving policy makers an improved lever when navigating difficult trade-offs. Additionally, explicitly considering *multiple* dimensions of targeting simultaneously – activities *and* age groups – could overturn some of the prevailing insight that specific age groups should uniformly face stricter confinements. However, such granular policies are more difficult to implement, and could lead to discriminatory and potentially unfair measures. Given that some amount of targeting of activities and age groups is already in place in existing real-world

policy implementations, it seems critical to quantify the relative merits of a policy that (i) targets both age-based population groups *and* activities, and (ii) identifies *optimal* interventions.

We propose a rigorous modeling framework and develop a set of associated algorithms that compute optimally targeted interventions that target both age groups and activities. The framework provides quantifiable answers to the following natural research questions: How large are the health and economic benefits of dual-targeted confinements? Would dual targeting lead to significant synergies, and why? Could dual targeting reduce time in confinement for *every* age group? What is the relationship between the effectiveness and the level of targeting allowed across distinct groups?

We showcase a proof of concept for our framework through a case study calibrated on Île-de-France data – a region of France encompassing Paris with a population of approximately 12 million. The implementation, publicly available at http://insead.arnia.ro, is flexible and portable to other geographies.

Methods

Our framework relies on a flexible model that captures several important real-world considerations. We extend a version of the discretized SEIR (Susceptible-Exposed-Infectious-Recovered) epidemiological model (18; 21; 22) with multiple population groups that interact with each other (SI §2). We augment the model with controls that target based on (i) *age groups*, and (ii) types of *activities* that individuals engage in. Different policy interventions can be embedded as controls: we focus on time-dependent, targeted confinements, but extend the model to mass testing in the SI; vaccinations can also be accommodated. Interventions modulate the rate of social contacts and the economic value generated, and the objective of the control problem is to minimize a combination of health and economic losses caused by deaths, illness, and activity restrictions. The model captures important resource constraints (such as hospital and ICU), and allows explicitly controlling the amount of targeting through "limited disparity" constraints that limit the difference in the extent of confinement imposed on distinct age groups.

Epidemiological Model and Controls. We segment the population by age into nine groups $g \in \mathscr{G}$; the youngest group g captures individuals with age 0-9 and the oldest those aged 80 or above. For each g, the compartmental model includes states for susceptible, exposed, infectious, quarantined infectious, recovered, and deceased individuals. We also reserve separate states for individuals who are hospitalized due to being infected, in either general hospital wards or in intensive care units (ICU). We use T to denote the time horizon of the control problem, and $\boldsymbol{X}_t$ to denote the entire vector of epidemiological states at time $0 \leq t \leq T$.

Individuals interact in activities belonging to the set $\mathscr{A} = \{\mathsf{work}, \mathsf{transport}, \mathsf{leisure}, \mathsf{school}, \mathsf{home}, \mathsf{other}\}$. These interactions generate *social contacts* which drive the rate of new infections.

We control the SEIR dynamics by adjusting the confinement intensity in each group-activity pair over time: we let $\ell_g^a(t) \in [0,1]$ denote the activity level allowed for group g and activity a at time t, expressed as a fraction of the activity level under *normal course* of life (no confinement). We denote $\boldsymbol{\ell}_g(t) = [\ell_g^a(t)]_{a \in \mathscr{A}}$, $\boldsymbol{u}_t = [\boldsymbol{\ell}_g(t)]_{g \in \mathscr{G}}$ and $\boldsymbol{u}_{t:t'} = [\boldsymbol{u}_t, \ldots, \boldsymbol{u}_{t'}]$.

We propose a parametric model to map activity levels to social contacts. We use $c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h)$ to denote the mean number of total daily contacts between an individual in group g and individuals in group h across all activities when their activity levels are $\boldsymbol{\ell}_g, \boldsymbol{\ell}_h$, respectively. Varying the activity levels changes the social contacts according to

$$c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h) = \sum_{a \in \mathscr{A}} C_{g,h}^a \cdot (\ell_g^a)^{\alpha_1} \cdot (\ell_h^a)^{\alpha_2}, \tag{1}$$

where $C_{g,h}^a$ denote the mean number of daily contacts in activity a under normal course (i.e., without confinement), and $\alpha_1, \alpha_2 \in \mathbb{R}$ are parameters. We retrieve values for $C_{g,h}^a$ from the data tool (23), which is based on the French social contact survey data in (24), and we estimate α_1, α_2 from health outcome data (25) and Google mobility data (26).

When the number of patients requiring hospitalization or ICU treatment exceeds the respective capacity of available beds, we assume that patients are turned away from each age group according to a proportional rule.[1]

Objective. Our objective captures two criteria. The first quantifies the total deaths directly attributable to the pandemic, which we denote by $\mathsf{Total\ Deaths}(\boldsymbol{u}_{0:T-1})$ to reflect the dependency on the specific confinement policy $\boldsymbol{u}_{0:T-1}$ followed.

The second captures the economic losses due to the pandemic, denoted by $\mathsf{Economic\ Loss}(\boldsymbol{u}_{0:T-1})$ and stemming from restrictions of activity (SI §2). For example, lowering $\ell_{30-39\text{ y.o.}}^{\mathsf{work}}(t)$ causes lost wages in $g = 30-39$ y.o., while lowering $\ell_{10-19\text{ y.o.}}^{\mathsf{school}}(t)$ causes lost schooling costs for $g = 10-19$ y.o.

To allow policy makers to weight the importance of the two criteria, we associate a cost χ to each death, which we express in multiples of French GDP per capita. Our framework can capture a multitude of policy preferences by considering a wide range of χ values, from completely prioritizing economic losses ($\chi = 0$) to completely prioritizing deaths ($\chi \to \infty$).

[1] Our framework allows implementing any turn-away rule or even optimizing over these decisions.

Optimization Problem. We seek control policies $\boldsymbol{u}_{0:T-1}$ that minimize the sum of economic and mortality costs. More formally, we solve:

$$\min_{\boldsymbol{u}_{0:T-1}} \left[\text{Economic Loss}(\boldsymbol{u}_{0:T-1}) + \chi \cdot \text{Total Deaths}(\boldsymbol{u}_{0:T-1})\right] \tag{2}$$

subject to constraints that (i) the state trajectory follows the SEIR dynamics, and (ii) the controls and states do not exceed the available capacities of ICU beds and hospital beds.

Re-Optimization with Linearized Dynamics — ROLD. Minimizing (2) is challenging to solve to optimality even for small time horizons due to the non-linear and non-convex nature of the SEIR dynamics. To overcome this, we design a model predictive control (27) approximate algorithm that remains tractable for realistic time horizons and problem sizes.

The algorithm proceeds sequentially over the time horizon. When the SEIR state is $\boldsymbol{X}_t$ at time t, it builds an approximation of the SEIR dynamics from t to T that is *linear* in the controls $\boldsymbol{u}_{t:T-1}$. This linearity allows us to find the best control $\boldsymbol{u}^*_{t:T-1}$ with respect to these linearized dynamics. The algorithm uses $\boldsymbol{u}^*_t$ as the intervention in the *real* SEIR system to advance to state $\boldsymbol{X}_{t+1}$, and iterates by updating the linearization around the new state. Re-solving at each time t is aimed to prevent the linearization errors from growing too large (details in SI §3).

Parametrization and Model Calibration. We adopt the confidence regions for SEIR parameters reported in (22) for the Île-de-France region, which we complement with Google mobility data to approximate the mean effective lockdowns for activities during the horizon of interest. We estimate values of all parameters by comparing the simulated SEIR model against real data on hospital and ICU utilization and deaths. We calibrate our economic model using data from France (and where available Île-de-France) on full time equivalent wages and employment rates, and sentiment surveys on business activity levels during confinement.

Experimental Setup. We use a horizon of $T = 90$ in the experiments reported in the main paper, and allow up to $T = 360$ days in additional experiments reported in SI. We allow the confinement decisions to change every two weeks. We use a capacity of 2900 beds for ICU in Île-de-France, and infinite beds for general hospital wards. We consider several ROLD policies that differ in the level of targeting allowed, which we compare over a wide range of values for χ, from 0 to 1000$\times$ the annual GDP per capita in France[2]. In SI §7 we also discuss the additional benefits of targeted administration of viral tests across age groups.

We provide all the details for the ROLD heuristic, calibration, parameter specification, and experimental setup in the SI, where we also report results from sensitivity and robustness analyses on the fitted parameters.

Results

We use this framework to address our main research questions.

How large are the gains from dual targeting? To isolate the benefits of each type of targeting, we compare four versions of ROLD that differ in the level of targeting allowed: no targeting whatsoever ("NO-TARGET"), targeting age groups only ("AGE"), activities only ("ACT"), or both ("AGE-ACT", or simply "ROLD" when no confusion can arise). Figure 1a records each policy's performance in several problem instances parameterized by the cost of death χ. A striking feature is that each of the targeted policies actually *Pareto-dominates* the NO-TARGET policy, and the improvements are significant: relative to NO-TARGET and for same number of deaths, economic losses are reduced by EUR 0-2.9B (0%-35.9%) in AGE, by EUR 0.4B-2.1B (4.5%-49.8%) in ACT, and by EUR 3.3B-5.3B (35.7%-80.0%) in AGE-ACT. This Pareto-dominance is unexpected, since it is not a property that we explicitly ask for in our optimization procedure, and it underlines that any form of targeting can lead to significant improvements in terms of *both* health and economic outcomes.

When comparing the different types of targeting, neither AGE nor ACT Pareto-dominate each other, and neither policy dominates in terms of the total loss objective (SI Figure S3). In contrast and crucially, AGE-ACT Pareto-dominates *all* other policies, and moreover leads to super-additive improvements in almost all cases: for the same number of deaths, AGE-ACT reduces economic losses by more than AGE and ACT *added together* (SI Figure S4). This suggests that substantial complementarities may be unlocked through targeting *both* age groups and activities, which may not be available under less granular targeting.

To confirm the significance of these gains, we also compare ROLD AGE-ACT with various practical benchmark policies in Figure 1b. Benchmarks ICU-t and Hybrid-t AND/Hybrid-t OR mimic implementations in the U.S. Austin area (28) and, respectively, France (29). These policies switch between a stricter and a relaxed confinement level based on conditions related to hospital occupancy and the rate of new infections (SI §5). We also consider two extreme benchmarks corresponding to enforcing "full confinement" (FC) or remaining "fully open" (FO); these can be expected to perform well when completely prioritizing one of the two metrics of interest, with FC minimizing the number of deaths and FO ensuring low economic losses.

[2]We quantify the cost of death χ as a multiple of the annual GDP per capita in France, and use the shorthand notation $n\times$ to denote a value of n times this annual GDP per capita.

ROLD Pareto-dominates all these benchmarks, decreasing economic losses by EUR 5.3B-16.9B (71.0%-82.6%) relative to Hybrid-t AND, by EUR 7.1B-11.6B (62.2%-82.8%) relative to Hybrid-t OR, and by EUR 5.4B-11.6B (62.2%-78.0%) relative to ICU-t for the same number of deaths. Additionally, ROLD meets or exceeds the performance of the two extreme policies: for a sufficiently large χ, ROLD exactly recovers the FC policy, resulting in 890 deaths and economic losses of EUR 27.6B; for a sufficiently low χ, ROLD actually Pareto-dominates the FO policy, reducing the number of deaths by 16,688 (76.7%) and reducing economic losses by EUR 1.6B (65.3%). The latter result, which may seem surprising, is driven by the natural premise captured in our model that deaths and illness generate economic loss because of lost productivity; thus, a smart sequence of confinement decisions can actually improve the economic loss relative to FO. Among all the policies we tested, ROLD AGE-ACT was *the only one* capable of Pareto-dominating the FO benchmark, confirming that dual targeting is critical and powerful.

The Pareto-dominance of ROLD AGE-ACT implies its dominance in terms of the total loss objective (SI Figure S3). These results are robust under more problem instances (SI Figure S5).

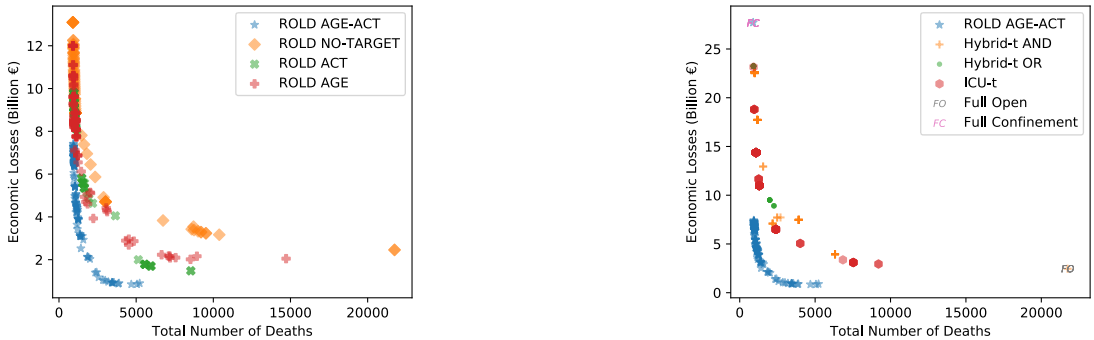

(a) Comparison between four ROLD policies with different levels of targeting

(b) Comparison between ROLD with dual targeting and benchmark policies

Figure 1. The total number of deaths and the economic losses generated by ROLD policies with different levels of targeting and by the benchmark policies. Panel **(a)** compares the four versions of ROLD that differ in the level of targeting allowed. Panel **(b)** compares the ROLD policy that targets age groups and activities with the benchmark policies. Each marker corresponds to a different problem instance parametrized by the cost of death χ. We include 128 distinct values of χ from 0 to 990×, and panel (b) also includes a very large value ($\chi = 10^{16}\times$).

How do gains arise from dual targeting? We examine the structure of the optimal ROLD AGE-ACT confinement decisions. We focus our discussion on the value $\chi = 50\times$, which is in the mid-range of estimates used in the economics literature on COVID-19 (30) and is representative of the behavior we observe across all experiments (SI §7). Figure 2a visualizes the optimized confinement policy.

Generally, the ROLD policy maintains high activity levels for those groups with a high ratio of marginal economic value to total social contacts in the activity, i.e., a high

$$\text{“econ-to-contacts-ratio”} := \frac{dv_g(\boldsymbol{u}_t)/d\ell_g^a(t)}{\sum_{h\in\mathcal{G}} C_{h,g}^a},$$

where $v_g(\boldsymbol{u}_t)$ is the economic value created by an individual in age group g when activity levels are $\boldsymbol{u}_t$. For example, in work, ROLD completely opens up the 40-69 y.o. groups, while confining the 20-39 y.o. groups during the first two weeks and the 10-19 y.o. groups for the first ten weeks. This is explainable since the 40-69 y.o. age groups produce the highest econ-to-contacts-ratio in work, while the younger groups have progressively lower ratios. Similarly, ROLD prioritizes activity in transport, then other, then leisure, in accordance with the relative econ-to-contacts-ratio of these activities.

To confirm the robustness of this insight, we also conduct a more thorough study where we compute optimal ROLD policies for several problem instances, and then train regression decision trees to predict the optimal ROLD activity levels as a function of several features (SI §7). The results are captured in Figure 2b and SI Figure S10. These simple trees can predict the optimal ROLD activity levels quite well (with root MSE values in the range 0.10-0.22), and they confirm our core insight that the econ-to-contacts ratio is the most salient feature when targeting confinements, as it is used as a split variable in the root node of each tree, with higher ratios leading to higher activity levels in all activities considered.

To understand how *complementarities* arise in this context, note that the ability to separately target age groups and activities allows the ROLD policy to fully exploit the fact that distinct age groups may be responsible for the largest econ-to-contacts

ratio in different activities. As an example, the 20-69 y.o. groups have the highest ratio in work, whereas the 0-19 y.o. and 70+ y.o. groups have the highest ratio in leisure. ROLD coordinates confinements to account for this: groups 20-69 y.o. remain more open in work but face confinement in leisure for the first ten weeks, whereas the 10-19 y.o. group is confined in work for a long period while remaining open in leisure, as remain the 70+ y.o. groups. These complementary confinement schedules allow ROLD to reduce both the number of deaths and economic losses, with the added benefit that no age group is completely confined.

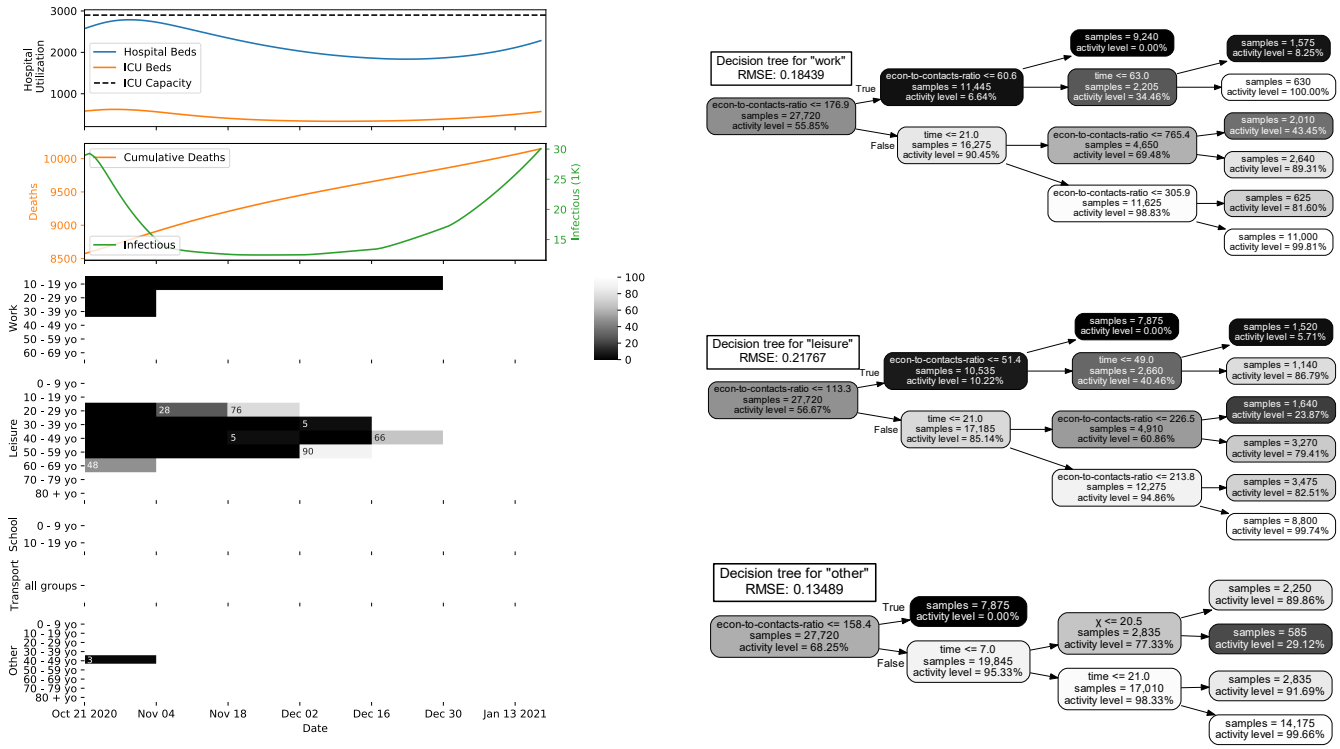

(a) Optimized ROLD confinement policy and its estimated impact

(b) Decision trees for work, leisure and other activities

Figure 2. The optimized ROLD AGE-ACT policies for problem instances with a 90-day optimization horizon starting on October 21, 2020 (see SI for optimized ROLD policies with a 360-day optimization horizon). Figure **(a)** corresponds to a problem instance where the cost of death χ is 50×. The seven panels depict the time evolution for the occupation of hospital and ICU beds (top panel), the number of actively infectious individuals and the cumulative number of deceased individuals in the population (panel 2), and the confinement policy imposed by ROLD in each age group and activity (panels 3-7). In panels 3-7, the values correspond to the activity levels allowed for the respective age group, and are color-coded so that darker shades capture a stricter confinement. Figure **(b)** depicts decision trees approximating the optimized ROLD confinement decisions for work, leisure and other (trained with 27,720 samples), with a horizon of $T = 90$ days. Each node in the tree records several pieces of information: a logical condition based on which all the training samples in the node are split, with the upper sub-tree corresponding to the logical condition being true (e.g., "econ-to-contacts-ratio $\leq$ 176.9" for the root node in the work tree), the number of training samples falling in the node ("samples"), and the average activity level for all the samples in the node. The nodes are color-coded based on the activity level, with darker colors corresponding to stricter confinement.

Can dual targeting reduce time in confinement for each age group? We calculate the fraction of time spent by each age group in confinement under each ROLD policy, averaged over the activities relevant to that age group (SI §7). The results are visualized in Figure 3, which depicts boxplots for the fractions of time in confinement across all problem instances parameterized by χ.

We find that the dual-targeted AGE-ACT policy is able to reduce the confinement time quite systematically for *every* age group, relative to *all* other policies. Specifically, it results in the lowest confinement time for every age group in 70% of all problem instances when compared with NO-TARGET, in 60% of instances when compared with AGE, in 83% of instances when compared with ACT, and in 50% of instances when compared with *all* other policies. Moreover, the fraction of confinement time achieved by AGE-ACT is within 5% (in absolute terms) from the lowest confinement time achieved by any policy for every age group, in 76% of all instances; within 10% in 80% of the instances; and within 14% in all instances.

Thus, even when the dual-targeted policy confines certain age groups more, it does not do so by much. These outcomes are quite unexpected as they are not something that the ROLD framework explicitly optimizes for, but rather a *by-product* of a dual-targeted confinement policy that minimizes the total loss objective (2).

It is worth noting that although ROLD AGE-ACT reduces confinements for every age group compared to less targeted policies, it does not do so uniformly. Instead, it can lead to a larger discrepancy in the confinements faced by different age groups: those aged 10-59 are generally more confined than those aged 0-9 or 60+.

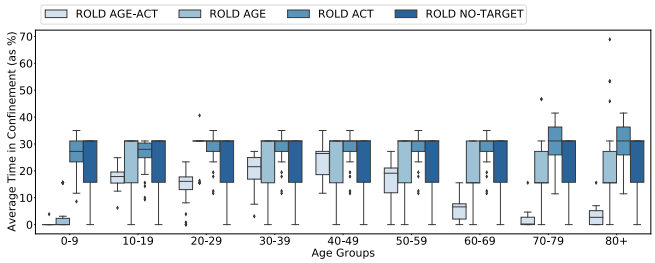

Figure 3. Average time in confinement for the ROLD policies with different targeting types. Each boxplot depicts the fraction of time the age group spends in confinement under the respective policy averaged over the activities relevant to that age-group, for different problem instances parameterized by the cost of death χ.

The impact of limited disparity requirements. That targeted policies confine some age groups more than others could be perceived as disparate or unfair treatment, so it is important to quantify how an intervention's effectiveness is impacted when requiring less differentiation across age groups. We embed a set of "limited disparity" constraints in ROLD that allow the activity levels of distinct age groups to differ by at most Δ in absolute terms, in each activity and at any point of time (SI §7). $\Delta = 0$ corresponds to a strictly non-discriminatory policy, whereas a larger value of Δ allows more targeting, with $\Delta = 1$ corresponding to a fully targeted policy. For every value of Δ, we record the total loss incurred by a ROLD policy with the limited disparity constraints, and calculate the increase in total loss relative to a fully targeted ROLD policy. We repeat the experiment for different problem instances parametrized by χ, and Figure 4 depicts boxplots of all the relative increases in total loss, as a function of Δ.

The results suggest that limited disparity requirements may be costly: on average, completely eliminating disparity in confinements would increase the total losses by EUR 1.2B (21.6%) and produce an additional 506 deaths (16.6%) and an extra EUR 0.5B of economic losses (18.9%) compared to a fully targeted policy. In certain instances, the increase in total loss could be as high as 63%. The high losses persist even when some limited discrepancy is allowed, dropping at an initially slow rate as Δ increases from 0 and eventually at a slightly faster rate as it approaches 1. This suggests that to fully leverage the benefits of targeting, a high level of disparity must be accepted, but reasonable trade-offs can be achieved with some intermediate disparity.

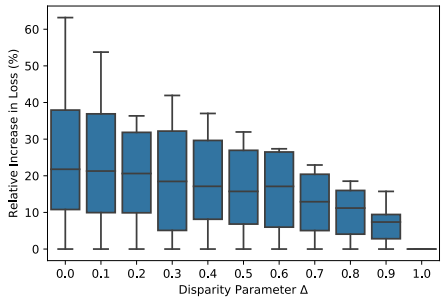

Figure 4. The impact of limited disparity requirements. The plot shows the relative (%) increase in total loss generated by a ROLD policy, compared to a fully targeted policy, as a function of the disparity parameter Δ that measures the maximum allowed difference in activity levels for distinct age groups. The experiments are run using several values of χ, which are used to generate each of the boxplots. Eleven values of Δ are tested, ranging from 0 to 1.

Discussion

Our case study suggests that an optimized intervention targeting both age groups *and* activities carries significant promise for alleviating a pandemic's health, economic and even psychological burden, but also points to certain challenges that require care in a real-world setting.

Why consider optimized dual-targeted interventions? The first reason are the significantly better health *and* economic outcomes: for the same or a lower number of deaths, dual- targeted confinements can reduce economic losses more than any of the simpler interventions that uniformly confine age groups or activities. Furthermore, the super-additive gains imply that significant synergies can be generated through finer targeting, with the ability to target along activities improving the effectiveness of targeting along age groups, and vice-versa.

The second reason is the intuitive nature of the optimized targeted confinement policy, which is consistent with a simple "bang-for-the-buck" rule: impose less confinement on group-activity pairs that generate a relatively high economic value prorated by (activity-specific) social contacts. This simple intuition combined with the reliance on just a few activity levels are appealing practical features, as they provide transparency into how targeted confinement decisions could be made.

The simple "bang-for-the-buck" intuition is also related to the third benefit of dual targeting: the ability to impose less restrictive confinements across all age groups. As different age groups may be responsible for generating a larger economic value prorated by social contacts in distinct activities, dual-targeted confinements may enable all age groups to remain more active, resembling normal life more closely compared to less differentiated confinements. This could result in more socially acceptable restrictions, and a more appealing policy intervention overall.

Lastly, we note that although dual targeting allows and can result in differences in confinements across age groups, such interventions are actually not far from many real-world policy implementations, which have been more or less explicit in their age-based discrimination. Dual targeting can arise *implicitly* in interventions that only seem to target activities. As an example, France implemented a population-wide 6 p.m. to 6 a.m. curfew (31), while maintaining school and work activities largely de-confined. This is effectively implementing restrictions similar to ROLD AGE-ACT: since a typical member of the 20-65 y.o. group is engaged in work until the start of the curfew, their leisure and other activities are implicitly limited; moreover, since most individuals aged above 65 are not in active employment, they are not that restricted in these last two activities.

These examples show that some amount of targeting of activities *and* age groups is already in place and is perhaps unavoidable for effective pandemic management. Given this state of affairs, our framework highlights the significant benefits in explicitly and transparently modelling targeting and identifying the interventions that rigorously optimize overall societal welfare, given some allowable amount of differentiation across age groups.

Challenges and limitations. An immediate practical challenge is data availability. Social contact matrices by age group and activity may be available from surveys on social behavior, which have been conducted in a number of countries; however, further data collection might be required to obtain these matrices for more refined population group or activity definitions. Similarly, economic data is reported by industry activities, but we are not aware of a dataset that splits economic value into separate (group, activity) contributions. Disparate data sources may be difficult to align: for example, social contact surveys and Google mobility reports use different activity categories, which requires non-trivial fitting (SI §4).

Availability of data also constrains our model's structure in several ways. In our model social contacts between age groups only depend on confinements in the same activity, since the available social contacts dataset (24) only reports interactions in the same activity. However, contacts occur as individuals are engaged in different activities (e.g., a professional in the services industry interacts during work with individuals who are engaged in leisure activities). A more refined contact mixing model that captures such interactions would be more appropriate for this study, provided that relevant data are available. Our economic model similarly ignores cross effects, such as young age groups engaged in school producing value in conjunction with educational staff engaged in work.

Another challenge with targeted interventions is the perception that they lead to unfair outcomes, as certain age groups face more confinement than others. Such discrimination does arise in the optimized dual-targeted policies; our framework can partially address these concerns through explicit constraints that limit the disparities across groups. Our requirement that limited disparity hold for every time period and every activity is quite strict, and a looser requirement based, e.g., on time-average confinements could lead to smaller incremental losses. Alternatively, one can impose fairness requirements based on the intervention's *outcomes*, e.g., requiring that the health or economic losses faced by different groups satisfy certain axiomatic fairness properties (32).

Although we focus on confinement policies, a direction for future research is to investigate how these can be optimally combined with other types of targeted interventions. SI §7 reports experiments where we optimize a targeted policy based on confinements and randomized testing and quarantining. The framework is sufficiently flexible to accommodate interventions such as contact tracing and also vaccinations, although a careful implementation would require work beyond the scope of this article.

Acknowledgments

The authors thank Ciprian Bădescu, Florin Georgescu, Alexandru Hulea, Andrei Leica, and Vlad Păunescu from Arnia Software for their efforts in implementing the framework and hosting it at http://insead.arnia.ro.

Authors' contributions

DFC, DAI and SIZ contributed equally to the research design, conceptualisation, modeling, analysis, and writing. SC and XSW contributed equally to the modeling, software implementation, data analysis and creating visualizations. DFC, DAI and SIZ verified all the data used in this study.

Declaration of interests

All authors declare no competing interests.

References

1. Acemoglu, D., Chernozhukov, V., Werning, I. & Whinston, M. D. Optimal targeted lockdowns in a multi-group SIR model. Working Paper 27102, NBER (2020). DOI: 10.3386/w27102.

2. Matrajt, L., Eaton, J., Leung, T. & Brown, E. R. Vaccine optimization for COVID-19, who to vaccinate first? (2020).

3. Goldstein, J. R., Cassidy, T. & Wachter, K. W. Vaccinating the oldest against COVID-19 saves both the most lives and most years of life. *Proc. Natl. Acad. Sci.* **118**, DOI: 10.1073/pnas.2026322118 (2021). https://www.pnas.org/content/118/11/e2026322118.full.pdf.

4. Bertsimas, D. *et al.* Optimizing vaccine allocation to combat the COVID-19 pandemic (2020).

5. Favero et al., C. A. Restarting the economy while saving lives under COVID-19 (2020).

6. Birge, J. R., Candogan, O. & Feng, Y. Reducing economic losses with targeted closures. Working Paper 2020-57, University of Chicago, Becker Friedman Institute for Economics (2020).

7. Chang, S. *et al.* Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* **589**, 1–6 (2020).

8. Evgeniou, T. *et al.* Epidemic models for personalised COVID-19 isolation and exit policies using clinical risk predictions, DOI: 10.1101/2020.04.29.20084707 (2020). https://www.medrxiv.org/content/early/2020/05/12/2020.04.29.20084707.full.pdf.

9. Giordano, G. *et al.* Modeling vaccination rollouts, SARS-CoV-2 variants and the requirement for non-pharmaceutical interventions in Italy. *Nat. Medicine* **27**, 993–998, DOI: 10.1038/s41591-021-01334-5 (2021).

10. Bastani, H. *et al.* Deploying an artificial intelligence system for COVID-19 testing at the Greek border. *Available at SSRN* (2021).

11. Tiirinki, H. *et al.* COVID-19 pandemic in Finland–preliminary analysis on health system response and economic consequences. *Heal. policy technology* **9**, 649–662 (2020).

12. Harrison, S. Coronavirus: Ireland's restrictions eased for over 70s (2020). Access. 01/12/2021.

13. Magid, J. Cabinet approves removal of age restriction for those returning to work (2020). Accessed January 12, 2021.

14. Foy, H. Russian businesses prepare for fresh lockdowns as Covid-19 cases soar (2020). Accessed January 12, 2021.

15. Reuters Staff. Bosnian region eases lockdown on seniors, children after court ruling (2020). Accessed January 12, 2021.

16. Kanbur, N. & Ankgül, S. Quaranteenagers: A single country pandemic curfew targeting adolescents in Turkey. *J. Adolesc. Heal.* **67** (2020).

17. Kucharski, A. J. *et al.* Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of sars-cov-2 in different settings: a mathematical modelling study. *The Lancet Infect. Dis.* **20**, 1151–1160, DOI: 10.1016/S1473-3099(20)30457-6 (2020).

18. Prem, K. *et al.* The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study. *The Lancet Public Heal.* **5**, e261–e270, DOI: 10.1016/S2468-2667(20)30073-6 (2020).

19. Di Domenico, L., Pullano, G., Sabbatini, C. E., Boëlle, P.-Y. & Colizza, V. Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC medicine* **18**, 1–13 (2020).

20. Aguilera, J. Some supermarkets are launching senior-only hours during the coronavirus pandemic. Not all retailers think that's a good idea (2020). Accessed January 12, 2021.

21. Anderson, R. M. & May, R. M. *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press, 1992).

22. Salje, H. *et al.* Estimating the burden of SARS-CoV-2 in France. *Science* **38**, DOI: 10.1126/science.abc3517 (2020).

23. Wille, L. *et al.* SOCRATES: an online tool leveraging a social contact data sharing initiative to assess mitigation strategies for COVID-19. *BMC Res. Notes* **13** (2020).

24. Béraud, G. *et al.* The French connection: The first large population-based contact survey in France relevant for the spread of infectious diseases. *PLOS ONE* **10**, 1–22, DOI: 10.1371/journal.pone.0133203 (2015).

25. French Government. Données hospitalières relatives à l'épidémie de COVID-19. https://www.data.gouv.fr/fr/datasets/donnees-hospitalieres-relatives-a-lepidemie-de-covid-19/ (2020). Accessed October 21, 2020.

26. Google. COVID-19 community mobility report. https://www.google.com/covid19/mobility/ (2020). Accessed October 21, 2020.

27. Borrelli, F., Bemporad, A. & Morari, M. *Predictive control for linear and hybrid systems* (Cambridge University Press, 2017).

28. Duque, D. *et al.* Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proc. Natl. Acad. Sci.* **117**, 19873–19878 (2020).

29. Lehot, M. & Borgne, B. L. Covid-19 : ces chiffres qui montrent que Paris a dépassé le seuil d'alerte maximale depuis le 25 septembre (2020). Accessed October 5, 2020.

30. Alvarez, F. E., Argente, D. & Lippi, F. A simple planning problem for COVID-19 lockdown. Working Paper 26981, National Bureau of Economic Research (2020). DOI: 10.3386/w26981.

31. Reuters. France goes under nationwide 6pm curfew as Covid-19 death toll surpasses 70,000 (2021). Accessed January 12, 2021.

32. Young, H. P. *Equity* (Princeton University Press, 1994).

Supplementary Information for

Quantifying The Benefits of Targeting for Pandemic Response

Sergio Camelo, Dragos F. Ciocan, Dan A. Iancu, Xavier S. Warnes, and Spyros I. Zoumpoulis

Emails of authors: camelo@stanford.edu, florin.ciocan@insead.edu, daniancu@stanford.edu, xwarnes@stanford.edu, spyros.zoumpoulis@insead.edu

This PDF file includes:

Supplementary text
Figs. S1 to S10
Tables S1 to S8
SI References

Supporting Information Text

1. Some Notation

We denote scalars by lower-case letters, as in v, and vectors by bold letters, as in $\boldsymbol{v}$. We use square brackets to denote the concatenation into vectors: $\boldsymbol{v} := [v_0, v_1]$. For a time series of vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$, we use the notation $\boldsymbol{v}_{i:j} := [\boldsymbol{v}_i, \ldots, \boldsymbol{v}_j]$ to denote the concatenation of vectors $\boldsymbol{v}_i$ through $\boldsymbol{v}_j$. Lastly, we use $\boldsymbol{v}^\top$ to refer to the transpose of $\boldsymbol{v}$.

2. Model and Optimization Problem

In this section, we propose a modified version of the discretized SEIR (Susceptible - Exposed - Infectious - Recovered) population-based epidemiological model with multiple population groups, which interact with each other. In our model, the SEIR dynamics are controllable via (i) choosing confinement patterns for different groups and (ii) choosing how to apportion testing among the groups.

Then, we build a economic model on top of the SEIR system, which factors in the cost of lost output, together with the cost of deaths. We conclude by formulating the overall optimization problem.

A. Epidemiological Model.

Time. Time is discrete, indexed by $t = 0, 1, \ldots, T$, and measured in days. We assume that no infections are possible beyond day T.

Age Groups. We split the population into age groups. We use $\mathcal{G}$ to denote the collection of considered age groups. In our case study for Île-de-France we use nine age groups, split in 10-year buckets, where the youngest group are the 0-9 year olds (y.o.), and the oldest group are the 80+ year olds. Since individuals in separate groups have social contacts with each other, the dynamics of each group depend on the number of infections in other groups.

Compartmental Model and States. At any given time, the population of an age group is divided into susceptible, exposed, infectious, recovered, and deceased as in a classical SEIR model. Figure S1 represents the compartmental model and the SEIR transitions for a specific group g.

All the states represent the number of individuals in a compartment of the model in the *beginning* of the time period. State $S_g(t)$ is the number of individuals in group g that are susceptible to be infected at time t. State $E_g(t)$ is the number of individuals in group g that have been exposed to the SARS-CoV-2 virus at time t, but are not yet infectious. $I_g(t)$ is the number of individuals in group g that are infectious at time t, and have not yet been tested, or recovered, or transferred to the hospital or the ICU.

Susceptible individuals get infected and transition to the exposed state at a rate determined by the number of social contacts as well as transmission rate $\beta(t)$. Exposed individuals transition to the infectious state at a rate σ.

$I_g^{\mathrm{q}}(t)$ is the number of individuals in group g who have been confirmed to be infectious through viral testing at time t, and are thus quarantined. We subdivide $I_g^{\mathrm{q}}(t)$ into $I_{j,g}^{\mathrm{q}}(t)$ for $j \in \{a, ps, ms, ss\}$ to model different degrees of severity of symptoms: infectious individuals can be asymptomatic, paucisymptomatic, have mild symptoms, or have severe symptoms. We assume that an infectious individual in group g will exhibit symptoms of degree j with probability $p_{j,g}$. Infectious individuals transition out of the infectious state at a rate μ.

Individuals with severe symptoms need hospitalization. $H_g(t)$ is the number of individuals in group g that are in the general hospital wards at time t, and $ICU_g(t)$ is the number of individuals in group g that are in the intensive care unit at time t. We assume that an infectious individual will need to be hospitalized in the general wards (go to ICU) with probability p_g^H (p_g^{ICU}). We have $p_{ss,g} = p_g^H + p_g^{ICU}$. On average, patients who are treated in the general hospital wards spend λ_H^{-1} days in the hospital; patients who are treated in the ICU spend λ_{ICU}^{-1} days in the ICU.

Infectious individuals with severe symptoms may decease. $D_g(t)$ is the number of individuals in group g that have died from COVID-19 up until time t. We assume that an infectious individual with severe symptoms in group g will decease with probability p_g^D (and recover with probability $p_g^R = 1 - p_g^D$).

Infectious individuals usually recover. $R_g(t)$ is the number of individuals in group g who have recovered until time t, but have not been confirmed to have recovered. $R_g^{\mathrm{q}}(t)$ is the number of individuals in group g who have recovered until time t, and have been confirmed to have recovered and to have had the virus either through testing, or because they recovered during their stay in the hospital or the ICU.

We keep track of the total number of individuals in group g who are not confirmed infectious, are not confirmed recovered, are not in the hospital or the ICU, and have not died: $N_g(t) := S_g(t) + E_g(t) + I_g(t) + R_g(t)$. We let

$$\boldsymbol{X}(t) = \left[S_g(t), I_g(t), ..., D_g(t)\right]_{g \in \mathcal{G}}$$

denote the full state of the system at time $t = 0, 1, \ldots, T-1, T$, i.e., the number of individuals in each of the states, across groups. We denote the number of compartments by $|\mathcal{X}|$. Then the dimension of $\boldsymbol{X}(t)$ is $|\mathcal{G}||\mathcal{X}| \times 1$.

We assume that $D_g(0) = 0$ for all $g \in \mathcal{G}$.

Fig. S1. Compartmental SEIR model for a specific group g

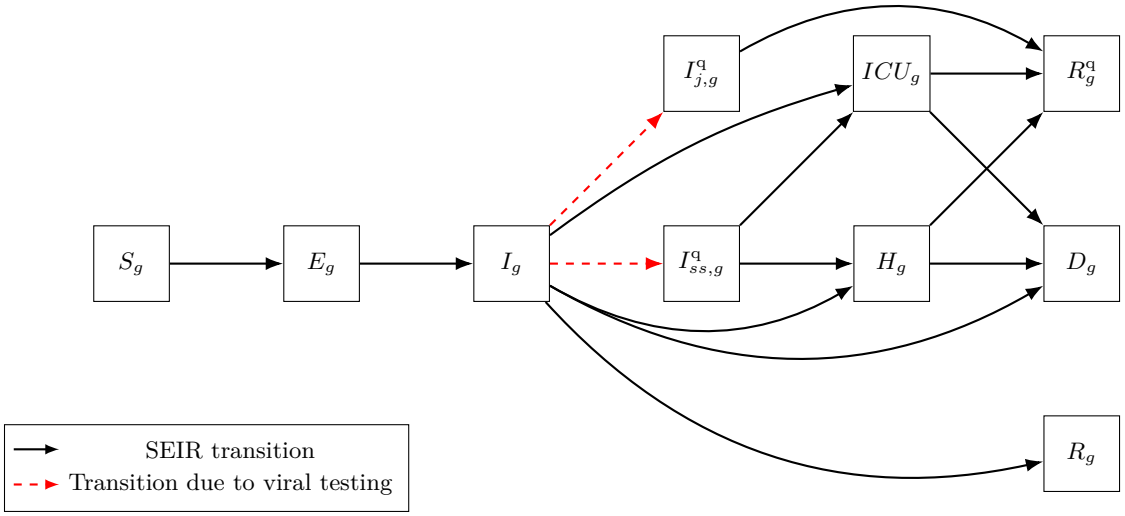

A.1. Controls/Decision Variables. We allow for three levers for the policy maker's response: targeted testing, targeted confinements, and targeted patient turn-away decisions.

Confinements and social mixing model. We refer to the activity level decision for a group g as the vector

$$\boldsymbol{\ell}_g(t) := \left[\ell_g^{\text{home}}(t), \ell_g^{\text{work}}(t), \ell_g^{\text{school}}(t), \ell_g^{\text{transport}}(t), \ell_g^{\text{leisure}}(t), \ell_g^{\text{other}}(t)\right] \in [0,1]^6$$

that captures the level of each activity that an individual in group g is allowed to perform in period t, in relation to the normal activity level under normal course (no confinement), as a result of imposed confinement restrictions. We denote the set of considered activities by

$$\mathcal{A} := \{\text{home}, \text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}.$$

Normal course, i.e., no confinement, is modeled using $\boldsymbol{\ell}_g = \mathbf{1} = [1,1,1,1,1,1]$. As an example, $\ell_g^{\text{work}} = 0.7$ can be interpreted as a restriction for individuals in group g to spend at the workplace only 70% of the baseline time they spend at the workplace. We provide more details about how the baseline is defined for our numerical experiments and calibration in Section 4 B. When unambiguous, we use the terms activity levels and confinements to refer to the same underlying decision of $\boldsymbol{\ell}_g(t)$. The confinement decision variables for the policy maker are the activity levels of each group, in each time period: $\{\boldsymbol{\ell}_g(t)\}_{g\in\mathcal{G}, t=0:T-1}$. In our proposed implementation for Île-de-France we allow the confinement decisions to update every two weeks.

Confinements affect the amount of social contacts between individuals of different groups. We build a model to reflect this dependence which we describe below.

For activity $a \in \mathcal{A}$ and groups g, h, we denote $C_{g,h}^a$ the number of contacts in a day between an individual in group g and individuals in group h in activity a under normal course (no confinement). We retrieve the contact matrices per activity $\boldsymbol{C}^a, a \in \mathcal{A}$ from the social contact data tool by (1), using the survey on social contacts in France by (2).

Let $c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h)$ denote the mean number of contacts per day between an individual in group g and individuals in group h, when groups g, h are confined according to confinement patterns $\boldsymbol{\ell}_g, \boldsymbol{\ell}_h$, respectively. We use the following multiplicative extrapolation of social mixing matrices, with parameter $\alpha_1, \alpha_2 \in \mathbb{R}$:

$$c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h) = \sum_{a\in\mathcal{A}} C_{g,h}^a \cdot \left(\ell_g^a\right)^{\alpha_1} \cdot \left(\ell_h^a\right)^{\alpha_2}. \qquad [1]$$

We remark that this parametrization is similar to a Cobb-Douglas production function (3), using the confinement patterns as inputs, and the number of social contacts as output. In Section 4 B, we describe our procedure for fitting the social mixing parameters from data on health outcomes and mobility data.

Testing. We model viral diagnostic tests. Individuals who are found to be infected through a diagnostic test are quarantined. We assume that there is access to perfect tests (i.e., the sensitivity and specificity are 100%). We assume that a viral test yields a positive result when conducted on an infectious individual (compartment I_g), and a negative result when conducted on an individual who is in a non-infectious state.

We use $N_g^{\text{Vtest}}(t)$ to denote the number of viral tests allocated to group g in period t. The testing decisions for the policy maker are then $\{N_g^{\text{Vtest}}(t)\}_{g\in\mathcal{G}, t=0:T-1}$.We assume that, given a testing allocation decision across groups, the individuals to receive the allocated tests in group g are randomly chosen among individuals in $N_g(t)$. For example, the outflow from performing $N_g^{\text{Vtest}}(t)$ viral tests in group g, which happens out of compartment $I_g(t)$, is equal to

$$\frac{I_g(t)}{N_g(t)} \cdot N_g^{\text{Vtest}}(t). \qquad [2]$$

In our proposed implementation for Île-de-France we allow the testing decisions to change weekly. Figure S1 represents the compartmental model and the flows resulting from testing from one compartment to the other for a specific group g.

Turning patients away. When the patient inflow into the hospital or the ICU exceeds the remaining number of available beds, then the decision maker needs to decide how to prioritize the admission of patients to the available beds. Our general formulation allows the decision maker to optimize the number of patients from each group turned away from the general hospital wards and from the ICU, at each time period. We denote the number of patients from age group g who need to be admitted to the general hospital wards but are turned away in period t by $B_g^H(t)$. We denote the number of patients from age group g who need to be admitted to the ICU but are turned away in period t by $B_g^{ICU}(t)$. The turn-away decision variables for the policy maker are then $\{B_g^H(t), B_g^{ICU}(t)\}_{g\in\mathcal{G}, t=0:T-1}$. We assume that all patients who are turned away decease immediately. In our experiments we use a proportional turn-away rule, which is defined in Section A.3.

Let

$$\boldsymbol{u}(t) = \left[\boldsymbol{\ell}_g(t), N_g^{\text{Vtest}}(t), B_g^H(t), B_g^{ICU}(t)\right]_{g\in\mathcal{G}}$$

denote the vector of all the decisions/controls at time $t = 0, 1, \ldots, T-1$, i.e., the confinement and viral test decisions for all the groups, as well as the decisions on turning patients away. We denote the number of different decisions for a given group at a given time by $|\mathcal{U}|$. Then the dimension of $\boldsymbol{u}(t)$ is $|\mathcal{G}||\mathcal{U}| \times 1$.

A.2. Dynamics of the Controlled SEIR Model. Having defined the states and controls, we write down a set of discrete time dynamics for the controlled SEIR model. We use notation $\Delta X(t)$ to denote $X(t+1) - X(t)$. For all groups $g \in \mathcal{G}$, we write

$$\Delta N_g(t) = -N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)} - \underbrace{\mu \cdot (p_g^H + p_g^{ICU}) \cdot I_g(t)}_{\text{infected going to ICU, H}} \quad [3]$$

$$\Delta S_g(t) = -\beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \frac{I_h(t)}{N_h(t) + R_h^{\text{q}}(t)} \right) \quad [4]$$

$$\Delta E_g(t) = \beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \frac{I_h(t)}{N_h(t) + R_h^{\text{q}}(t)} \right) - \sigma E_g(t) \quad [5]$$

$$\Delta I_g(t) = \sigma E_g(t) - \underbrace{\mu \cdot (p_g^H + p_g^{ICU}) \cdot I_g(t)}_{\text{go to ICU, H}} - \underbrace{\mu \cdot (1 - p_g^H - p_g^{ICU}) \cdot I_g(t)}_{\text{recover}} - \underbrace{N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)}}_{\text{outflow to } I_{\cdot,g}^{\text{q}} \text{ from positive V tests}} \quad [6]$$

$$\Delta R_g(t) = \mu \cdot (1 - p_g^H - p_g^{ICU}) \cdot I_g(t) \quad [7]$$

$$\Delta I_{j,g}^{\text{q}}(t) = \underbrace{p_{j,g} \cdot N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)}}_{\text{inflow from } I_g \text{ from positive V tests}} - \underbrace{\mu \cdot I_{j,g}^{\text{q}}(t)}_{\text{recover}}, \qquad \forall j \in \{a, ps, ms\} \quad [8]$$

$$\Delta I_{ss,g}^{\text{q}}(t) = \underbrace{p_{ss,g} \cdot N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)}}_{\text{inflow from } I_{ss,g} \text{ from positive V tests}} - \underbrace{\mu \cdot I_{ss,g}^{\text{q}}(t)}_{\text{go to ICU, H, R, or D}} \quad [9]$$

$$\Delta R_g^{\text{q}}(t) = \underbrace{\mu \cdot \sum_{j \in \{a,ps,ms\}} I_{j,g}^{\text{q}}(t) + \lambda_g^H \cdot p_g^R \cdot H_g(t) + \lambda_g^{ICU} \cdot p_g^R \cdot ICU_g(t)}_{\text{recovery of known infected cases}} \quad [10]$$

$$\Delta H_g(t) = -\lambda_g^H H_g(t) + \mu \cdot \left(p_g^H \cdot I_g(t) + \frac{p_g^H}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\text{q}}(t) \right) - B_g^H(t) \quad [11]$$

$$\Delta ICU_g(t) = -\lambda_g^{ICU} ICU_g(t) + \mu \cdot \left(p_g^{ICU} \cdot I_g(t) + \frac{p_g^{ICU}}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\text{q}}(t) \right) - B_g^{ICU}(t) \quad [12]$$

$$\Delta D_g(t) = \lambda_g^H \cdot p_g^D \cdot H_g(t) + \lambda_g^{ICU} \cdot p_g^D \cdot ICU_g(t) + B_g^H(t) + B_g^{ICU}(t). \quad [13]$$

In Eq. (7) and Eq. (10), note that we do not have terms for the population turned away from hospital/ICU which may eventually recover. Instead, we assume the turned away patients will go into the deceased state. In Eq. (13), we are assuming that if a patient is turned away from the ICU, they transition into deceased, instead of being allocated a hospital bed if one is available.

We can summarize the SEIR dynamics with

$$\boldsymbol{X}(t+1) = \boldsymbol{X}(t) + F_t\left(\boldsymbol{X}(t), \boldsymbol{u}(t)\right), \quad [14]$$

where the function F_t captures the dynamics in Eq. (3)-Eq. (13).

We now provide justification for how we account for social contacts and, in particular, for the expressions in Eq. (4) and Eq. (5). We assume that individuals in R_g^{q} can interact with members of S_g, E_g, I_g and R_g but cannot infect.

Fix a person i in age group $g \in \mathcal{G}$, in state S_g. Then we can write

$$Pr\left(\bigcup_{h\in\mathcal{G}}\{i \text{ got infected through socializing with age group } h\}\right) \tag{15}$$

$$= Pr\left(\bigcup_{h\in\mathcal{G}}\{i \text{ got infected through socializing with individuals in } S_h, E_h, I_h \text{ or } R_h\}\right) \tag{16}$$

$$= 1 - Pr\left(\{i \text{ did not get infected through socializing with individuals in any } S_h, E_h, I_h \text{ or } R_h \text{ for any } h \in \mathcal{G}\}\right) \tag{17}$$

$$= 1 - \prod_{h\in\mathcal{G}} Pr\left(\{i \text{ did not get infected through socializing with individuals in } S_h, E_h, I_h \text{ or } R_h\}\right) \tag{18}$$

$$= 1 - \prod_{h\in\mathcal{G}}\left(1 - \beta(t)\frac{I_h(t)}{N_h(t)}\right)^{\frac{N_h(t)}{N_h(t)+R_h^{\mathrm{q}}(t)}c_{g,h}} \tag{19}$$

$$\approx 1 - \prod_{h\in\mathcal{G}}\left(1 - \beta(t)\cdot\frac{N_h(t)}{N_h(t)+R_h^{\mathrm{q}}(t)}\cdot c_{g,h}\cdot\frac{I_h(t)}{N_h(t)}\right) \tag{20}$$

$$\approx \beta(t)\sum_{h\in\mathcal{G}} c_{g,h}\frac{I_h(t)}{N_h(t)+R_h^{\mathrm{q}}(t)}, \tag{21}$$

where in Eq. (19) we use the following reasoning. Having fixed person i in age group g, the number of her contacts in age group h is given by $c_{g,h} = c_{g,h}(\ell_g(t), \ell_h(t))$. The number of contacts $c_{g,h}$ reflects the number of contacts of individual i from group h, assuming all individuals in age group h are put under the same confinement pattern. However, recall that individuals in $R_h^{\mathrm{q}}(t)$ cannot infect individuals in S_h, E_h, I_h or R_h. Therefore, we scale $c_{g,h}$ by the ratio of $(S, E, I, R, R^{\mathrm{q}})$ individuals in age group h that are in states (S, E, I, R), i.e., $\frac{N_h(t)}{N_h(t)+R_h^{\mathrm{q}}(t)}$, to account for the role of the relative size of groups of individuals in different confinement patterns within an age group. Finally, getting infected as the result of any contact within group h is considered to be an independent event.

By taking the expectation of random variable

$$\sum_{i\in S_g}\mathbb{1}\{i \text{ got infected through socializing}\},$$

we retrieve the expressions in Eq. (4) and Eq. (5).

A.3. Resources and Constraints. We use $K^{\mathrm{H}}(t)$ ($K^{\mathrm{ICU}}(t)$) to denote the capacity of beds in the general hospital wards (in the ICU) on day t. When the patient inflow into the hospital or the ICU exceeds the remaining number of available beds, then the policy maker needs to decide how many patients to turn away from each group. Although our framework allows to optimize over the decisions of turning away patients, we mainly look at a specific turn-away rule: let in patients from each age group proportionally to the arrivals from that age group, up to capacity, and turn away the remaining patients. This specific turn-away rule can be expressed by setting

$$B_g^H(t) = \underbrace{\frac{\mu\left(p_g^H\cdot I_g(t) + \frac{p_g^H}{p_g^H+p_g^{ICU}}\cdot I_{ss,g}^{\mathrm{q}}(t)\right)}{\sum_h \mu\left(p_h^H\cdot I_h(t) + \frac{p_h^H}{p_h^H+p_h^{ICU}}\cdot I_{ss,h}^{\mathrm{q}}(t)\right)}}_{\text{proportion of inflow into H from group } g}$$
$$\cdot\left(\underbrace{\sum_h \mu\left(p_h^H\cdot I_h(t) + \frac{p_h^H}{p_h^H+p_h^{ICU}}\cdot I_{ss,h}^{\mathrm{q}}(t)\right)}_{\text{total inflow into H}} - \underbrace{\left(K^{\mathrm{H}}(t) - \sum_h\left(1-\lambda_g^H\right)H_h(t)\right)}_{\text{available beds in H}}\right)^+ \tag{22}$$

$$B_g^{ICU}(t) = \frac{\mu\left(p_g^{ICU}\cdot I_g(t) + \frac{p_g^{ICU}}{p_g^H+p_g^{ICU}}\cdot I_{ss,g}^{\mathrm{q}}(t)\right)}{\sum_h \mu\left(p_h^{ICU}\cdot I_h(t) + \frac{p_h^{ICU}}{p_h^H+p_h^{ICU}}\cdot I_{ss,h}^{\mathrm{q}}(t)\right)}$$
$$\cdot\left(\sum_h \mu\left(p_h^{ICU}\cdot I_h(t) + \frac{p_h^{ICU}}{p_h^H+p_h^{ICU}}\cdot I_{ss,h}^{\mathrm{q}}(t)\right) - \left(K^{\mathrm{ICU}}(t) - \sum_h\left(1-\lambda_g^{ICU}\right)ICU_h(t)\right)\right)^+ \tag{23}$$

We further assume a given capacity for viral tests each day, which we denote by $K^{\mathrm{Vtest}}(t)$ on day t. We assume that viral tests used to test individuals with severe symptoms that enter the ICU or hospital, as well as viral tests that test hospitalized

individuals to confirm they have recovered, come from a different pool of tests and do not consume the capacity for viral testing in the non-hospitalized population.

We can now write the following constraints for the optimization problem:

$$\text{(“Hospital capacity”)} \quad \sum_g H_g(t) \leq K^{\mathrm{H}}(t), \quad \forall\, t \tag{24}$$

$$\text{(“Hospital turn away”)} \quad B_g^H(t) - \mu\left(p_g^H \cdot I_g(t) + \frac{p_g^H}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) \leq 0, \quad \forall\, g, t \tag{25}$$

$$\text{(“ICU capacity”)} \quad \sum_g ICU_g(t) \leq K^{\mathrm{ICU}}(t), \quad \forall\, t \tag{26}$$

$$\text{(“ICU turn away”)} \quad B_g^{ICU}(t) - \mu\left(p_g^{ICU} \cdot I_g(t) + \frac{p_g^{ICU}}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) \leq 0, \quad \forall\, g, t \tag{27}$$

$$\text{(“Viral testing capacity”)} \quad \sum_g N_g^{\mathrm{Vtest}}(t) \leq K^{\mathrm{Vtest}}(t), \quad \forall\, t \tag{28}$$

$$\text{(“Fractional lockdown”)} \quad \ell_g^a(t) \leq 1, \quad \forall\, g, a, t \tag{29}$$

$$\text{(“Non-negative decisions”)} \quad N_g^{\mathrm{Vtest}}(t), \ell_g^a(t) \geq 0, B_g^H(t) \geq 0, B_g^{ICU}(t) \geq 0 \quad \forall\, g, a, t \tag{30}$$

We denote by $\mathcal{C}(\boldsymbol{X}(t))$ the feasible set described by Eq. (24) - Eq. (30) for the vector of decisions $\boldsymbol{u}(t)$ at time t.

B. Economic Model and Optimization Objective. Here, we describe in more detail the objective we use in our optimization problem. This objective seeks to combine two different losses due to the pandemic: (i) the measurable economic loss stemming from quarantining, confinements and wages lost due to death and illness as well as (ii) the mortality from the pandemic.

We first describe how we account for economic loss in our model.

B.1. Economic Loss. These losses come from three separate sources:

Effect of quarantining. We capture the effect of quarantine and hospitalization by assuming that if at some time period an individual in group g is in one of the SEIR chambers $I_{j,g}^{\mathrm{q}}$, $\forall j \in \{a, ps, ms, ss\}, H_g$, ICU_g, then they generate no economic value. At the same time, we assume that individuals in R_g^{q} generate economic value as they would under no confinement.

Effect of confinement. To account for confinement in the non-quarantined population, we make the economic value per day generated by an individual in group g in the remaining (non-quarantined) SEIR chambers explicitly depend on the enforced confinement in the population. Recall that for a group g, the activity levels $\boldsymbol{\ell}_g$ specify the level of each activity allowed for that group as compared to normal course, and $\boldsymbol{\ell} = \{\boldsymbol{\ell}_g\}_{g \in \mathcal{G}}$. We denote the economic value generated by a member of g per day by $v_g(\boldsymbol{\ell})$. We remark that $v_g(\mathbf{1})$ corresponds to the economic value generated by an individual under normal circumstances.

The $v_g(\boldsymbol{\ell})$ specific to a group can be of two types: (a) wages from employment and (b) future wages from employment due to schooling. Naturally, depending on the age group, both, one, or neither of these will actually contribute to economic value. Distinguishing whether the specific group is comprised of school age, employable or retired population, we define

$$v_g(\boldsymbol{\ell}) := \begin{cases} v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{0-9 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) + v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{10-19 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) & \text{if } g = \text{20-29, 30-39, 40-49, 50-59, 60-69 y.o.} \\ 0 & \text{otherwise.} \end{cases} \tag{31}$$

We break down the definitions of $v_g^{\text{schooling}}(\boldsymbol{\ell})$ and $v_g^{\text{employment}}(\boldsymbol{\ell})$ below:

- **Value from employment $v_g^{\text{employment}}(\boldsymbol{\ell})$.** The value generated from employment is a function of the confinement level in the work activity, but also of the confinement levels in leisure, transport, as well as other activities. As an example, we expect the economic value of those employed in restaurants, retail stores, etc. to depend on foot traffic levels, which in turn are driven by the confinement levels in leisure, transport and other activities across all groups.

  Our model for employment value is a linear parametrization of these confinement decisions; specifically, $v_g^{\text{employment}}(\boldsymbol{\ell})$ is linear in ℓ_g^{work} and the average of $\boldsymbol{\ell}^{\text{transport}}, \boldsymbol{\ell}^{\text{leisure}}$ and $\boldsymbol{\ell}^{\text{other}}$ across these three activities and all groups $g \in \mathcal{G}$:

$$v_g^{\text{employment}}(\boldsymbol{\ell}) := w_g \cdot \left(\nu^{\text{work}} \cdot \ell_g^{\text{work}} + \nu^{\text{other activities}} \cdot \left(\frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} \frac{\boldsymbol{\ell}_h^{\text{transport}} + \boldsymbol{\ell}_h^{\text{leisure}} + \boldsymbol{\ell}_h^{\text{other}}}{3}\right) + \nu^{\text{fixed}}\right). \tag{32}$$

  Additionally, $\nu^{\text{work}}, \nu^{\text{other activities}}$ and ν^{fixed} are activity level sensitivity parameters such that $\nu^{\text{work}} \cdot 1 + \nu^{\text{other activities}} \cdot 1 + \nu^{\text{fixed}} = 1$; under fully open activity, they induce a multiplier of 1 in Eq. (32). Then, w_g measures the overall employment value of a typical member of group g under no confinement, and is equal to $v_g^{\text{employment}}(\mathbf{1})$.

We estimate the coefficients of this model from data, as we describe in detail in Section 4.

- **Value from schooling $v_g^{\text{schooling}}(\boldsymbol{\ell})$.** A day of schooling for the individuals in relevant groups results in economic value, equal to a day of wages that members of these groups would gain in the future.

  We use the salary of the 20-29 year-old group, multiplied by a factor, and we discount for a number of years corresponding to the difference between the midpoint of the age group and the beginning of the 20-29 year-old group. For instance, for the 0-9 year olds, we would discount over 15 years, and for 10-19 year olds we discount over five years. That is, the discounting factor we apply is

$$\delta_g := \begin{cases} (1+r)^{-15} & \text{if } g = \text{0-9 y.o.} \\ (1+r)^{-5} & \text{if } g = \text{10-19 y.o.,} \\ 0 & \text{otherwise,} \end{cases}$$

  where r is the discount rate. We further multiply the wage by f_g, which is the fraction of group g that is in school.* Lastly, we also use a multiplicative factor θ. We use this for sensitivity analysis: θ reflects that an additional day of schooling may have a multiplier effect in future wages, as well as the fact that schooling can be continued online during lockdowns. We provide ranges for θ in Section 6.

  Thus, the definition for value of school days becomes

$$v_g^{\text{schooling}}(\boldsymbol{\ell}) := \theta \cdot f_g \cdot \delta_g \cdot v_{\text{20-29 y.o}}^{\text{employment}}(\mathbf{1}) \cdot \ell_g^{\text{school}}. \qquad [33]$$

Effect of lost wages due to death. We account for a deceased individual's loss of wages which they would have earned from their current age until retirement age, given the prevailing wage curve under status quo conditions given by $\{v_g(\mathbf{1})\}_{g\in\mathcal{G}}$. For group g, we set the current age to the mid-point of the age group. We discount the resulting cash flows by an annualized interest rate. We denote the resulting lost wages amount by v_g^{life}.

For instance, for someone in age group 30-39 y.o., we calculate this cash flow by†

$$\begin{aligned} v_{30-39\text{ y.o.}}^{\text{life}} := \sum_{\tau=35}^{69} \frac{1}{(1+r)^{\tau-35}} \\ \cdot \Big(\mathbb{1}\{35 \le \tau \le 39\} \cdot v_{30-39\text{ y.o.}}^{\text{employment}}(\mathbf{1}) + \mathbb{1}\{40 \le \tau \le 49\} \cdot v_{40-49\text{ y.o.}}^{\text{employment}}(\mathbf{1}) \\ + \mathbb{1}\{50 \le \tau \le 59\} \cdot v_{50-59\text{ y.o.}}^{\text{employment}}(\mathbf{1}) + \mathbb{1}\{60 \le \tau \le 69\} \cdot v_{60-69\text{ y.o.}}^{\text{employment}}(\mathbf{1}) \Big). \end{aligned} \qquad [34]$$

We are now ready to define the economic loss component of the objective. We first define a quantity V which represents the economic value that would be generated in all groups $g \in \mathcal{G}$ under a "no pandemic" scenario. More precisely, to calculate V we assume that at time $t = 0$ all the infected or exposed population is instantaneously healed and able to generate full economic value $v_g(\mathbf{1})$. Thus,

$$V := \sum_{t=0}^{T-1} \sum_{g\in\mathcal{G}} v_g(\mathbf{1}) \cdot \left(N_g(0) + \sum_{j\in\{a,ps,ms,ss\}} I_{j,g}^{\text{q}}(0) + H_g(0) + ICU_g(0) + R_g^{\text{q}}(0) \right). \qquad [35]$$

We note that this term is a constant and does not depend on the policy followed by the policy maker.

The economic loss is then defined as the difference between V and the (potentially negative) value generated during the pandemic, and which depends on the policy followed by the policy maker:

$$\text{Economic Loss}(\boldsymbol{u}_{0:T-1}) := V - \sum_{t=0}^{T-1} \sum_{g\in\mathcal{G}} \Big(v_g\left(\boldsymbol{\ell}(t)\right) \cdot \left(S_g(t) + E_g(t) + I_g(t) + R_g(t)\right) + v_g(\mathbf{1}) \cdot R_g^{\text{q}}(t) \Big) + \sum_{g\in\mathcal{G}} v_g^{\text{life}} \cdot D_g(T), \qquad [36]$$

B.2. Total Deaths. We associate each death with a non-pecuniary cost parametrized by χ. Similarly to (4), our approach is to characterize the frontier between deaths and economic losses. For any cost of death χ, maximizing the objective function for that χ will yield a particular point on the frontier. One can trace the frontier by varying χ. We focus on characterizing the frontier, rather than selecting an optimal point along the frontier. Selecting an optimum would entail determining a choice for the value of life, which is not desirable, as there is disagreement about the right value.

We define

$$\text{Total Deaths}(\boldsymbol{u}_{0:T-1}) := \sum_{g\in\mathcal{G}} D_g(T). \qquad [37]$$

*This is due to the fact that a small fraction of the members of the 10-19 year old group are already in workforce. We do not count the value of lost schooling for them.

†Note that we are counting the entire value for the 60-69 y.o. age group; this is due to the fact that, as we explain in Section 4, this value has already incorporated that only a fraction of the population in the 60-69 y.o. age group have work-eligible ages (60-64 y.o.), while the rest are retired.

C. Optimization Problem. We can now formulate the overall optimization problem we are solving. Our objective captures the trade-off between the cost of the pandemic in terms of lost lives, and the cost in terms of foregone economic output (the economic loss). In addition, we require that the state trajectory follow the SEIR dynamics, and that the controls and state trajectory respect the capacity and feasibility constraints discussed above. We can then formulate this problem as

$$\min_{\boldsymbol{u}_{0:T-1}} \text{Economic Loss}(\boldsymbol{u}_{0:T-1}) + \chi \cdot \text{Total Deaths}(\boldsymbol{u}_{0:T-1}) \quad [38]$$

$$\text{s.t.} \quad \boldsymbol{X}(t) + F_t\left(\boldsymbol{X}(t), \boldsymbol{u}(t)\right), \forall 0 \leq t \leq T-1 \quad [39]$$

$$\boldsymbol{u}(t) \in \mathcal{C}(\boldsymbol{X}(t)), \qquad \forall 0 \leq t \leq T-1. \quad [40]$$

3. ROLD: Re-Optimization with Linearized Dynamics

In this section, we describe our approach to solving problem Eq. (38)-Eq. (40) via a linearization heuristic. Before delving into this, we make a few comments on why is it difficult to solve this problem to optimality.

Observe that a key dynamic in any SEIR-type model is the rate of new infections that multiplies the current susceptible population in a given group with the infected population in another group and which already introduces non-linearity in the state trajectory. This can be illustrated with our dynamic for the evolution of the susceptible population in group g, Eq. (4), which we replicate here:

$$\Delta S_g(t) = -\beta(t) \cdot S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \cdot \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)} \right).$$

It can be easily seen that expanding out $S_g(t)$ produces a nested fraction of polynomials in all the past decisions $\ell(\tau), N^{\mathrm{Vtest}}(\tau)$ for $0 \leq \tau \leq t-1$[‡]. This function has no identifiable structure that would make the resulting optimization problem tractable via convex optimization techniques. Similarly, the objective also involves products of states and controls, and suffers from the same non-linearity. Moreover, this issue persists even if one focuses on holding testing fixed and optimizing over the confinement decisions, or vice-versa.

With this in mind we focus on developing heuristics that can tractably yield good policies, and propose an algorithm, which we call Re-optimization with Linearized Dynamics, or ROLD, that builds a control policy by incrementally solving linear approximations of the true SEIR system.

A. Linearization. The key idea here is to solve the problem in a shrinking-horizon fashion, where at each time step $k = 0, \ldots, T$ we linearize the system dynamics and objective (over the remaining horizon), determine optimal confinement and testing decisions for all $k, \ldots, T$, and only implement the decisions for the current time-step.

We first describe the linearization of the dynamics. As a slight abuse of notation which helps readability, we index time with subscripts, i.e. we write $\boldsymbol{X}_t := \boldsymbol{X}(t)$ and similarly, $\boldsymbol{u}_t := \boldsymbol{u}(t)$. Note that we can write the true evolution of our dynamical system as:

$$\frac{\Delta \boldsymbol{X}_t}{\Delta t} = F_t(\boldsymbol{X}_t, \boldsymbol{u}_t), \quad [41]$$

where the (multi-)function F_t is given by Eq. (3)-Eq. (13) and is non-linear. The typical approach in dynamical systems is to linearize the system dynamics around a particular "nominal" trajectory. More precisely, assume that at time k we have access to a nominal control $\hat{\boldsymbol{u}}_{k:T-1}$ and let $\hat{\boldsymbol{X}}_{k:T}$ denote the resulting nominal system trajectory under the true dynamic Eq. (41) and under $\hat{\boldsymbol{u}}_{k:T-1}$. We approximate the original dynamics through a Taylor expansion around $\left(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t\right)$:

$$\frac{\Delta \boldsymbol{X}_t}{\Delta t} \approx F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) + \nabla_X F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t) + \nabla_u F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{u}_t - \hat{\boldsymbol{u}}_t), \quad [42]$$

where $\nabla_X F_t$ and $\nabla_u F_t$ denote the Jacobians with respect to $\boldsymbol{X}_t$ and $\boldsymbol{u}_t$, respectively. Note that these Jacobians are evaluated at points on the nominal trajectory, so Eq. (42) is indeed a linear expression of $\boldsymbol{X}_t$ and $\boldsymbol{u}_t$. Additionally, by induction, every state $\boldsymbol{X}_t$ under dynamic Eq. (42) will be a linear function of $\boldsymbol{u}_\tau$ for $\tau < t$. Similarly, we can approximate the objective Eq. (38) as a linear expression of $\boldsymbol{u}_\tau$ for $\tau \leq T-1$.

The controls $\boldsymbol{u}_t$ must also satisfy feasibility constraints Eq. (24) - Eq. (30). Since our algorithm works using linear approximations, it could try to apply the update $F_t(\boldsymbol{X}_t, \boldsymbol{u}_t)$ with a control $\boldsymbol{u}_t$ for which the turn-away variables violate their respective capacity constraints. To prevent this, the algorithm always first caps $\boldsymbol{u}_t$ to ensure feasibility. We use the notation $\pi(\boldsymbol{X}_t, \boldsymbol{u}_t)$ to denote the capped controls and describe this function in Section 8 A.

Linearization procedure. We use the following heuristic for obtaining an approximate control at time k, for $k = 0, \ldots, T-1$:

1. Given the current $\boldsymbol{X}_k$, use the true dynamic Eq. (41) to build a forecast for the trajectory $\hat{\boldsymbol{X}}_t$, $t \in \{k+1, \ldots, T\}$ using some nominal control $\hat{\boldsymbol{u}}_{k:T-1}$. Apply the capping function $\pi(\cdot)$ to ensure the trajectory is feasible.

[‡] In our social contacts model, $c_{g,h}(\boldsymbol{\ell}_g(t), \boldsymbol{\ell}_h(t))$ is a polynomial.

2. Use Eq. (42) to approximate the state dynamic for periods $t \in \{k, \dots, T\}$ around the forecasted trajectory $\hat{\boldsymbol{X}}_t$. The calculation of the Jacobians is provided in Section 8 B. Note that under this approximation, all the states $\boldsymbol{X}_t$ will be **linear** functions of the controls $\boldsymbol{u}_k, \dots, \boldsymbol{u}_{T-1}$.

3. Consider the objective-to-go in the sub-problem over periods $k, \dots, T$. Since $v_g(\boldsymbol{\ell}(t))$ is linear in $\boldsymbol{u}_t$ and the states $\boldsymbol{X}_g(t)$ are also linear in controls $\boldsymbol{u}_k, \dots, \boldsymbol{u}_{t-1}$, this objective will contain **bilinear** terms. More precisely, the objective can be written compactly as

$$\sum_{t=k}^{T-1} \left(\boldsymbol{u}_t^\top \boldsymbol{M} \boldsymbol{X}_t + \boldsymbol{\gamma}^\top \boldsymbol{X}_t \right) + \boldsymbol{\eta}^\top \boldsymbol{X}_T, \qquad [43]$$

for some matrix $\boldsymbol{M}$ with dimensions $|\mathcal{G}||\mathcal{U}| \times |\mathcal{G}||\mathcal{X}|$, and vectors $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ of dimensions $|\mathcal{G}||\mathcal{X}| \times 1$, which we precisely compute in Appendix C. We also linearize it using a Taylor approximation, i.e., we consider the following objective instead:

$$\sum_{t=k}^{T-1} \left(\hat{\boldsymbol{u}}_t^\top \boldsymbol{M} \hat{\boldsymbol{X}}_t + \hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top (\boldsymbol{u}_t - \hat{\boldsymbol{u}}_t) + \hat{\boldsymbol{u}}_t^\top \boldsymbol{M} (\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t) + \boldsymbol{\gamma}^\top \boldsymbol{X}_t \right) + \boldsymbol{\eta}^\top \boldsymbol{X}_T. \qquad [44]$$

This objective is now linear in all the decisions $\boldsymbol{u}_k, \dots, \boldsymbol{u}_{T-1}$.

4. Determine decision variables $\boldsymbol{u}^{*,(k)}_{k:T-1}$ to maximize the objective-to-go Eq. (44) subject to all the relevant constraints. Similarly, all constraints are linear in these decisions.

5. Update the true states using the optimal $\boldsymbol{u}^*_k := \boldsymbol{u}^{*,(k)}_k$ and using the proportional turn-away rule Eq. (22) - Eq. (23).

The linearization described above is run for all periods $k = 0, \dots, T-1$ sequentially to output a full control policy $\boldsymbol{u}^*_{k:T-1}$. We emphasize that although we allow the turn-away variables $B^H_g(t)$ and $B^{ICU}_g(t)$ to be optimized when building and optimizing the linearized model above, when we apply $\boldsymbol{u}^*_t$ in the real dynamical system, we discard the turn-away values generated by the linearization and instead set the turn-away variables to their proportional values from Eq. (22), Eq. (23) as reported in Section 2 A.3.[§]

Our algorithm needs to set an initialization for the nominal control $\hat{\boldsymbol{u}}_{k:T-1}$ at period k. In the algorithmic details that follow (Section 3 B), we specify a choice for the period k initialization, which we denote by $\hat{\boldsymbol{u}}^{(k)}_{k:T-1}$.

Trust region implementation. In our experiments, we have found that the linearized model described above may diverge significantly from the real dynamical system when the controls diverge sufficiently from the nominal-controls. This linearization error leads to a large sensitivity to the initialization $\hat{\boldsymbol{u}}^{(k)}_{0:T-1}$ around which we construct the Taylor approximations. For example, if we construct the Taylor approximation around a policy of full confinement, the linearized model would underestimate the effect of completely relaxing all the confinements, with respect to the real dynamics. Intuitively this is because the linearization errors are increasing in $||\boldsymbol{u}^{(k)}_{0:T-1} - \hat{\boldsymbol{u}}^{(k)}_{0:T-1}||$.

We overcome this difficulty by employing an iterative procedure for optimizing within each period k; this procedure resembles a trust region optimization method. The key idea is that, in order to avoid the large approximation errors, we run the linearization iteratively, with each linearization only allowed to take a small step towards the optimum within a trust region of an ϵ-ball around $\hat{\boldsymbol{u}}_{k:T-1}$, and with the updated control of each linearization being used as seed for the next. This leads to a procedure that is much more robust to the initial guess of control sequence, albeit at the expense of increased computation time.

B. Algorithmic Details. In this section, we clarify the algorithmic details of the linearization procedure described above. We first focus on how we build a linear model given $k, \boldsymbol{X}_k$ and $\hat{\boldsymbol{u}}_{k:T-1}$.

Linearized Dynamics. In step 2, the algorithm builds a approximation of the state dynamics that is linear in the controls $\boldsymbol{u}_k, \dots, \boldsymbol{u}_{T-1}$. Here, we compute the coefficients in front of each $\boldsymbol{u}_t$ explicitly. Let us introduce the notation:

$$A_t := \mathbb{I} + \Delta t \cdot \nabla_X F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \qquad [45]$$

$$B_t := \Delta t \cdot \nabla_u F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \qquad [46]$$

$$\boldsymbol{c}_t := \Delta t \cdot \left[F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) - \nabla_X F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \cdot \hat{\boldsymbol{X}}_t - \nabla_u F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \cdot \hat{\boldsymbol{u}}_t \right], \qquad [47]$$

where matrix A_t has dimensions $|\mathcal{G}||\mathcal{X}| \times |\mathcal{G}||\mathcal{X}|$, matrix B_t has dimensions $|\mathcal{G}||\mathcal{X}| \times |\mathcal{G}||\mathcal{U}|$, and vector c_t has dimensions $|\mathcal{G}||\mathcal{X}| \times 1$. With this, we have

$$\boldsymbol{X}_{t+1} = A_t \boldsymbol{X}_t + B_t \boldsymbol{u}_t + \boldsymbol{c}_t, \quad t = 0, \dots T-1.$$

[§] We note that Eq. (22) - Eq. (23) are non-linear dependencies which we could also linearize in the heuristic, at the cost of increased complexity. We instead chose a simpler architecture for the heuristic, which keeps the turn-away variables free in the linearization.

We can then express the state $\boldsymbol{X}_t$ as:[¶]

$$\boldsymbol{X}_t = \left(\prod_{\tau=t-1,t-2,\ldots,k} A_\tau\right)\boldsymbol{X}_k + \sum_{\tau=k}^{t-1}\left(\prod_{i=t-1,t-2,\ldots,\tau+1} A_i\right)B_\tau \boldsymbol{u}_\tau + \sum_{\tau=k}^{t-1}\left(\prod_{i=t-1,t-2,\ldots,\tau+1} A_i\right)\boldsymbol{c}_\tau. \quad [48]$$

It is now possible to express both objective and constraints linearly in the decision $\boldsymbol{u}_t$.

Constraint coefficients. We can write each of the constraints Eq. (24), Eq. (26) and Eq. (28) in the form

$$L_t \leq K(t), \text{ where } L_t := {\boldsymbol{\gamma}_x}^\top \cdot \boldsymbol{X}_t + {\boldsymbol{\gamma}_u}^\top \cdot \boldsymbol{u}_t, \forall\, t \in \{k, k+1, \ldots, T-1\}, \quad [49]$$

for some (time-invariant) $\boldsymbol{\gamma}_x, \boldsymbol{\gamma}_u$. Since $\boldsymbol{X}_t$ is linear in $\boldsymbol{u}_k, \ldots, \boldsymbol{u}_{t-1}$, to represent one such constraint we just need to store the coefficients corresponding to all decision variables (i.e., $\boldsymbol{u}_k, \ldots, \boldsymbol{u}_{T-1}$) and the free terms/constants that appear in L_t.

In particular, in the LHS ${\boldsymbol{\gamma}_x}^\top \cdot \boldsymbol{X}_t + {\boldsymbol{\gamma}_u}^\top \cdot \boldsymbol{u}_t$ of such a constraint, the decision $\boldsymbol{u}_\tau$, for $k \leq \tau \leq t$, will have coefficients

$$\begin{cases} {\boldsymbol{\gamma}_u}^\top, & \text{for } \tau = t \\ {\boldsymbol{\gamma}_{\boldsymbol{x}}}^\top \left(\prod_{i=t-1,t-2,\ldots,\tau+1} A_i\right) B_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases} \quad [50]$$

To make calculations efficient, we note that the coefficients can be obtained recursively as in the Calculate ConstraintCoefficients function defined in Algorithm 1.

Algorithm 1 CalculateConstraintCoefficients function.

Require: $\bar{X}_k := \boldsymbol{X}_k$
  for $t = k, k+1, \ldots, T-1$ **do**
    Calculate the constants in L_t as: ${\boldsymbol{\gamma}_x}^\top \bar{X}_t$
    Set $\bar{X}_{t+1} := A_t \bar{X}_t + c_t$
    Calculate coefficients for $\boldsymbol{u}_\tau$ in L_t as:

$$\begin{cases} {\boldsymbol{\gamma}_u}^\top, & \text{for } \tau = t \\ {\boldsymbol{\gamma}_x}^\top \bar{A}_{t,\tau} B_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases}$$

  where

$$\bar{A}_{t,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = t-1 \\ \bar{A}_{t,\tau+1} A_{\tau+1}, & \text{for } \tau = t-2, t-3, \ldots, k. \end{cases} \quad [51]$$

  return Coefficients for $\boldsymbol{u}_\tau$ in L_t in each $t = k, k+1, \ldots, T-1$.

Objective coefficients. Up to constants, Eq. (44) can be written as

$$\sum_{t=k}^{T} \left({\boldsymbol{d}_t}^\top \boldsymbol{X}_t + {\boldsymbol{e}_t}^\top \boldsymbol{u}_t\right) \quad [52]$$

with

$${\boldsymbol{d}_t}^\top = \begin{cases} \hat{\boldsymbol{u}}_t^\top \boldsymbol{M} + \boldsymbol{\gamma}^\top, & \text{if } t < T \\ \boldsymbol{\eta}^\top, & \text{if } t = T, \end{cases} \qquad {\boldsymbol{e}_t}^\top = \begin{cases} \hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top, & \text{if } t < T \\ 0, & \text{if } t = T. \end{cases} \quad [53]$$

In Eq. (44) the decisions $\boldsymbol{u}_t$, for $k \leq t \leq T-1$, will have objective coefficients:

$$\hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top + \sum_{\tau=t+1}^{T-1} \left(\hat{\boldsymbol{u}}_\tau^\top \boldsymbol{M} + \boldsymbol{\gamma}^\top\right)\left(\prod_{i=\tau-1,\tau-2,\ldots,t+1} A_i\right) B_t + \boldsymbol{\eta}^\top \left(\prod_{i=T-1,T-2,\ldots,t+1} A_i\right) B_t. \quad [54]$$

Then, the objective coefficients in front of each $\boldsymbol{u}_t$ can be calculated in a recursive manner similar to the one we use for the constraint coefficients. The detailed function CalculateObjectiveCoefficients is defined in Section 8 C.1.

Specifics of Iterative Linearization Procedure.

Algorithm 2 LINEARIZATION function.

Require: time period k, starting state $\boldsymbol{X}_k$, nominal control initialization $\hat{\boldsymbol{u}}_{k:T-1}$

Calculate forecast $\hat{\boldsymbol{X}}_{k:T}$ from $\boldsymbol{X}_k$ and $\hat{\boldsymbol{u}}_{k:T-1}$

Run CALCULATEOBJECTIVECOEFFICIENTS and CALCULATECONSTRAINTCOEFFICIENTS to compute linear coefficients of decisions $\boldsymbol{u}_t,\ t = k, \ldots, T-1$

Solve resulting linear program with additional trust region constraints $\boldsymbol{u}_{k:T-1} \in B_\epsilon^\infty(\hat{\boldsymbol{u}}_{k:T-1})$ to obtain controls $\boldsymbol{u}_{k:T-1}^{*,(k)}$

return $\boldsymbol{u}_{k:T-1}^{*,(k)}$.

Having defined these functions, the LINEARIZATION function is described in Algorithm 2. This function builds the linear approximation for the remaining trajectory of the system, and optimizes it in a trust region of an infinity-norm ϵ-ball around the initial nominal control $\hat{\boldsymbol{u}}_{k:T-1}$. We denote that ϵ-ball by $B_\epsilon^\infty(\hat{\boldsymbol{u}}_{k:T-1})$.

Within each period k, we call the LINEARIZATION function iteratively up to a termination condition, using the output control to initialize the nominal control and the trust region for the next call of the function. This still requires us to choose an initialization of the $k = 0$ nominal control $\hat{\boldsymbol{u}}_{0:T-1}^{(0)}$; in our experiments we initialize this with a heuristic solution generated via a gradient method. The pseudo-code for this procedure is Algorithm 3.

For the termination conditions of the initial iterations, we combine a fixed upper bound on the number of iterations with a condition that we do not repeat control sequences $\boldsymbol{u}_k^{*,(k)}$, in order to avoid cycles. The fixed upper bound on the number of iterations is set so as to ensure that for each k, every confinement decision in $\boldsymbol{u}_{k:T-1}^{*,(k-1)}$ can be changed to any value in $[0, 1]$ with ϵ-length steps. In particular, for all the runs presented, we take the upper bound of steps to be $\frac{mult}{\epsilon}$, with $\epsilon = 0.05$ and $mult = 2$, resulting in an upper bound of 40 runs for the inner loop.

We experimented with different values of ϵ between 0.01 and 0.5 and values of $mult$ between 1 and 5. As expected, lower values of ϵ resulted in a more stable and higher performing heuristic, while higher values of $mult$ improved the heuristic only up to around 2, after which point the non-cycling termination condition was triggered almost always. On the other hand, reducing ϵ did have a significant impact on the run-time of the linearization algorithm. We chose the combination of ϵ and $mult$ that gave us the best trade-off between the quality of the solution and the total run-time.

Algorithm 3 Iterative linearization algorithm.

Require: initial state $\boldsymbol{X}_0$ and initial nominal control for first period $\hat{\boldsymbol{u}}_{0:T-1}^{(0)}$
for $k = 0, \ldots, T-1$ **do**
Initialize

$$\hat{\boldsymbol{u}}_{k:T-1}^{(k)} = \begin{cases} \hat{\boldsymbol{u}}_{0:T-1}^{(0)}, & \text{if } k = 0 \\ \boldsymbol{u}_{k:T-1}^{*,(k-1)}, & \text{otherwise} \end{cases} \quad [55]$$

$$\hat{\boldsymbol{X}}_k = \begin{cases} \boldsymbol{X}_0, & \text{if } k = 0 \\ \text{Calculate } \boldsymbol{X}_k \text{ from } \hat{\boldsymbol{X}}_{k-1} \text{ and } \boldsymbol{u}_{k-1}^{*,(k-1)}, & \text{otherwise} \end{cases} \quad [56]$$

while termination condition? **do**
$\boldsymbol{u}_{k:T-1}^{*,(k)} = \text{LINEARIZATION}(k, \hat{\boldsymbol{X}}_k, \hat{\boldsymbol{u}}_{k:T-1}^{(k)})$
$\hat{\boldsymbol{u}}_{k:T-1}^{(k)} = \boldsymbol{u}_{k:T-1}^{*,(k)}$

Set $\boldsymbol{u}_k^*$ equal to the confinement and testing decisions of $\boldsymbol{u}_k^{*,(k)}$, and use the proportional turn-away rule defined in Eq. (22) - Eq. (23)

4. Parametrization and Calibration of the Model for Île-de-France

We parametrize the model using SEIR parameter estimates from (5). We calibrate the model using data on health outcomes, Google mobility data, as well as economic output data. Overall, we use data specific to the the Île-de-France region, or specific to France when Île-de-France-specific data is not available.

A. Basic SEIR Model Parameters. The SEIR model parameters that pertain across the age groups are summarized in Table S1. The age-group specific parameters are reported in Table S2. We start with the parameters as reported in (5),[‖] and then we allow the values of parameters $\left\{\beta, \lambda_H^{-1}, \lambda_{ICU}^{-1}, \{p_g^{ICU}\}_{g\in\mathcal{G}}, \{p_g^D\}_{g\in\mathcal{G}}\right\}$ to change through time, as detailed in Section 4 B. We allow these parameters to change in time in order to model changes in the way hospitals manage COVID patients, as well as changes in mandates for using masks and other measures that reduce transmission.

For R_0 and λ_H, the reported uncertainty ranges are 95% confidence intervals. For σ^{-1} (i.e., the mean stay in compartment E), the uncertainty range is calculated as $4 \pm 0.8 \cdot 0.6$, where 0.6 is half the width of the 95% confidence interval for the incubation period reported in (6), and 0.8 accounts for the fact that the stay in compartment E is 4/5 of the mean incubation time in (5). For μ^{-1} (i.e., the mean stay in an infectious state), the uncertainty range is calculated as 4 ± 0.43, where 0.43 is half the width of the 95% confidence interval for the serial interval reported by (7).[**] For the average stay in the ICU, we add to the mean stay of 20.46 days for Île-de-France, another 1.5 day, which is the mean time spent in hospital prior to ICU admission (5).

[¶] In the expressions in Eq. (48), Eq. (50), Eq. (54), we follow the convention that a product of matrices over an empty set of indices results in the identity matrix.

[‖] We retrieve the parameter values as reported before (5) got updated on July 8, 2020.

[**] We note that (7) estimate the serial interval, and not the infectious period, to be 3.96 days. We borrow their confidence interval for the serial interval estimation and use it as an uncertainty range for our infectious period, which is of about the same length as their estimated serial interval.

Table S1. SEIR model parameters

Variable	Description	Value [Uncertainty Range]	Reference
R_0	Basic reproduction number	2.9 [2.8 - 2.99]	(5)
σ^{-1}	Latency period	4.0 days [3.57 - 4.43]	(5, 6)
μ^{-1}	Infectious period	4.0 days [3.52 - 4.48]	(5, 7)
$(\lambda^H)^{-1}$	Average time spent in Hospital	14.94 days [14.55 - 15.32]	(5)
$(\lambda^{ICU})^{-1}$	Average time spent in ICU	21.96 days [21.38 - 22.6]	(5)

Table S2. Age-group specific SEIR model probability parameters

Age group g	$p_{ss,g}$ Prob. of severe symptoms given infection	p_g^{ICU} Prob. of ICU given infection	p_g^H Prob. of Hospital given infection	p_g^D Prob. of Death given severe symptoms
0-9	0.002 [0.001 - 0.003]	0.000444 [0.000192 - 0.000765]	0.001556 [0.000235 - 0.002808]	0.006 [0.003 - 0.013]
10-19	0.002 [0.001 - 0.003]	0.000444 [0.000192 - 0.000765]	0.001556 [0.000235 - 0.002808]	0.006 [0.003 - 0.013]
20-29	0.006 [0.004 - 0.01]	0.00069 [0.000404 - 0.00132]	0.00531 [0.00268 - 0.009596]	0.011 [0.007 - 0.016]
30-39	0.013 [0.008 - 0.02]	0.002067 [0.001168 - 0.00346]	0.010933 [0.00454 - 0.018832]	0.019 [0.015 - 0.023]
40-49	0.017 [0.01 - 0.027]	0.003774 [0.0021 - 0.006345]	0.013226 [0.003655 - 0.0249]	0.033 [0.029 - 0.037]
50-59	0.035 [0.021 - 0.054]	0.00966 [0.005565 - 0.015498]	0.02534 [0.005502 - 0.048435]	0.065 [0.06 - 0.07]
60-69	0.071 [0.042 - 0.11]	0.021868 [0.012516 - 0.03498]	0.049132 [0.00702 - 0.097484]	0.126 [0.12 - 0.132]
70-79	0.113 [0.067 - 0.175]	0.028137 [0.016147 - 0.04515]	0.084863 [0.02185 - 0.158853]	0.21 [0.203 - 0.218]
80+	0.32 [0.19 - 0.494]	0.01792 [0.01007 - 0.029146]	0.30208 [0.160854 - 0.48393]	0.316 [0.309 - 0.324]

Calculating the transmission rate β from R_0. We obtain β by linearizing the dynamics for E_g, I_g around a point where $S_h \approx N_h, I_h \approx 0, \forall\, h$. More precisely, we have:

$$\frac{\partial E_g}{\partial t} \approx \beta(t) \sum_h c_{gh} I_h(t) - \sigma E_g(t)$$
$$\frac{\partial I_g}{\partial t} \approx \sigma E_g(t).$$

Then, with $\boldsymbol{Y}(t) := \left[E_1(t), E_2(t), \ldots, E_{|\mathcal{G}|}, I_1(t), \ldots, I_{|\mathcal{G}|}\right]^T$, we can write $\dot{\boldsymbol{Y}}(t) = (\boldsymbol{\Phi} + \boldsymbol{\Lambda})\boldsymbol{Y}(t)$, where

$$\boldsymbol{\Phi} = \beta \cdot \left[\begin{array}{c|c} \mathbf{0} & [c_{gh}]_{g,h\in\mathcal{G}} \\ \hline \mathbf{0} & \mathbf{0} \end{array}\right] \qquad [57]$$

and

$$\boldsymbol{\Lambda} = \left[\begin{array}{c|c} -\mathbf{diag}(\boldsymbol{\sigma}) & \mathbf{0} \\ \hline \mathbf{diag}(\boldsymbol{\sigma}) & -\mathbf{diag}(\boldsymbol{\mu}) \end{array}\right] \qquad [58]$$

Then R_0 can be identified as the spectral radius (i.e., the largest absolute value of the eigenvalues) of the matrix $-\boldsymbol{\Phi}\boldsymbol{\Lambda}^{-1}$ (8, 9). Since the eigenvalues of a matrix $\beta \cdot \boldsymbol{A}$ are simply β multiples of the eigenvalues of $\boldsymbol{A}$, we can therefore determine β as R_0 divided by the spectral radius of the matrix $(-\boldsymbol{\Phi}/\beta)\,\boldsymbol{\Lambda}^{-1}$.

B. Epidemiological Model Parameter Fitting Using Health Outcomes and Mobility Data. We use data on health outcomes from the French Public Health Agency (10), as well as Google mobility data (11), to estimate the unknown parameters in our model. The data on health outcomes includes counts for individuals who are in the hospital, in the ICU, and have died, by age group, and is maintained and updated daily by the French Public Health Agency (Santé publique France). The Google mobility data reports changes in activity at different places compared to a baseline, and is calculated using aggregate and anonymized data. For both health outcomes and mobility, we use data specific to the Île-de-France region.

The calibration exercise has two purposes: (a) to further refine the SEIR parameters reported in the literature to the data observed in Île-de-France and (b) to estimate our other parameters for which we do not have existing references. We then use the values of the estimated parameters in our experiments and simulations.

We first describe the set of parameters to be estimated, which we denote by $\boldsymbol{\mathcal{P}}$.

- Date of patient zero. We assume that the SEIR process starts with an infected individual of the 40-49 y.o. age group (12). We wish to estimate the date when this infection occurs.

- Epidemiological parameters. We use the epidemiological parameters of (5) to initialize the SEIR model. We allow these parameters to change in time in order to model changes in the way hospitals manage COVID patients, as well as changes in mandates for using masks and other measures that reduce transmission. We assume that on date d, each parameter from the set $\mathcal{S} = \left\{\beta, \lambda_H^{-1}, \lambda_{ICU}^{-1}, \{p_g^{ICU}\}_{g\in\mathcal{G}}, \{p_g^D\}_{g\in\mathcal{G}}\right\}$ changes with respect to its initial value (as reported in Section A), according to the relationship
$$s_{\text{after_}d} = m_s \cdot s_{\text{before_}d}, \quad s \in \mathcal{S}, \tag{59}$$
where m_s is a multiplier pertaining to parameter s. We assume the same multiplier $m_{p_g^{ICU}}$ for all groups $g \in \mathcal{G}$, and similarly for $m_{p_g^D}$. We seek to determine the date of change d as well as the multipliers $m_s, s \in \mathcal{S}$.

- Confinement patterns. To estimate activity levels for the activities in our social mixing model, we use Google mobility data (11). The mobility data reports changes of activity (visits and length of stay) for each day, compared to a baseline value. The baseline used corresponds to the median value for the corresponding day of the week, during the five-week period January 3 - February 6 2020.

  We fix the "home" activity level to be equal to 1, throughout time. We estimate the level of the other activities using the corresponding activities from the Google mobility data, as shown in Table S3.

Table S3. Mapping between the activities in our model and the activities in the Google mobility data

Activity in Our Model	Activity in Google Mobility Data
Work	Workplaces
Transport	Transit stations
Leisure	Retail & recreation
Other	$\alpha_{\text{other}}\cdot$Retail & recreation $+(1-\alpha_{\text{other}})\cdot$ Grocery & pharmacy

  What remains to be estimated in the calibration is the weight parameter α_{other}, as well as the school activity levels. We calibrate the level of schooling activity for four different time periods. These periods are chosen to reflect (i) the dates when the French government closed down schools, and (ii) the French school calendar and summer recess.

- Social mixing parameter. To reduce the number of parameters to be calibrated, we simplify our mixing model in Eq. (1) by constraining $\alpha_1 = \alpha_2$. We seek to determine this mixing parameter.

We next describe the details of the fitting procedure that we set up in order to retrieve an optimal parameter fitting. The mixing dynamics of the SEIR model are driven by the vector of activity levels $\boldsymbol{\ell}_g$ of each age group (as described in Section A). Data on activity levels can be noisy; we model this uncertainty by assuming the vector of activity levels is a random vector, distributed as follows:
$$\tilde{\boldsymbol{\ell}}_g(t) = \bar{\boldsymbol{\ell}}_g(t)(\mathbf{1} + Z^t), \quad Z^t \sim N(\mathbf{0}, \sigma^2\mathbb{I}), \quad \sigma = 10^{-2}. \tag{60}$$
The value $\bar{\boldsymbol{\ell}}_g(t)$ is obtained from the Google activity data at time t; this dataset does not differentiate activity by age, so $\bar{\boldsymbol{\ell}}_g(t) = \bar{\boldsymbol{\ell}}_h(t)$ for all $g, h \in \mathcal{G}$; in other words, all groups are assigned the same activity level. Recall that $H_g(t), ICU_g(t), D_g(t)$ denote the hospital utilization, ICU utilization and cumulative number of deaths according to the SEIR model. We denote these quantities with $H_g^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}), ICU_g^{\boldsymbol{\mathcal{P}}}(t, \tilde{\ell}), D_g^{\mathcal{P}}(t, \tilde{\ell})$ to emphasize the dependence of the SEIR process on parameters $\boldsymbol{\mathcal{P}}$ and the vector of activity levels $\tilde{\boldsymbol{\ell}}$. We aggregate these quantities over all age groups:
$$H^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \quad := \quad \sum_{g\in\mathcal{G}} H_g^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \tag{61}$$
$$ICU^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \quad := \quad \sum_{g\in\mathcal{G}} ICU_g^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \tag{62}$$
$$D^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \quad := \quad \sum_{g\in\mathcal{G}} D_g^{\boldsymbol{\mathcal{P}}}(t, \tilde{\boldsymbol{\ell}}) \tag{63}$$

Denote $H^{obs}(t), ICU^{obs}(t), D^{obs}(t)$ as the general ward hospital beds utilization, ICU beds utilization, and cumulative deaths, respectively, at time t, as observed from the Île-de-France real data from (10). We calculate the relative fitting error of the SEIR model at time t for each of these three quantities as

$$\mathcal{E}_H^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) := \left|\frac{H^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) - H^{obs}(t)}{H^{obs}(t)}\right| \quad [64]$$

$$\mathcal{E}_{ICU}^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) := \left|\frac{ICU^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) - ICU^{obs}(t)}{ICU^{obs}(t)}\right| \quad [65]$$

$$\mathcal{E}_D^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) := \left|\frac{D^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) - D^{obs}(t)}{D^{obs}(t)}\right|. \quad [66]$$

We define the total expected fitting error as a sum of these errors over different time intervals:

$$\begin{aligned}\mathcal{E}^{\mathcal{P}} :=& \mathbb{E}_{\tilde{\boldsymbol{\ell}}}\Bigg[\sum_{t\in\text{Peak_H}} w^{\text{Peak}}\cdot\mathcal{E}_H^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + \sum_{t\in\text{Peak_ICU}} w^{\text{Peak}}\cdot\mathcal{E}_{ICU}^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + \sum_{t\in\text{Peak_D}} w^{\text{Peak}}\cdot\mathcal{E}_D^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) \\ &+ \sum_{t\in\text{Tail}}\left(w_H^{\text{Tail}}\cdot\mathcal{E}_H^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + w_{ICU}^{\text{Tail}}\cdot\mathcal{E}_{ICU}^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + w_D^{\text{Tail}}\cdot\mathcal{E}_D^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}})\right) \quad [67] \\ &+ \sum_{t\in\text{Main}}\left((1-w_H^{\text{Tail}})\cdot\mathcal{E}_H^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + (1-w_{ICU}^{\text{Tail}})\cdot\mathcal{E}_{ICU}^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}}) + (1-w_D^{\text{Tail}})\cdot\mathcal{E}_D^{\mathcal{P}}(t,\tilde{\boldsymbol{\ell}})\right)\Bigg], \quad [68]\end{aligned}$$

where the expectation is taken with respect to random vector $\tilde{\boldsymbol{\ell}}$, and where the time intervals are defined in Table S4 and comprise the entire period between March 17 2020 and October 21 2020. Our approach penalizes the errors at the peak times for hospitalizations, beds utilization and deaths; it also penalizes errors over the last 14 days of the considered period, to ensure an accurate fit at the end of the calibration horizon. We use different weights to account for the different errors. We use $w^{\text{Peak}} = 1/6$ to account for the relatively smaller period of the Peaks. We use $w_H^{\text{Tail}} = w_D^{\text{Tail}} = 1/3$, but set w_{ICU}^{Tail} to $2/3$. The reason for this choice is that we are interested in low error at the tail predictions of ICU utilization, as this number plays an important role in the dynamics of the model right after October 21st.

Table S4. Time intervals used in the calibration

Label	Period
Main	March 17 - October 7 2020 (except for April 8, April 14, May 14)
Peak_ICU	April 8 2020
Peak_H	April 14 2020
Peak_D	May 14 2020
Tail	October 8 - 21 2020

We seek to determine the set of parameters $\boldsymbol{\mathcal{P}}$ that minimize $\mathcal{E}^{\mathcal{P}}$ in Eq. (67). We approximate the expectation in Eq. (67) through a Monte Carlo sample-average approximation, using 100 samples. The set $\boldsymbol{\mathcal{P}}$ contains both discrete and continuous parameters. To minimize $\mathcal{E}^{\mathcal{P}}$, we first do a grid search over all possible combinations of the discrete parameters, and then for each such combination, we perform a gradient descent procedure over the space of the continuous parameters.

Our calibration procedure yields the parameter fitting summarized in Table S5. We compare the fitted values of the SEIR model with the values reported by the French National Agency of Public Health in Figure S2.

C. Economic Model Parameter Fitting. We obtain data on population, employment, and wages from the French National Institute of Statistics and Economic Studies (INSEE). Where relevant, we discount all cash flows at 3% annualized rate. We set the retirement age to be 65 (i.e., 64 is the last working year of age.)

We first obtain the initial population data $N_g(0)$ for each age group in Île-de-France at the end of 2019 from (13).

Estimation of w_g. Recall that w_g in Eq. (32) corresponds to the employment value for a member of group g, under normal conditions. To estimate w_g, we use two datasets from INSEE:

- Yearly full time equivalent (FTE)[††] wages and employed population count for Île-de-France in 2016, broken up into the age groups "under 26 years old", "26 to 49 years old" and "more than 50 years old" (14).
- FTE employment rates across the entire economy for the fourth quarter of 2019, bucketed by age groups "15 to 24 years old", "25 to 49 years old", "50 to 64 years old", and "55 to 64 years old" (15).

[††] This is a normalization to account for employees doing part time work.

Table S5. Calibration: fitted epidemiological and economic parameters and sensitivity analysis ranges

Parameter	Estimated Value	Estimation Interval
Date of patient zero	Jan 1 2020	Dec 8 2019 - Jan 27 2020
Date d when epidemiological parameters change	April 16 2020	Mar 1 2020 - Apr 30 2020
Multiplier m_β	0.60	0.5-2.0
Multiplier $m_{p_g^{ICU}}$	1.66	0.5-2.0
Multiplier $m_{p_g^D}$	1.13	0.5-2.0
Multiplier m_{λ_H}	1.50	0.5-2.0
Multiplier $m_{\lambda_{ICU}}$	2.00	0.5-2.0
Weight for activities α_{other}	0.96	0.0-1.0
School activity ℓ^{school} Mar 17 - May 10 2020	0.0	0.0-0.1
School activity ℓ^{school} May 11 - Jun 30 2020	0.0	0.0-1.0
School activity ℓ^{school} Jul 1 - Aug 31 2020	0.2	0.0-0.2
School activity ℓ^{school} Sep 1 - Oct 21 2020	1.0	0.5-1.0
Social mixing $\alpha_1 = \alpha_2$	0.39	0.1-2.0
Sensitivity of econ. value ν^{work}	0.50	
Sensitivity of econ. value ν^{fixed}	0.47	

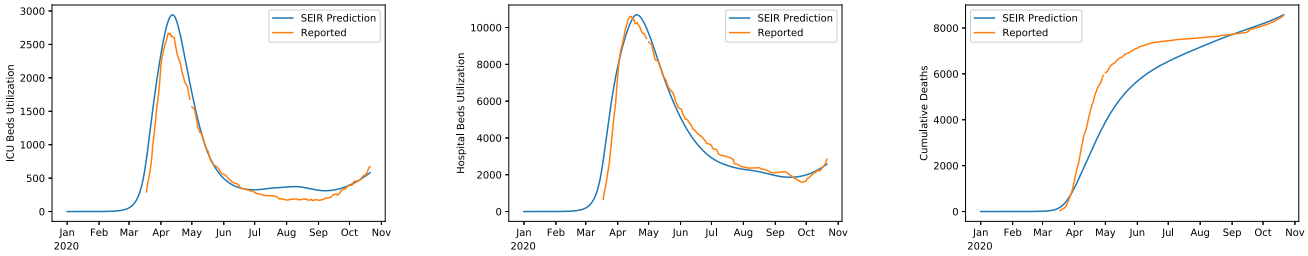

(a) ICU Beds Utilization **(b)** Hospital Beds Utilization **(c)** Cumulative Deaths

Fig. S2. Predictions of the fitted SEIR model vs. reported values by the French Public Health Agency for three different quantities: **(a)** ICU Beds Utilization, **(b)** Hospital Beds Utilization and **(c)** Cumulative Deaths.

Since we do not have a consolidated data source for economic data split by our exact age group definitions, we use the above datasets to interpolate values for w_g. At a high level, we derive wage curves across age ranges.

We next explain the general procedure, as well as the additional assumptions we have made for the interpolation. First, for the construction of wage curve by age bucket:

- We assume that the national level employment rates from (15) are equal to those of the Île-de-France region. Because the age bucketing for our age groups is finer than the age bucketing in the data, we use interpolation. Specifically, we fit a piece-wise linear model (consisting of three pieces) to the four employment rates reported for the "15 to 24 years old", "25 to 49 years old", "50 to 64 years old", and "55 to 64 years old" groups. We take the mid-point of the age group as the x value of the datapoint; for example, for "50 to 64 years old" we use a midpoint of 57.5.

  With this model, we can infer an employment rate for any arbitrary age and construct an employment rate curve.

- We perform a similar fitting procedure for the age group wage information from (14); since the wage progression by age is much smoother, we use simple linear regression to construct a wage curve for each one of our age buckets.

- The previous wage curve only accounts for the employed population, whereas our age groups count the entire population. We thus combine the wage curve with employment rate and population data to arrive at a wage number blended across an entire age group's population.

  When doing this, we treat the 10-19 y.o. and 60-69 y.o. age groups specially by assuming the employment rates are reported only with respect to the work-eligible population in that bucket (15-19 and 60-64 year olds, respectively). We also set the work-eligible population for the 0-9, 70-79, and 80+ age buckets to 0. The formula we use is

$$\text{average_wage}_g = \text{employed_average_wage}_g \times \text{employment_rate}_g \times \text{fraction_work_eligible_population}_g \qquad [69]$$

- The interpolations we use introduce errors: in particular, if we aggregate the wages inferred by our constructed curve across the entire population, we overestimate the real total wages by 5.12%. We shade down all wages average_wage$_g$ proportionally so as to retrieve the real total wage amount w_g.

Table S6 summarizes the year-based employment contribution parameters we use in the objective function per age group. We note that we divide these year-based values by 365 when using them in the objective of the optimization problem, in order to capture employment value on a daily basis.

Table S6. Year-based economic value parameters under normal activity, per age group (in €)

Age Group g	$w_g \times 365$
0-9	0.00
10-19	2084.79
20-29	11 743.20
30-39	25 799.29
40-49	31 746.68
50-59	32 573.79
60-69	12 640.83
70-79	0.00
80+	0.00

Estimation of $\nu^{\text{work}}, \nu^{\text{other activities}}, \nu^{\text{fixed}}$. We move on to the estimation of parameters $\nu^{\text{work}}, \nu^{\text{other activities}}, \nu^{\text{fixed}}$ in Eq. (32). These measure the sensitivity of economic value to the confinement pattern $\boldsymbol{\ell}(t)$. We estimate them from data on lost economic output during the first lockdown phase employed in Île-de-France, and in particular using the month of April 2020. We break up the approach into a few steps:

- We use survey data of French managers regarding business activity during the lockdown starting March 17 2020. This is sentiment data where managers are asked to compare current business conditions to normal conditions for the same relevant time period (16, 17). These data are reported by industry, and we aggregate them into a single number, weighting by industry size. We use FTE wages and employed population count for the Île-de-France region in 2016 (18) to figure out the appropriate weights to use in the aggregation.

  We then use these monthly readings as proxies for the economic activity level due to confinements in the month of April 2020, as compared to normal activity. The economic activity level for the month of April is 58.51%.

- A requirement for our estimation are the precise levels of confinement in April 2020. We retrieve these from Google mobility data (11), as explained in Section 4 B. To simplify estimation, we set $\nu^{\text{other activities}} = 0$ and then determine parameters $\nu^{\text{work}}, \nu^{\text{fixed}}$ solving the system of equations

$$\nu^{\text{work}} + \nu^{\text{fixed}} = 1 \qquad [70]$$

$$\nu^{\text{work}} \ell^{\text{work}}_{\text{April}} + \nu^{\text{fixed}} = 0.5851, \qquad [71]$$

  where $\ell^{\text{work}}_{\text{April}}$ corresponds to the average value of $\ell^{\text{work}}(t)$ through the month of April 2020. In our experiments, we also test our algorithm in alternative scenarios where we set $\nu^{\text{other activities}} > 0$, keep the value for ν^{fixed} from the system Eq. (70)-Eq. (71), and adjust $\nu^{\text{work}} = 1 - \nu^{\text{other activities}} - \nu^{\text{fixed}}$. The specific values we test are $\nu^{\text{other activities}} \in \{0, 0.1, 0.2\}$.

5. Benchmark Policies

Benchmark policies for lockdown. We compare ROLD to several simpler classes of policies drawing inspiration by real life confinement management rules:

- **ICU admissions trigger policy — ICU-t**. This class of policies is similar in spirit to the trigger rule proposed by (19) for the Austin metropolitan area. This rule places all age groups and activities (except home) at a strict level of confinement when the average seven-day inflow of hospital admissions exceeds a pre-determined threshold, and then relaxes the confinement to a relaxed level when the average seven-day hospital admissions and the hospital utilization rates drop below pre-determined thresholds.

  Since (19) does not differentiate between hospital and ICU beds, we define our policy class on ICU admissions and utilization instead of hospitalizations. Specifically, the ICU admissions trigger policy is defined as in Algorithm 4.

  We optimize over the parameters ℓ_{strict}, ℓ_{relaxed}, $\tau_{\text{admissions}}$ and $\tau_{\text{utilization}}$ via grid search with the goal of minimizing the objective in Eq. (38) corresponding to the total economic and life loss due to the pandemic, and we report the performance of the best policy.

Algorithm 4 ICU Admissions Trigger Policy — ICU-t.

Require: parameters ℓ_{strict}, ℓ_{relaxed}, $\tau_{\text{admissions}}$, $\tau_{\text{utilization}}$

Initialize $\ell_g^a(0) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$

for $t \in \{1 \dots T-1\}$, $g \in \mathcal{G}$ **do**

  if 7 day average ICU admissions $> \tau_{\text{admissions}}$ **then**

    Set $\ell_g^a(t) = \ell_{\text{strict}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$

  else if (7 day average ICU admissions $\leq \tau_{\text{admissions}}$) AND $\left(\sum_g ICU_g(t)/K^{ICU}(t) \leq \tau_{\text{utilization}}\right)$ **then**

    Set $\ell_g^a(t) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$

  else

    Set $\ell_g^a(t) = \ell_g^a(t-1)$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$

- **Hybrid trigger policy — Hybrid-t**. This policy resembles the rule used in France for declaring a region in "maximum alert"[‡‡] (20). Like the previous policy, this policy also switches between uniform strict versus relaxed confinement levels, but the trigger condition combines ICU utilization with infection prevalence rates in the population. Specifically, the policy switches to strict confinement if the seven-day incidence rate in the population, defined[§§] as $\sum_{t-6\leq\tau\leq t}\sum_g I_g(\tau)/\sum_g N_g(0)$, is greater than a threshold $\rho_{\text{incidence}}$, *and* the incidence rate in age groups corresponding to the population that is 60 y.o. and above $\sum_{t-6\leq\tau\leq t}\sum_{g\geq 60\text{ y.o.}} I_g(\tau)/\sum_g N_g(0)$ is greater than a threshold $\rho_{\text{incidence_60+}}$, *and* ICU utilization rate is greater than $\rho_{\text{ICU_utilization}}$. The specific rule is described in Algorithm 5. We optimize over all parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{incidence}}$, $\rho_{\text{incidence_60+}}$ and $\rho_{\text{ICU_utilization}}$ with the goal of minimizing the total loss objective in Eq. (38), and we report the performance of the best policy.

  This is the Hybrid-t AND policy, and it is described in Algorithm 5. We also test a stricter version of this policy that takes the logical *or* of the three conditions (Hybrid-t OR) instead of taking the *and*.

Algorithm 5 HYBRID TRIGGER POLICY — Hybrid-t AND.

Require: parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{incidence}}$, $\rho_{\text{incidence_60+}}$, $\rho_{\text{ICU_utilization}}$
  Initialize $\ell_g^a(0) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$
  for $t \in \{1 \dots T-1\}$, $g \in \mathcal{G}$ **do**
    if $\left(\sum_{t-6\leq\tau\leq t}\sum_g I_g(\tau)/\sum_g N_g(0) > \rho_{\text{incidence}}\right)$ AND $\left(\sum_{t-6\leq\tau\leq t}\sum_{g\geq 60y.o.} I_g(\tau)/\sum_g N_g(0) > \rho_{\text{incidence_60+}}\right)$ AND $\left(\sum_g ICU_g(t)/K^{ICU}(t) > \rho_{\text{ICU_utilization}}\right)$ **then**
      Set $\ell_g^a(t) = \ell_{\text{strict}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$
    else
      Set $\ell_g^a(t) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$

- **Fully open — FO.** This corresponds to the normal conditions where $\ell_g^a(t) = 1$ for all $a \in \{\text{home}, \text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}$, $g \in \mathcal{G}$ and $t \in \{0, \dots, T-1\}$.

- **Full confinement — FC.** In this policy, all activities except home are fully restricted. That is, we set $\ell_g^a(t) = 0$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}$, $g \in \mathcal{G}, t \in \{0, \dots, T\}$, and $\ell_g^{\text{home}}(t) = 1$ for all $g \in \mathcal{G}$ and $t \in \{0, \dots, T-1\}$

Benchmark policies for testing. We compare the testing decisions of ROLD against a proportional testing policy, which allocates testing proportionally to each age group's census population, up to capacity.

6. Experimental and Optimization Setup

We run experiments over a range of values for the parameters of our model. Table S7 summarizes the values we use for each parameter in our experimental setup, as well as the details of the optimization setup. The "Baseline Value" column reports the values of the parameters used in our baseline setting, as reported in results in the main paper and the SI, unless specified otherwise. We also run sensitivity analysis on the calibrated parameters: we run experiments over a range for each of the fitted parameters, as reported in Table S8.

We make some further comments on ICU and hospital capacity, as well as testing capacity. For the ICU capacity, we do a sensitivity analysis using values that range up to 3200 beds. We note that the ICU capacity in Île-de-France started at 1,200 beds (21) and was enhanced during the crisis, with a reported peak ICU utilization during the spring of 2020 of 2,668 beds, on April 8 (22). In our experiments we set the capacity for general ward beds to infinity. When we run experiments with viral testing, we consider a baseline capacity of 60,000 tests daily for Île-de-France (23, reported on May 11, 2020).

We run experiments for a range of values for the cost of death χ, accounted as multiples of the GPD per capita of France. For the GDP per capita of France, we use the figure for 2019, converting US dollars to euros using the exchange rate on June 17, 2020.

7. Additional Results

A. Benefits of Dual Targeting. We consider four versions of ROLD that differ in the level of targeting of confinements allowed: no targeting whatsoever (NO-TARGET), targeting age groups only (AGE), targeting activities only (ACT), or targeting both (AGE-ACT). To obtain each of the four variants (NO-TARGET, AGE, ACT, AGE-ACT), we run constrained versions of the optimization problem.

[‡‡] For reference, we provide here the French rule. For a geographic "department" in France to be declared under maximum alert, three criteria need to be satisfied. First criterion: the incidence rate, i.e. the number of positive cases per 100,000 inhabitants over seven days, must be greater than 250 per 100,000 inhabitants. Second criterion: the incidence rate among those over 65, the most vulnerable population, must exceed 100 cases per 100,000 inhabitants. Third criterion: more than 30% of resuscitation beds must be occupied by patients with COVID-19.

[§§] For the incidence rate, we count new cases using the Infectious state (and not the Exposed state) of the SEIR model; this is aligned with the assumption that a positive case can be detected through a viral test if and only if the individual is in the Infectious state.

Table S7. Parameter values for experimental and optimization setup

Parameter Description	Notation	Values in Experiments	Baseline Value
Cost of death	χ	121 values in $[0, 990]\times$ GDP per capita of France	50 $\times$
GDP per capita of France		€37199.03	
ICU capacity	K^{ICU}	2000, 2300, 2600, 2900, 3200	2900
Hospital capacity	K^{H}	∞	
Viral testing capacity	K^{Vtest}	243 values in the range of 0 to 1.2 million daily	0
Sensitivity of econ. value on confinement	$\nu^{\text{other activities}}$	0, 0.1, 0.2	0.1
Discount rate	r	0.03	
Fraction going to school	f_g	1 for $g =$ 0-9 y.o. 0.907 for $g =$ 10-19 y.o.	
Mult. factor for value of schooling	θ	0.5, 1, 5	0.5
Starting time for optimization		October 21 2020	
Optimization horizon	T	$\{90, 180, 360\}$ days	90 days
Frequency of confinement decisions		14 days	
Frequency of testing decisions		7 days	

$$\text{NO-TARGET:}\quad \ell_g^a(t) = \ell_{g'}^{a'}(t)\,\forall g, g' \in \mathcal{G},\ \forall a, a' \in \mathcal{A},\ \forall t \tag{72}$$

$$\text{AGE:}\quad \ell_g^a(t) = \ell_g^{a'}(t)\,\forall g \in \mathcal{G}, \forall a, a' \in \mathcal{A},\ \forall t \tag{73}$$

$$\text{ACT:}\quad \ell_g^a(t) = \ell_{g'}^{a}(t)\,\forall g, g' \in \mathcal{G}, \forall a \in \mathcal{A},\ \forall t \tag{74}$$

We impose no additional constraints for AGE-ACT. All four variants are allowed to change the confinement policy through time. We assume no testing for all four variants. For each variant, ROLD gets initialized at the solution of the gradient descent heuristic with the corresponding constraints. Figure S3a plots the total losses as a function of the cost of death χ for each of the four variants.

Figure S3b compares the total losses of ROLD with dual targeting (AGE-ACT) against the total losses of the benchmark policies, as a function of the cost of death χ.

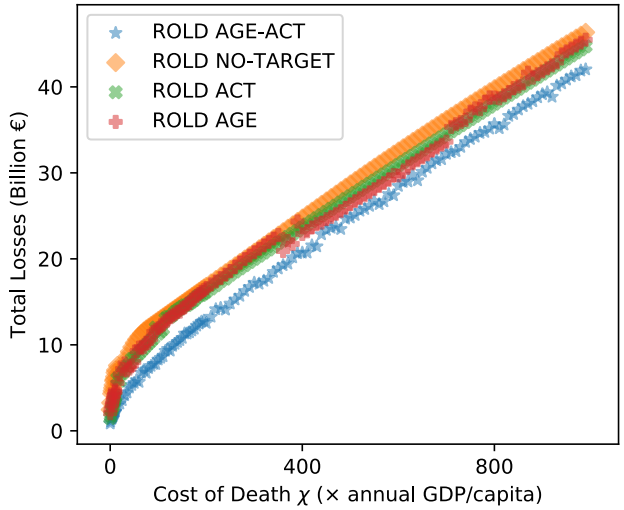

(a) Comparison between four ROLD policies with different levels of targeting

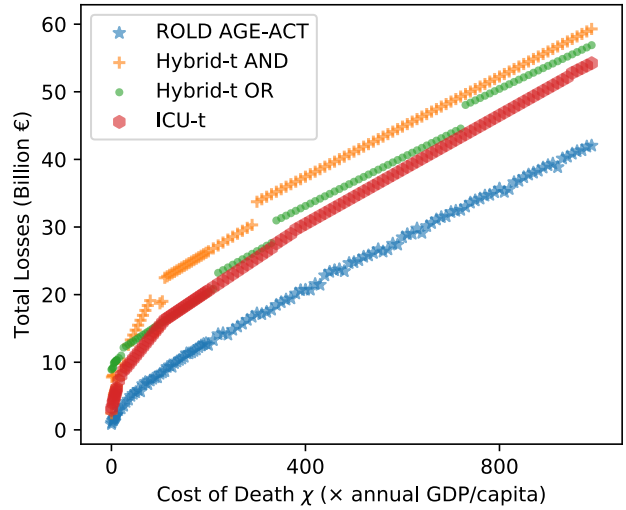

(b) Comparison between ROLD with dual targeting and benchmark policies

Fig. S3. Total losses generated by each policy at different values of the cost of death χ. **(a)** The impact of targeting confinements by age group and activity. The figure depicts the performance of four versions of ROLD that differ in the level of targeting allowed: no targeting whatsoever (NO-TARGET), targeting age-groups only (AGE), targeting activities only (ACT), or targeting both (AGE-ACT). The performance is calculated for different values of the parameter χ corresponding to the cost of death. A total of 128 values of χ are tested, ranging from 0 to 990 times the annual GDP per capita. **(b)** Comparison of the ROLD policy that targets age groups and activities with the benchmark policies. A total of 128 values of χ are tested, ranging from 0 to 990 times the annual GDP per capita. In both panels, all policies implement zero viral tests, so only confinement decisions are optimized and compared.

AGE-ACT not only Pareto-dominates AGE, ACT, and NO-TARGET, but also leads to super-additive improvements in almost all cases: for the same number of deaths, AGE-ACT reduces economic losses (with respect to NO-TARGET) by more than AGE and ACT *added together*. Furthermore, in almost all cases, for a given χ, AGE-ACT reduces total losses (with respect to NO-TARGET) by more that AGE and ACT *added together*. Figure S4 visualizes these comparisons.

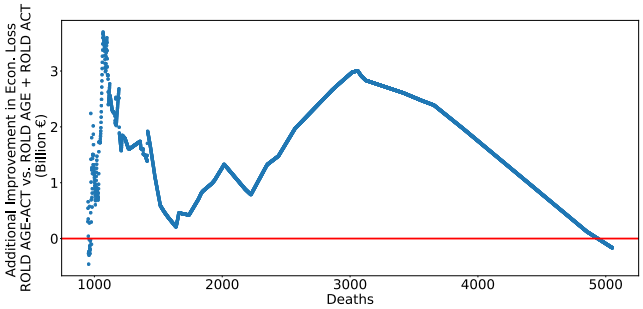

(a) The difference between the improvement in economic losses from ROLD AGE-ACT and the sum of the improvements from ROLD AGE and ROLD ACT, for the same number of deaths

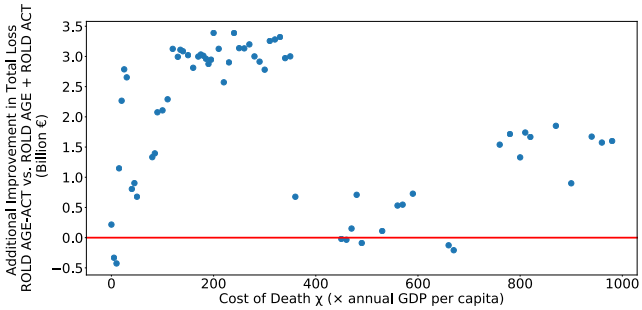

(b) The difference between the improvement in total losses from ROLD AGE-ACT and the sum of the improvements from ROLD AGE and ROLD ACT, for the same value of χ

Fig. S4. The super-additivity of ROLD AGE-ACT. The figures compare the improvement from AGE-ACT with the sum of the improvements from AGE and ACT. All improvements are with respect to NO-TARGET. Panel **(a)** compares the improvements in economic losses for the same number of deaths. Panel **(b)** compares the improvements in total losses for the same cost of death χ.

B. Robustness Analyses. We analyze additional problem instances by changing the value of each of 13 estimated parameters within a sensitivity range, as shown in Table S8. For each parameter, we sample 40 values uniformly at random from its specified sensitivity range. In each problem instance, one parameter is changed from its estimated value, for a total of $13 \times 40 = 520$ problem instances.

Table S8. Robustness analysis: parameters and sensitivity analysis ranges

Parameter	Estimated Value	Sensitivity Range
Multiplier m_β	0.60	0.45-0.75
Multiplier $m_{p_g^{ICU}}$	1.66	1.51-1.81
Multiplier $m_{p_g^D}$	1.13	0.98-1.28
Multiplier m_{λ_H}	1.50	1.35-1.65
Multiplier $m_{\lambda_{ICU}}$	2.00	1.85-2.15
Weight for activities α_{other}	0.96	0.76-1.0
School activity ℓ^{school} Mar 17 - May 10 2020	0.0	0.0-0.2
School activity ℓ^{school} May 11 - Jun 30 2020	0.0	0.0-0.2
School activity ℓ^{school} Jul 1 - Aug 31 2020	0.2	0.0-0.4
School activity ℓ^{school} Sep 1 - Oct 21 2020	1.0	0.8-1.0
Social mixing $\alpha_1 = \alpha_2$	0.39	0.19-0.59
Work activity $\ell^{\text{work}}_{\text{April}}$ April 2020	0.213	0.113-0.313
Economic activity level April 2020	58.51%	48.51%-68.51%

Figure S5 shows robustness results for six values of the economic cost of death χ: $[0, 15, 25, 50, 100, 150]\times$ the annual GDP per capita in France. The shown boxplots summarize results over the 520 problem instances, for each value of χ. These results reinforce our findings on the gains of dual targeting, as well as the observation that dual targeting unlocks complementarities which may not be available under targeting age groups or activities separately.

C. Testing and Quarantining.

Value of targeting testing. We next assess the value that can be generated by targeting testing based on age groups, with subsequent quarantining of positive cases. We compare the performance of two testing policies, while keeping all activities open, and varying the testing capacity: targeting testing to different age groups using ROLD, versus allocating testing proportionally to the census population of each age group. Figure S6 reports the performance of these two testing policies. Figure S7 reports the difference in deaths, economic losses, and total losses between targeted testing and proportional testing.

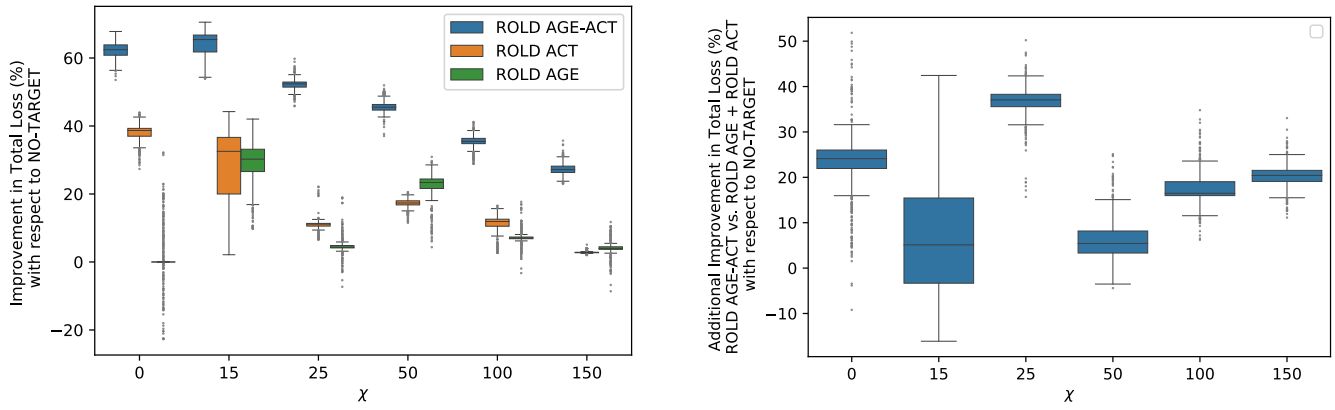

(a) Comparison between ROLD policies with different levels of targeting in terms of total losses

(b) The difference between the improvement in total losses from ROLD AGE-ACT and the sum of the improvements from ROLD AGE and ROLD ACT

Fig. S5. Robustness analyses showing the superiority of ROLD AGE-ACT over ROLD policies with less granular targeting, and the super-additive improvements of ROLD AGE-ACT over the sum of the improvements of ROLD AGE and ROLD ACT, for different values for the cost of death χ, over a wide set of problem instances. All improvements are with respect to ROLD NO-TARGET. For each value of χ, the boxplots summarize results over 520 different problem instances.

At a cost of death $\chi = 50\times$ the annual GDP per capita, the value of targeted testing with a capacity of 60,000 daily tests is an improvement by EUR 2.9B (7.8%) in terms of total losses (1488 or 7.8% less deaths; EUR 0.1B or 8.3% less economic loss). The improvement in deaths and in the total loss due to targeting testing would be maximized (3210 or 26.7% less deaths, EUR 6.3B or 26.9% less total losses) with a daily testing capacity of around 238,700 tests, corresponding to more than 1.9% of the Île-de-France population. This is more than four times larger than the real testing capacity in Île-de-France (24). Figures S6 and S7 report results for a wide range of testing capacities.

We note that for targeted testing, the improvement in total losses from higher testing capacity starts slowing down around 230,000 daily tests, as shown in Figure S6; that's also the testing capacity at which the difference between targeted and proportional testing peaks, as shown in Figure S7. At that testing capacity, there is no more infection in the community. There are still benefits from a larger testing capacity, as the point in time at which the infection is killed would come sooner, however the benefits are diminishing.

For proportional testing, the improvement in total losses from higher testing capacity starts slowing down around 400,000 daily tests, as shown in Figure S6. First, it makes sense that the proportional testing curve would start slowing down for a larger testing capacity than the targeted testing curve. Second, that's also the testing capacity at which the drop in the difference between targeted and proportional testing tapers off, as shown in Figure S7. At large testing capacities, for which the infection is killed both for targeted testing and for proportional testing, targeted testing is still able to extract more value from abundant testing capacity than proportional testing.

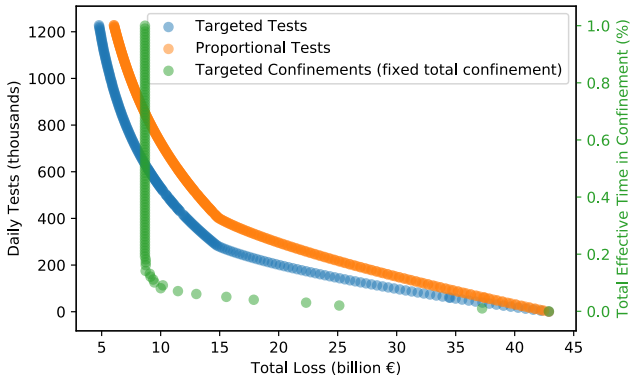

Fig. S6. Comparison of the performance of testing-and-quarantining-only interventions (with ROLD targeted testing vs. with proportional testing), and confinement-only interventions (with ROLD confinement decisions), for different testing capacities, and for different levels of effective confinement time. The experiments are run using $\chi = 50\times$ the annual GDP per capita and an ICU capacity of 2900 beds. The number of values tested is 243 and 541 values of testing capacity for the targeted and proportional testing policies respectively, and 99 values of the percentage of total effective time in confinement for the confinements-only policy. The range for the number of tests is from 0 tests to 10% of the population of Île-de-France daily. (The population of Île-de-France is 12278210.)

Value of confinement versus value of testing. We use our framework to compare (a) an intervention that only uses confinement and no testing to (b) an intervention that only uses testing and subsequent quarantining while keeping all

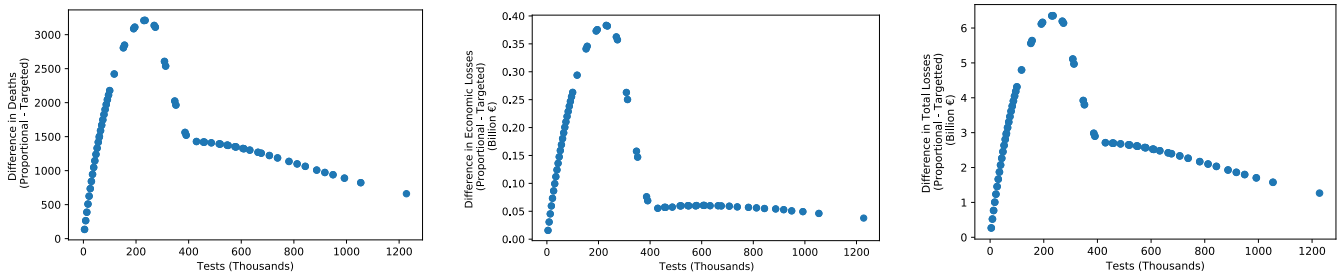

(a) Difference in deaths **(b)** Difference in econ. losses **(c)** Difference in total losses

Fig. S7. Comparison of targeted testing using ROLD versus proportional allocation of testing to age groups, in terms of deaths, economic losses, and total losses. The experiments are run while keeping all activities open, and using $\chi = 50\times$ the annual GDP per capita and an ICU capacity of 2900 beds. The number of data points in each plot is 100.

activities open. For the confinement-only intervention, we constrain ROLD to use a specified "budget" of total effective time in confinement. For the testing-only intervention, we constrain ROLD to use a specified testing capacity. We run experiments where we vary the effective confinement time for the confinement-only policy, as well as the testing capacity for the testing-only policy, and report the performance of ROLD. In both cases we allow ROLD to target its decisions to different age groups. Note that for the confinement-only policy, we do not allow ROLD to target activities, so as to keep the comparison between the confinement intervention and the testing intervention fair: both can only target age groups.

For the testing-and-quarantining-only intervention, we run ROLD assuming no confinement, for a range of testing capacities, allowing for optimized allocation of tests to different age groups. We run ROLD for a wide range of daily testing capacity K^{Vtest}, from no testing up to testing 10% of the population daily.

We also run ROLD assuming no testing, allowing for optimized targeting of confinements to different age groups, subject to a constraint of "total confinement time". To model total confinement time, we add the following constraint to ROLD:

$$(\text{“Total confinement time”}) \quad \sum_{t=0}^{T-1} \sum_{g} \ell_g^a(t) N_g(0) \geq (1 - p_{lock}) \cdot T \sum_{g} N_g(0), \; \forall a \in \mathcal{A}. \qquad [75]$$

where p_{lock} is the percentage of time spent in confinement. We provide numerical results for a range of values of $p_{\text{lock}} \in [0, 1]$.

Figure S6 reports the comparison in the performance (in terms of the total loss objective) between the confinement-only and testing-only interventions, for a cost of death $\chi = 50\times$ the annual GDP per capita. The plot shows that all the benefits from confinement can be accomplished with a total effective confinement time that corresponds to 19.2% of the total horizon window, or 17.3 days of effective confinement time within a 90-day horizon, as long as the confinement is optimally targeted to different age groups. To attain the same total loss (EUR 8.6B) without confinement, an extraordinary amount of (targeted) testing would be needed: more than 632,500 tests would need to be administered daily, which corresponds to more than 5.1% of the Île-de-France population. The plot shows that testing without confinement can in theory reach performance levels that are unattainable by confinement without testing, however this only happens when the daily testing capacity exceeds 5.1% of the Île-de-France population. This is about eleven times larger than the real testing capacity in Île-de-France (24).

D. The Proposed Policy.

Comments on the simplicity of the policy. We make some other observations about the simplicity of the optimized ROLD policy: it relies on only a few different values for activity levels, and it relaxes confinements as we get closer to the end of the optimization horizon. Furthermore, at any fixed time, as the cost of death χ increases, there is generally a nestedness in the confinement decisions: if a certain age group is confined in a particular activity at a value of χ, that group will face a confinement that is at least as strict at a larger value of χ.

The proposed policy with testing. We allow ROLD to also optimize for testing decisions. Figure S8a shows the ROLD confinement and testing policy, for a cost of death $\chi = 50\times$ and a 90-day optimization horizon, under a daily testing capacity of 60,000 tests. ROLD cleverly mixes confinement decisions and targeted testing. For example, it distributes the testing capacity to the 40-69 age groups, which are also kept open in work. Intuitively, this is because the work activity generates more social contacts (and thus new infections) than other activities. Moreover, of all the age groups (10-69) that are active in work, although the 40-69s do not have the largest number of social contacts, they do have the most contacts with the high risk 70+ age groups, so most of the testing capacity is focused on them. In other words, the tests are used to detect and quarantine the individuals which are most likely to create new infections in high risk individuals through their social contacts in open activities.

Adding in testing can substantially change the confinement policy by permitting less total confinement of the population. Figure S8a show that the testing reduces the overall confinement in work, leisure and other. Also, although testing is mostly concentrated in the 40-49 y.o. and 50-59 y.o., it also serves to reduce confinement in some groups which are not tested, such as the confinement in work for the 10-19 y.o. group.

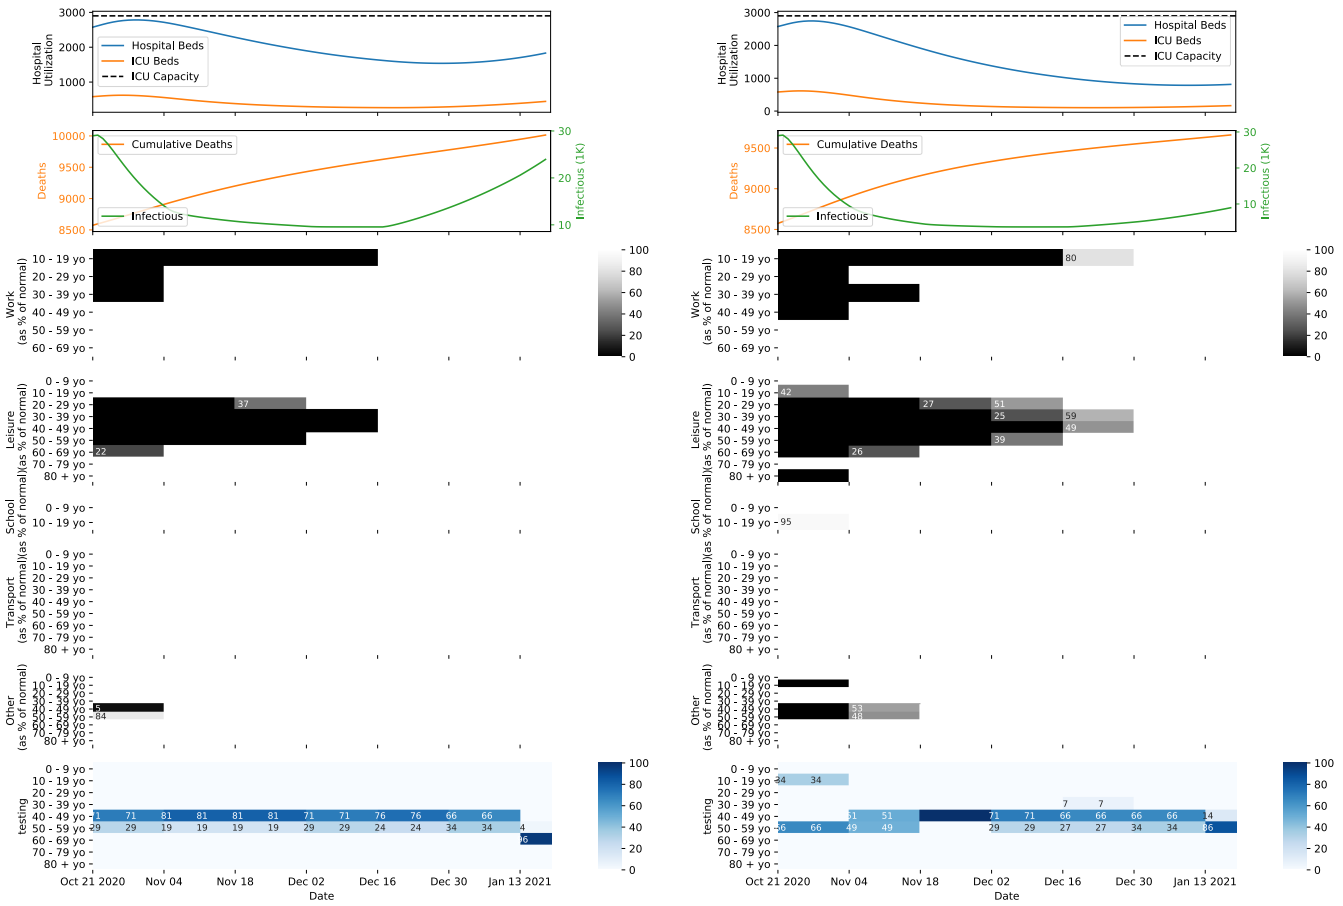

(a) $\chi = 50\times$ the annual GDP per capita

(b) $\chi = 150\times$ the annual GDP per capita

Fig. S8. The optimized ROLD confinement and testing decisions and their impact on hospitalizations, infections, and deaths, for different values of the cost of death, for a 90-day optimization horizon, with a testing capacity of 60,000 daily tests.

Dependence of the proposed policy on problem parameters. ROLD decisions generally depend quite strongly on the specific problem parameters, such as the cost of death χ. We make a few other observations on how the ROLD policies depend on critical problem parameters.

Cost of death χ: At any fixed time, as χ increases there is *generally* a nestedness in the confinement decisions followed by ROLD: if a certain age group is confined in a particular activity at a fixed time for a value of χ, that group will face a confinement that is at least as strict in that activity, at that time, for a larger value of χ. As a comparison to the policy for $\chi = 50\times$, Figure S8b reports the decisions produced by ROLD for a cost of death $\chi = 150\times$ the annual GDP per capita, using an optimization horizon of $T = 90$ days. This property can be very useful for decision making, especially in cases where a policy maker cannot or prefers not to assign a specific value for χ, but instead is comfortable providing a range for χ. It is important to note that this property of nested confinements does not immediately follow, e.g., as χ increases and a policy decides to impose more strict confinements on specific age groups, it could relax the confinements of other groups.

Time horizon T: In order to understand the effect of the time horizon, we also run experiments for an optimization horizon of $T = 360$ days. These reinforce the behavior observed for $T = 90$ days. Figure S9 reports the decisions produced by ROLD for an optimization horizon of $T = 360$ days, using a cost of death $\chi = 50\times$ the annual GDP per capita. Comparing the policy for $T = 360$ days with the policy for $T = 90$ days, we see that for $T = 360$ days the confinement policy is not monotonic over time. However, the policy still relaxes confinements as we get close to the end of the horizon. Furthermore, the confinement for a fixed (age group, activity) pair and time t will generally be at least as strict as the optimization horizon increases. Exceptions to this general "rule" can occur at early times of a long optimization horizon.

E. Additional Details for the Regression Trees for the ROLD Policy. To understand how the ROLD policies target confinements across different age groups and activities, we train an explainable machine learning model – a regression decision tree – to predict the optimal ROLD confinement in each activity as a function of several features.

To build a training set, we first create a larger set of problem instances by considering different values for the cost of death (χ), the relative importance of each activity in generating economic value (captured by the parameters θ, ν^{work}, $\nu^{\text{other activities}}$), and the ICU capacity. All instances use a horizon of $T = 90$ days, adopted to reduce the computational burden of the

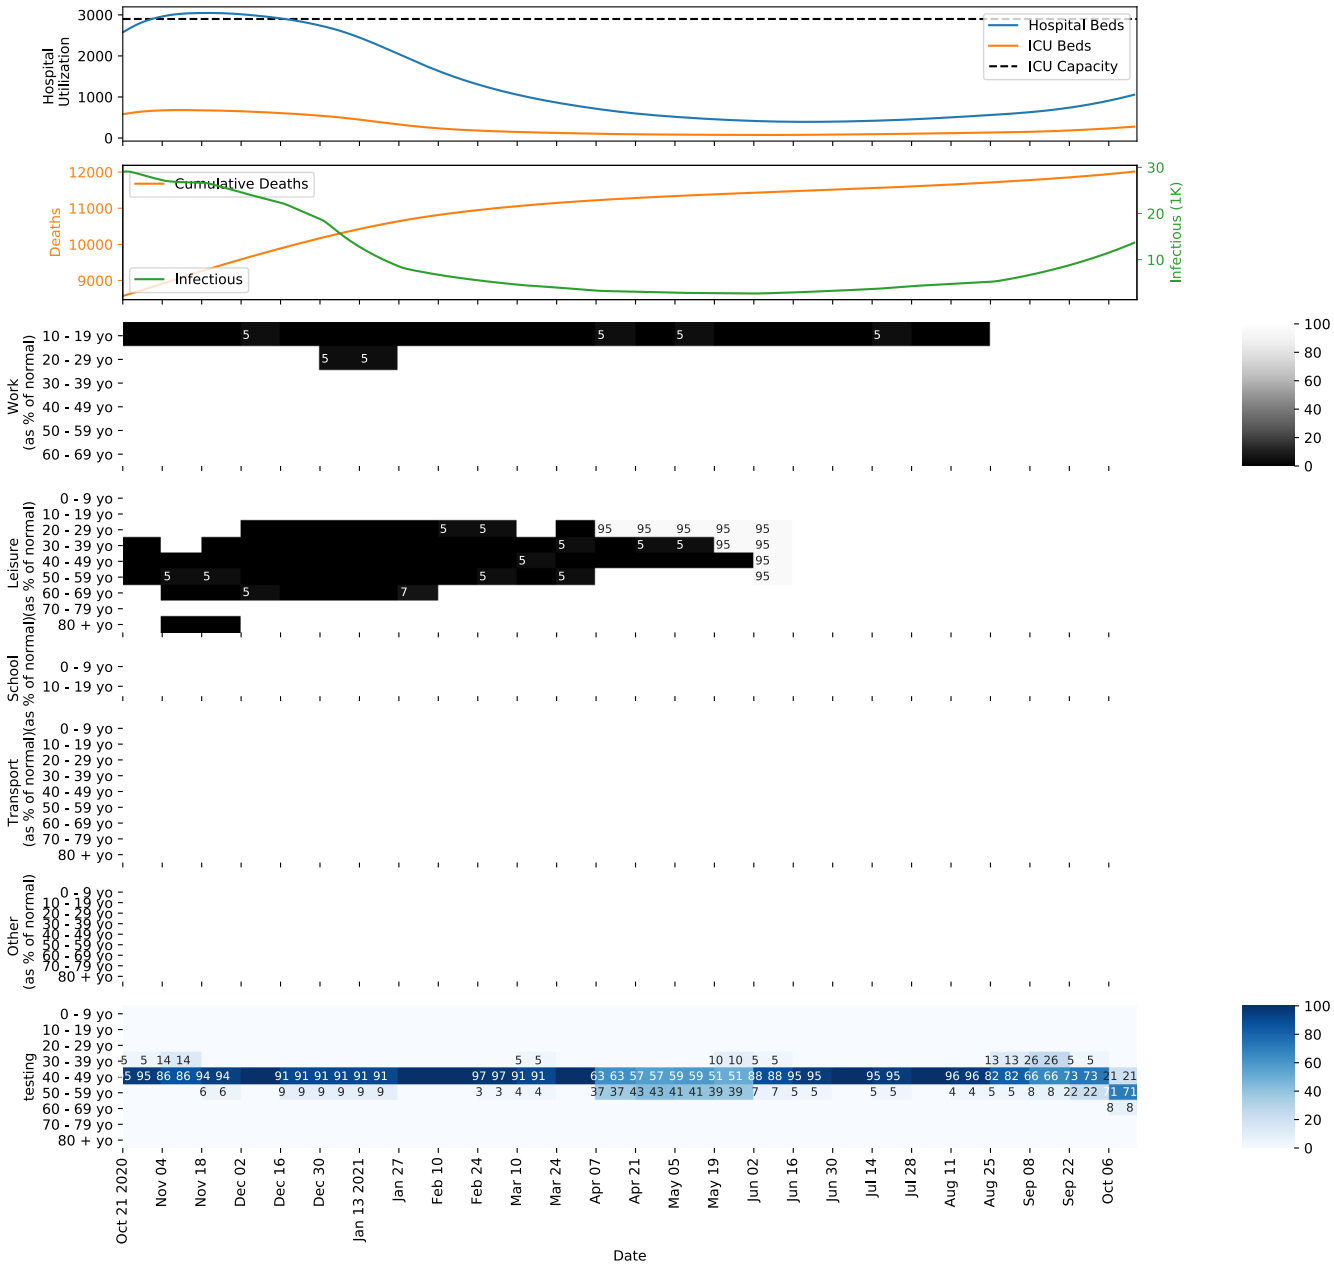

Fig. S9. The optimized ROLD confinement and testing decisions and their impact on hospitalizations, infections, and deaths for a 360-day optimization horizon, in a problem instance with cost of death $\chi = 50\times$ and a testing capacity of 60,000 daily tests.

experiments.

For each problem instance, we compute the optimal ROLD activity levels, i.e., the decisions $\ell_g^a(t)$ for $g \in \mathcal{G}, a \in \mathcal{A}, 0 \leq t \leq T-1$, which we then use to simulate the SEIR dynamics and calculate all the corresponding values of the SEIR states. Lastly, we create a set of training set samples — with one sample for each time period t, each age group g, and each activity a relevant to that age group — where we include several direct and derived features based on the parameter values characterizing the instance and the induced SEIR states, and a target corresponding to the optimal ROLD decision $\ell_g^a(t)$. Specifically, we include the following as features:

- the activity name a, the age group g, the parameters χ, ν^{other}, ν^{work}, v_g^{life}, the ICU capacity, the time t
- "econ-gradient(g)": the gradient of the economic value with respect to the activity level $\ell_g^a(t)$
- "I-perc(t)": the percentage of the population estimated by the model to be in the infectious state at time t, i.e., $\sum_{g\in\mathcal{G}} I_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)]$

- "R-perc(t)": the percentage of the population estimated by the model to be in the recovered state at time t, i.e., $\sum_{g\in\mathcal{G}} R_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)]$

- "H-perc(t)": the percentage of the population in the hospital due to COVID-19 at time t, i.e., $\sum_{g\in\mathcal{G}} H_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)]$

- "ICU-utilization(t)": percentage of ICU beds occupied at time t

- "contacts-received(g)": the number of social contacts received by group g from other groups in activity a, i.e., $\sum_{h\in\mathcal{G}} C_{g,h}^a$

- "contacts-received-into-deaths(g)": contacts received by group g multiplied by their probability of dying conditional on being infected, i.e., $\sum_{h\in\mathcal{G}} C_{g,h}^a \cdot p_g^D$;

- "infections-received(g)": contacts received by group g multiplied by the probability of those contacts coming from infected people, i.e., $\sum_{h\in\mathcal{G}} C_{g,h}^a \cdot I_h(t)/[\sum_{g\in\mathcal{G}} N_h(t) + R_h^q(t)]$;

- "deaths-received(g)": "infections-received(g)" multiplied by the group g's probability of dying conditional on being infected by those contacts, i.e., $\sum_{h\in\mathcal{G}} C_{g,h}^a \cdot I_h(t)/[\sum_{g\in\mathcal{G}} N_h(t) + R_h^q(t)] \cdot p_g^D$;

- "deaths-econ-received(g)": "deaths-received(g)" multiplied by the economic value of the losses of those deaths $\chi + v_g^{\text{life}}$, i.e., $\sum_{h\in\mathcal{G}} C_{g,h}^a \cdot I_h(t)/[\sum_{g\in\mathcal{G}} N_h(t) + R_h^q(t)] \cdot p_g^D \cdot [\chi + v_g^{\text{life}}]$;

- "contacts-given(g)": number of social contacts given to other groups by group g in the specific activity a, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a$;

- "contacts-given-into-deaths(g)": the sum of terms of the form $C_{h,g}$ multiplied by group h's probability of dying conditional on being infected, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a \cdot p_h^D$;

- "infections-given(g)": the number of "contacts-given(g)" multiplied by the fraction of the population in age group g being infected, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a \cdot I_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)]$;

- "deaths-given(g)": infections-given(g) multiplied by the probability of dying conditional on being infected for each age group, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a \cdot I_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)] \cdot p_h^D$;

- "deaths-econ-given(g)": "deaths-given(g)" multiplied by the economic cost of those deaths $\chi + v_h^{\text{life}}$, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a \cdot I_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)] \cdot p_h^D \cdot [\chi + v_h^{\text{life}}]$;

- "death-prob(g)": group g's probability of dying conditional on being infected p_g^D;

- "econ-contacts-ratio": the ratio of "econ-gradient(g)" and "contacts-given(g)".

This results in a training set of 27,720 unique samples for each of the work, leisure and other activities.

Using this data, we then use the SCIKIT-LEARN Python library to train separate, activity-specific regression decision trees to predict the optimal ROLD confinement decisions in the respective activity as a function of the considered features. We train trees with different maximum depth ranging from three to six, and using the traditional mean-squared-error (MSE) criterion as a goodness of fit metric. The trees reported in the main body of the paper were all trained with a depth of three and using the econ-to-contacts-ratio, the cost of death χ, the age, and time as features.

Other than confirming that the econ-to-contacts ratio is the most salient feature when targeting confinements, the trees also confirm that time and the cost of death χ play an important role: the optimal ROLD policy tends to enforce stricter confinements in each activity in earlier periods and subsequently relax these through time, and the confinements become stricter for higher values of χ.

In addition, we also report in Figure S10 trees of depth three trained using *all* the features described above. Note that the "econ-to-contacts-ratio" continues to be the most relevant feature, used as a split variable in the root node of each activity tree, and also subsequently used for splits in several sub-trees.

F. Can Dual Targeting Reduce Time in Confinement for Each Age Group?. We calculate the fraction of time each age group spends in confinement under each ROLD policy, averaged over the activities relevant to that age-group. Specifically, these are

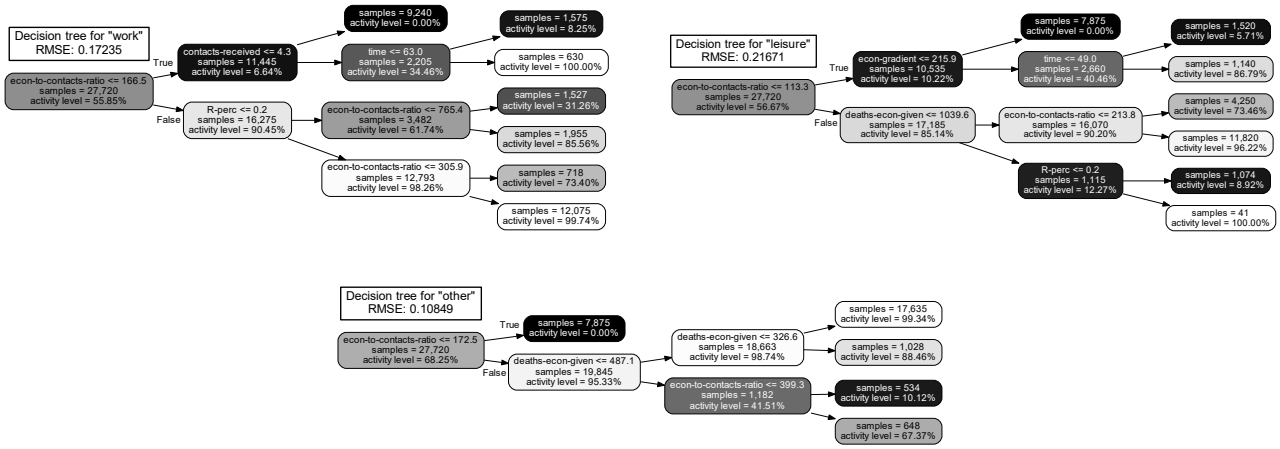

Fig. S10. Decision trees for work, leisure and other activities trained on the optimal ROLD policy for different values of problem parameters. Darker shading in a node corresponds to lower activity level decisions for samples in that node. Each tree is trained with a total of 27,720 unique samples.

defined as:

$$1 - \frac{1}{4T} \sum_{t=1}^{T} \sum_{a \in \{\text{school,leisure,other,transport}\}} \ell_g^a(t), \qquad \text{if } g = \text{0-9 y.o.} \tag{76}$$

$$1 - \frac{1}{5T} \sum_{t=1}^{T} \sum_{a \in \{\text{work,school,leisure,other,transport}\}} \ell_g^a(t), \qquad \text{if } g = \text{10-19 y.o.} \tag{77}$$

$$1 - \frac{1}{4T} \sum_{t=1}^{T} \sum_{a \in \{\text{work,leisure,other,transport}\}} \ell_g^a(t), \qquad \text{if } g = \text{20-29, 30-39, 40-49, 50-59, 60-69 y.o.} \tag{78}$$

$$1 - \frac{1}{3T} \sum_{t=1}^{T} \sum_{a \in \{\text{leisure,other,transport}\}} \ell_g^a(t), \qquad \text{otherwise.} \tag{79}$$

Further observations. We make some further comments about the results shown in Figure 3 in the main paper. It is worth noting that although ROLD AGE-ACT reduces confinements for every age group compared to less targeted policies, it does not do so uniformly, and it can lead to an even larger discrepancy in the confinements faced by different age groups: those aged 10-59 are generally more confined than those aged 0-9 or 60 and above. Such disparate treatment is non-existent in the NO-TARGET policy by definition, and is less apparent in the AGE and ACT policies. Interestingly, while AGE generally exhibits a similar discrimination profile to AGE-ACT, ACT exhibits an almost opposite profile, with more confinement applied to those aged 70 and above. This is explainable by considering again the "econ-to-contacts" ratio: for instance, since leisure is responsible for the smallest (population-weighted average) econ-to-contacts ratio among all activities, it is the first activity that ACT targets for confinement, and this disproportionately impacts the groups aged 70 and above, who are active in fewer activities overall than other age groups and have leisure as one of their activities.

G. Limited Disparity Requirements. We impose limited disparity constraints on confinement decisions across pairs of age groups, for a given activity. For work, we impose constraints for age groups 10-19 y.o., . . . , 60-69 y.o. For transport, leisure and other, we impose constraints for all age groups. No constraints are imposed for home and for school. Given a disparity tolerance $\Delta \geq 0$, we formalize the constraints as follows:

$$\begin{aligned}
&\ell_g^{\text{work}}(t) \leq \ell_{g'}^{\text{work}}(t) + \Delta,\ \forall\, t \in \{0, \ldots, T-1\}, \forall\, (g, g') : g \neq g',\ g, g' \notin \{\text{0-9 y.o., 70-79 y.o., 80+ y.o.}\} \\
&\ell_g^{\text{work}}(t) \geq \ell_{g'}^{\text{work}}(t) - \Delta,\ \forall\, t \in \{0, \ldots, T-1\}, \forall\, (g, g') : g \neq g',\ g, g' \notin \{\text{0-9 y.o., 70-79 y.o., 80+ y.o.}\} \\
&\ell_g^{a}(t) \leq \ell_{g'}^{a}(t) + \Delta,\ \forall\, t \in \{0, \ldots, T-1\}, \forall\, a \in \{\text{leisure, transport, other}\},\ \forall (g, g') : g \neq g' \\
&\ell_g^{a}(t) \geq \ell_{g'}^{a}(t) - \Delta,\ \forall\, t \in \{0, \ldots, T-1\}, \forall\, a \in \{\text{leisure, transport, other}\},\ \forall (g, g') : g \neq g'.
\end{aligned}$$

Smaller values of Δ imply more stringent non-disparity constraints. We provide numerical results for a range of values of Δ in $[0, 1]$.

8. Further Calculations and Algorithmic Details for ROLD

A. Capping the Turn-away Variables. Because our controls $\boldsymbol{u}_t$ include decisions on how many people are turned away for each group from the hospital and ICU, we need to guarantee that when calling $F_t(\boldsymbol{X}_t, \boldsymbol{u}_t)$, the turn-away decisions are *feasible*, i.e., that they satisfy the conditions in Eq. (25), Eq. (27). In order to ensure this, we perform a modification of the controls that accomplishes two goals:

1. Cap all turn-away decisions to the incoming flow to the hospital (ICU).

2. Increase the turned away proportionally to the excess contribution of each group if necessary to respect the capacity.

To accomplish the first goal, we simply cap all the controls by the inflow, $\mu\left(p_g^H \cdot I_g(t) + \frac{p_g^H}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right)$ (and equivalently for ICU). To accomplish the second goal, we define new controls $B_g^H(t)$ and $B_g^{ICU}(t)$, based on the given controls $\tilde{B}_g^H(t)$ and $\tilde{B}_g^{ICU}(t)$, as:[¶¶]

$$B_g^H(t) = \tilde{B}_g^H(t) + \frac{\mu\left(p_g^H \cdot I_g(t) + \frac{p_g^H}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) - \tilde{B}_g^H(t)}{\sum_h \mu\left(p_h^H \cdot I_h(t) + \frac{p_h^H}{p_h^H + p_h^{ICU}} \cdot I_{ss,h}^{\mathrm{q}}(t)\right) - \tilde{B}_h^H(t)} \tag{80}$$

$$\cdot \left(\left(\sum_h \left(1 - \lambda_h^{H,R} - \lambda_h^{H,D}\right) H_h(t) + \mu\left(p_h^H \cdot I_h(t) + \frac{p_h^H}{p_h^H + p_h^{ICU}} \cdot I_{ss,h}^{\mathrm{q}}(t)\right) - \tilde{B}_h^H(t)\right) - K^{\mathrm{H}}(t)\right)^+ \tag{81}$$

$$B_g^{ICU}(t) = \tilde{B}_g^{ICU}(t) + \frac{\mu\left(p_g^{ICU} \cdot I_g(t) + \frac{p_g^{ICU}}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) - \tilde{B}_g^{ICU}(t)}{\sum_h \mu\left(p_h^{ICU} \cdot I_h(t) + \frac{p_h^{ICU}}{p_h^H + p_h^{ICU}} \cdot I_{ss,h}^{\mathrm{q}}(t)\right) - \tilde{B}_h^{ICU}(t)} \tag{82}$$

$$\cdot \left(\left(\sum_h \left(1 - \lambda_h^{ICU,R} - \lambda_h^{ICU,D}\right) ICU_h(t) + \mu\left(p_h^{ICU} \cdot I_h(t) + \frac{p_h^{ICU}}{p_h^H + p_h^{ICU}} \cdot I_{ss,h}^{\mathrm{q}}(t)\right) - \tilde{B}_h^{ICU}(t)\right) - K^{\mathrm{ICU}}(t)\right)^+ \tag{83}$$

[¶¶] To simplify the notation, we set

$$\begin{aligned}
\lambda_g^{H,R} &= p_g^R \cdot \lambda_g^H \\
\lambda_g^{H,D} &= p_g^D \cdot \lambda_g^H \\
\lambda_g^{ICU,R} &= p_g^R \cdot \lambda_g^{ICU} \\
\lambda_g^{ICU,D} &= p_g^D \cdot \lambda_g^{ICU}.
\end{aligned}$$

B. Deriving the Jacobians. Here, we give the exact formulas for the Jacobians we use for the ROLD procedure.

$$f^{N,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := -N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)} - \mu \cdot (p_g^H + p_g^{ICU}) \cdot I_g(t) \tag{84a}$$

$$\frac{\partial f^{N,g}}{\partial N_g} = N_g^{\text{Vtest}}(t) \frac{I_g(t)}{N_g^2(t)} \tag{84b}$$

$$\frac{\partial f^{N,g}}{\partial I_g} = -\frac{N_g^{\text{Vtest}}(t)}{N_g(t)} - \mu \cdot (p_g^H + p_g^{ICU}) \tag{84c}$$

$$\frac{\partial f^{N,g}}{\partial N_g^{\text{Vtest}}(t)} = -\frac{I_g(t)}{N_g(t)} \tag{84d}$$

$$\tag{84e}$$

$$f^{S,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := -\beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell) \frac{I_h(t)}{N_h(t) + R_h^q(t)} \right) \tag{84f}$$

$$\frac{\partial f^{S,g}}{\partial N_h} = \beta(t) S_g(t) \cdot c_{g,h}(\ell) \frac{I_h(t)}{(N_h(t) + R_h^q(t))^2}, \forall\, h \tag{84g}$$

$$\frac{\partial f^{S,g}}{\partial S_g} = -\beta(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell) \frac{I_h(t)}{N_h(t) + R_h^q(t)} \right) \tag{84h}$$

$$\frac{\partial f^{S,g}}{\partial R_h^q} = \beta(t) S_g(t) \cdot \left(c_{g,h}(\ell) \frac{I_h(t)}{(N_h(t) + R_h^q(t))^2} \right), \forall\, h \tag{84i}$$

$$\frac{\partial f^{S,g}}{\partial I_h} = -\beta(t) S_g(t) \cdot \left(c_{g,h}(\ell) \frac{1}{N_h(t) + R_h^q(t)} \right), \forall\, h \tag{84j}$$

$$\frac{\partial f^{S,g}}{\partial \ell_{g,a}} = -\beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} \frac{\partial c_{g,h}(\ell)}{\partial \ell_{g,a}} \frac{I_h(t)}{N_h(t) + R_h^q(t)} \right), \forall\, a \tag{84k}$$

$$\frac{\partial f^{S,g}}{\partial \ell_{h,a}} = -\beta(t) S_g(t) \cdot \frac{\partial c_{g,h}(\ell)}{\partial \ell_{h,a}} \frac{I_h(t)}{N_h(t) + R_h^q(t)}, \forall\, h, \forall\, a \tag{84l}$$

$$\tag{84m}$$

$$f^{E,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := \beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell) \frac{I_h(t)}{N_h(t) + R_h^q(t)} \right) - \sigma E_g(t) \tag{84n}$$

$$\frac{\partial f^{E,g}}{\partial S_g} = -\frac{\partial f^{S,g}}{\partial S_g} \tag{84o}$$

$$\frac{\partial f^{E,g}}{\partial X_h} = -\frac{\partial f^{S,g}}{\partial X_h}, \forall\, X \in \{N, R^q, I\}, \forall\, h \tag{84p}$$

$$\frac{\partial f^{E,g}}{\partial \ell_{g,a}} = -\frac{\partial f^{S,g}}{\partial \ell_{g,a}}, \forall\, a \tag{84q}$$

$$\frac{\partial f^{E,g}}{\partial \ell_{h,a}} = -\frac{\partial f^{S,g}}{\partial \ell_{h,a}}, \forall h, \forall\, a \tag{84r}$$

$$\frac{\partial f^{E,g}}{\partial E_g} = -\sigma \tag{84s}$$

$$f^{I,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := \sigma E_g(t) - \mu \cdot I_g(t) - N_g^{\text{Vtest}}(t) \cdot \frac{I_g(t)}{N_g(t)} \tag{85a}$$

$$\frac{\partial f^{I,g}}{\partial E_g} = \sigma \tag{85b}$$

$$\frac{\partial f^{I,g}}{\partial I_g} = -\mu - \frac{N_g^{\text{Vtest}}(t)}{N_g(t)} \tag{85c}$$

$$\frac{\partial f^{I,g}}{\partial N_g} = N_g^{\text{Vtest}}(t) \frac{I_g(t)}{(N_g(t))^2} \tag{85d}$$

$$\frac{\partial f^{I,g}}{\partial N_g^{\text{Vtest}}(t)} = -\frac{I_g(t)}{N_g(t)} \tag{85e}$$

$$\tag{85f}$$

$$f^{R,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := \mu(1 - p_h^H - p_g^{ICU}) I_g(t) \tag{85g}$$

$$\frac{\partial f^{R,g}}{\partial I_g} = \mu(1 - p_g^H - p_g^{ICU}) \tag{85h}$$

$$\tag{85i}$$

$$f^{I,j,g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := p_{j,g} N_g^{\text{Vtest}}(t) \frac{I_g(t)}{N_g(t)} - \mu I_{j,g}^{\text{q}}(t), \quad \forall j \in as, ps, ms, ss \tag{85j}$$

$$\frac{\partial f^{I,j,g}}{\partial I_g} = p_{j,g} \frac{N_g^{\text{Vtest}}(t)}{N_g(t)} \tag{85k}$$

$$\frac{\partial f^{I,j,g}}{\partial N_g} = -p_{j,g} N_g^{\text{Vtest}}(t) \frac{I_g(t)}{(N_g(t))^2} \tag{85l}$$

$$\frac{\partial f^{I,j,g}}{\partial I_{j,g}^{\text{q}}} = -\mu \tag{85m}$$

$$\frac{\partial f^{I,j,g}}{\partial N_g^{\text{Vtest}}(t)} = p_{j,g} \frac{I_g(t)}{N_g(t)} \tag{85n}$$

$$\tag{85o}$$

$$f^{R^{\text{q}},g}(\boldsymbol{X}_t, \boldsymbol{u}_t) := \mu \cdot \sum_{j \in \{a,ps,ms\}} I_{j,g}^{\text{q}}(t) + \lambda_g^{H,R} \cdot H_g(t) + \lambda_g^{ICU,R} \cdot ICU_g(t) \tag{85p}$$

$$\frac{\partial f^{R^{\text{q}},g}}{\partial I_{j,g}^{\text{q}}} = \mu, \quad \forall j \tag{85q}$$

$$\frac{\partial f^{R^{\text{q}},g}}{\partial H_g} = \lambda_g^{H,R} \tag{85r}$$

$$\frac{\partial f^{R^{\text{q}},g}}{\partial ICU_{j,g}} = \lambda_g^{ICU,R} \tag{85s}$$

$$\tag{85t}$$

$$f^{H,g}\left(\boldsymbol{X}_t, \boldsymbol{u}_t\right) := -\left(\lambda_g^{H,R} + \lambda_g^{H,D}\right) H_g(t) + \mu \cdot \left(p_g^H \cdot I_g(t) + \frac{p_g^H}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) - B_g^H(t) \tag{86a}$$

$$\frac{\partial f^{H,g}}{\partial H_g} = -\left(\lambda_g^{H,R} + \lambda_g^{H,D}\right) \tag{86b}$$

$$\frac{\partial f^{H,g}}{\partial I_g} = \mu \cdot p_g^H \tag{86c}$$

$$\frac{\partial f^{H,g}}{\partial I_{ss,g}} = \mu \cdot \frac{p_g^H}{p_g^H + p_g^{ICU}} \tag{86d}$$

$$\frac{\partial f^{H,g}}{\partial B_g^H} = -1 \tag{86e}$$

$$\tag{86f}$$

$$f^{ICU,g}\left(\boldsymbol{X}_t, \boldsymbol{u}_t\right) := -\left(\lambda_g^{ICU,R} + \lambda_g^{ICU,D}\right) ICU_g(t) + \mu \cdot \left(p_g^{ICU} \cdot I_g(t) + \frac{p_g^{ICU}}{p_g^H + p_g^{ICU}} \cdot I_{ss,g}^{\mathrm{q}}(t)\right) - B_g^{ICU}(t) \tag{86g}$$

$$\frac{\partial f^{ICU,g}}{\partial ICU_g} = -\left(\lambda_g^{ICU,R} + \lambda_g^{ICU,D}\right) \tag{86h}$$

$$\frac{\partial f^{ICU,g}}{\partial I_g} = \mu \cdot p_g^{ICU} \tag{86i}$$

$$\frac{\partial f^{ICU,g}}{\partial I_{ss,g}^{q}} = \mu \cdot \frac{p_g^{ICU}}{p_g^H + p_g^{ICU}} \tag{86j}$$

$$\frac{\partial f^{ICU,g}}{\partial B_g^{ICU}} = -1 \tag{86k}$$

$$\tag{86l}$$

$$f^{D,g}\left(\boldsymbol{X}_t, \boldsymbol{u}_t\right) := \lambda_g^{H,D} \cdot H_g(t) + \lambda_g^{ICU,D} \cdot ICU_g(t) + B_g^H(t) + B_g^{ICU}(t) \tag{86m}$$

$$\frac{\partial f^{D,g}}{\partial H_g} = \lambda_g^{H,D} \tag{86n}$$

$$\frac{\partial f^{D,g}}{\partial ICU_g} = \lambda_g^{ICU,D} \tag{86o}$$

$$\frac{\partial f^{D,g}}{\partial B_g^H} = 1 \tag{86p}$$

$$\frac{\partial f^{D,g}}{\partial B_g^{ICU}} = 1. \tag{86q}$$

Partials of $c_{g,h}$ with respect to ℓ for the multiplicative mixing model:

$$\frac{\partial c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h)}{\partial \ell_u^a} = \begin{cases} (\alpha_1 + \alpha_2) C_{g,h}^{(a)} \cdot \left(\ell_g^a\right)^{\alpha_1 + \alpha_2 - 1} & \text{if } u = g = h \\ \alpha_1 C_{g,h}^{(a)} \cdot \left(\ell_g^a\right)^{\alpha_1 - 1} \left(\ell_h^a\right)^{\alpha_2} & \text{if } u = g \neq h \\ \alpha_2 C_{g,h}^{(a)} \cdot \left(\ell_g^a\right)^{\alpha_1} \left(\ell_h^a\right)^{\alpha_2 - 1} & \text{if } g \neq h = u \end{cases} \tag{87a}$$

C. Calculation of Objective Coefficients.

C.1. Recursion for the Objective Coefficients. Algorithm 6 calculates the coefficients of each decision $\boldsymbol{u}_t$.

C.2. Calculation of M, γ and η. We recall the objective

$$V - \sum_{t=0}^{T-1} \sum_{g \in \mathcal{G}} \left(v_g\left(\boldsymbol{\ell}(t)\right) \cdot \left(S_g(t) + E_g(t) + I_g(t) + R_g(t)\right) + v_g(\mathbf{1}) \cdot R_g^{\mathrm{q}}(t)\right) + \sum_{g \in \mathcal{G}} \left(v_g^{\text{life}} + \chi\right) \cdot D_g(T),$$

as well as the economic model:

$$v_g(\boldsymbol{\ell}) := \begin{cases} v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{0-9 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) + v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{10-19 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) & \text{if } g = \text{20-29, 30-39, 40-49, 50-59, 60-69 y.o.} \\ 0 & \text{otherwise,} \end{cases}$$

Algorithm 6 CalculateObjectiveCoefficients algorithm.

Initialization: $\bar{X}_k := \boldsymbol{X}_k$

for $t = k, k+1, \ldots, T-1$ **do**

Calculate coefficients for $\boldsymbol{u}_\tau$ in the period t summand of Eq. (52) as:

$$\begin{cases} \boldsymbol{e}_t{}^\top, & \text{for } \tau = t \\ \boldsymbol{d}_t{}^\top \bar{A}_{t,\tau} B_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases} \quad [88]$$

where

$$\bar{A}_{t,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = t-1 \\ \bar{A}_{t,\tau+1} A_{\tau+1}, & \text{for } \tau = t-2, t-3, \ldots, k. \end{cases} \quad [89]$$

for $t = T$ **do**

Calculate coefficients for $\boldsymbol{u}_\tau$ in the period T summand of Eq. (52) as:

$$\boldsymbol{d}_T{}^\top \bar{A}_{T,\tau} B_\tau, \quad \text{for } \tau = T-1, T-2, \ldots, k \quad [90]$$

where

$$\bar{A}_{T,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = T-1 \\ \bar{A}_{T,\tau+1} A_{\tau+1}, & \text{for } \tau = T-2, T-3, \ldots, k. \end{cases} \quad [91]$$

Sum up the terms attributed to a common u_τ:

$$u_\tau = e_\tau{}^\top + \sum_{t=\tau+1}^{T} d_t{}^\top \bar{A}_{t,\tau} B_\tau. \quad [92]$$

return Coefficients for $\boldsymbol{u}_\tau$ for each $\tau = k, k+1, \ldots, T-1$

where

$$
\begin{aligned}
v_g^{\text{employment}}(\boldsymbol{\ell}) :=& w_g \cdot \left(\nu^{\text{work}} \cdot \ell_g^{\text{work}} + \nu^{\text{other activities}} \cdot \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} \frac{\boldsymbol{\ell}_h^{\text{other}} + \boldsymbol{\ell}_h^{\text{leisure}} + \boldsymbol{\ell}_h^{\text{transport}}}{3} + \nu^{\text{fixed}} \right) \\
=& \ell_g^{\text{work}} w_g \nu^{\text{work}} && [93] \\
&+ \left(\sum_{h \in \mathcal{G}} \sum_{a \in \{\text{other,leisure,transport}\}} \ell_h^a \right) w_g \nu^{\text{other activities}} \frac{1}{3|\mathcal{G}|} && [94] \\
&+ w_g \nu^{\text{fixed}}, && [95]
\end{aligned}
$$

and

$$
v_g^{\text{schooling}}(\boldsymbol{\ell}) := \theta \cdot f_g \cdot \delta_g \cdot v_{\text{20-29 y.o}}^{\text{employment}}(\mathbf{1}) \cdot \ell_g^{\text{school}}.
$$

From the equations above, we can write $\boldsymbol{M}$ (where the rows are indexed by the controls and the columns indexed by the compartments) as

$$
M[\ell_g^{\text{work}}, S_g] = M[\ell_g^{\text{work}}, E_g] = M[\ell_g^{\text{work}}, I_g] = M[\ell_g^{\text{work}}, R_g] = \begin{cases} -w_g \nu^{\text{work}}, & \text{if } g = \text{10-19, ..., 60-69 y.o.} \\ 0, & \text{otherwise.} \end{cases}
$$

$$
M[\ell_h^a, S_g] = M[\ell_h^a, E_g] = M[\ell_h^a, I_g] = M[\ell_h^a, R_g] = \begin{cases} \frac{-w_g \nu^{\text{other activities}}}{3|\mathcal{G}|}, & \text{if } g = \text{10-19, ..., 60-69 y.o., } h \in \mathcal{G},\ a \in \{\text{other}, \text{leisure}, \text{transport}\} \\ 0, & \text{otherwise.} \end{cases}
$$

$$
M[\ell_g^{\text{school}}, S_g] = M[\ell_g^{\text{school}}, E_g] = M[\ell_g^{\text{school}}, I_g] = M[\ell_g^{\text{school}}, R_g] = \begin{cases} -\theta f_g \delta_g v_{\text{20-29 y.o}}^{\text{employment}}(\mathbf{1}), & \text{if } g = \text{0-9, 10-19 y.o.} \\ 0, & \text{otherwise.} \end{cases}
$$

$$
M[\cdot, \cdot] = 0, \qquad \text{otherwise.}
$$

Similarly, we can write $\boldsymbol{\gamma}$ (indexed by the compartments for each group) as

$$
\begin{aligned}
&\boldsymbol{\gamma}[S_g] = \boldsymbol{\gamma}[E_g] = \boldsymbol{\gamma}[I_g] = \boldsymbol{\gamma}[R_g] = \begin{cases} -w_g \nu^{\text{fixed}}, & \text{if } g = \text{10-19, ..., 60-69 y.o.} \\ 0, & \text{otherwise.} \end{cases} && [96] \\
&\boldsymbol{\gamma}[R_g^q] = -v_g(\mathbf{1}), \qquad \text{for } g \in \mathcal{G} && [97] \\
&\boldsymbol{\gamma}[\cdot] = 0, \qquad \text{otherwise.} && [98]
\end{aligned}
$$

Finally, we have

$$
\boldsymbol{\eta} = \left[\underbrace{0, \cdots, 0, (v_g^{\text{life}} + \chi)}_{\text{compartments of group } g}, 0, \cdots \right]^\top_{g \in \mathcal{G}},
$$

where the only non-zero values are in the indices corresponding to compartment D_g of each group g.

References

1. L Wille, et al., SOCRATES: an online tool leveraging a social contact data sharing initiative to assess mitigation strategies for covid-19. *BMC Res. Notes* **13** (2020).
2. G Béraud, et al., The French connection: The first large population-based contact survey in France relevant for the spread of infectious diseases. *PLOS ONE* **10**, 1–22 (2015).
3. A Mas-Colell, MD Whinston, JR Green, , et al., *Microeconomic theory.* (Oxford University Press New York) Vol. 1, (1995).
4. D Acemoglu, V Chernozhukov, I Werning, MD Whinston, Optimal targeted lockdowns in a multi-group SIR model, (NBER), Working Paper 27102 (2020).
5. H Salje, et al., Estimating the burden of SARS-CoV-2 in France. *Science* **38** (2020).
6. Q Bi, et al., Epidemiology and transmission of COVID-19 in 391 cases and 1286 of their close contacts in Shenzhen, China: a retrospective cohort study. *The Lancet Infect. Dis.* **20** (2020).
7. Z Du, et al., Serial interval of covid-19 among publicly reported confirmed cases. *Emerg. Infect. Dis.* **26** (2020).
8. O Diekmann, JAP Heesterbeek, MG Roberts, The construction of next-generation matrices for compartmental epidemic models. *J. The Royal Soc. Interface* **7**, 873–885 (2010).
9. A Perasso, An introduction to the basic reproduction number in mathematical epidemiology. *ESAIM: Proc. Surv.* **62**, 123–138 (2018).

10. French Government, Données hospitalières relatives à l'épidémie de COVID-19 (https://www.data.gouv.fr/fr/datasets/donnees-hospitalieres-relatives-a-lepidemie-de-covid-19/) (2020) Accessed October 21, 2020.
11. Google, COVID-19 community mobility report (https://www.google.com/covid19/mobility/) (2020) Accessed October 21, 2020.
12. H Mohammad, Coronavirus : un habitant de Bobigny considéré comme le nouveau "patient zéro" (France Bleu) (2020) Accessed January 9, 2021.
13. INSEE, Estimation de population par département, sexe et âge quinquennal - années 1975 à 2020 (https://www.insee.fr/fr/statistiques/fichier/1893198/estim-pop-dep-sexe-aq-1975-2020.xls) (2020) Accessed June 19, 2020.
14. INSEE, T404 : Salaire brut en équivalent temps plein, par tranche d'âge simplifiée, catégorie socioprofessionnelle simplifiée et région (https://www.insee.fr/fr/statistiques/fichier/4204500/T404.xls) (2016) Accessed June 19, 2020.
15. INSEE, Au quatrieme trimestre 2019, le taux de chômage passe de 8,5 % a 8,1 % (https://www.insee.fr/fr/statistiques/4309346) (2019) Accessed June 19, 2020.
16. Banque de France, Point sur la conjoncture française à fin avril 2020 (https://www.banque-france.fr/sites/default/files/media/2020/06/10/point-conjoncture_avril-2020.pdf) (2020) Accessed June 19, 2020.
17. Banque de France, Point sur la conjoncture française à fin mai 2020 (https://www.banque-france.fr/sites/default/files/media/2020/06/11/point-conjoncture-09_juin-2020-20200609.pdf) (2020) Accessed June 19, 2020.
18. INSEE, T402 : Salaire brut en équivalent temps plein, par secteur d'activité, région et département (https://www.insee.fr/fr/statistiques/fichier/4204500/T402.xls) (2016) Accessed June 19, 2020.
19. D Duque, et al., Timing social distancing to avert unmanageable covid-19 hospital surges. *Proc. Natl. Acad. Sci.* **117**, 19873–19878 (2020).
20. M Lehot, BL Borgne, Covid-19 : ces chiffres qui montrent que Paris a dépassé le seuil d'alerte maximale depuis le 25 septembre (2020) Accessed October 5, 2020.
21. S Godeluck, Coronavirus : plus de 2.000 patients en réanimation en Île-de-France (Les Echos) (2020) Accessed June 30, 2020.
22. C Sterlé, V Alexandre, S de Livonnière, Coronavirus en Île-de-France : 13 000 malades encore hospitalisés, déjà 40 000 tests positifs (Le Parisien) (2020) Accessed June 30, 2020.
23. Agence Régionale de Santé — Île-de-France, Déconfinement : stratégie de dépistage et recherche de cas-contacts (2020) Accessed June 30, 2020.
24. C Sterlé, Île-de-France : y aura-t-il assez de tests Covid avant Noël ? (https://www.leparisien.fr/societe/sante/ile-de-france-y-aura-t-il-assez-de-tests-covid-avant-noel-18-12-2020-8414952.php) (2021) Accessed January 28, 2021.