Quantifying the Benefits of Targeting for Pandemic Response

Sergio Camelo[a,1,2], **Dragos F. Ciocan**[b,1], **Dan A. Iancu**[b,c,1], **Xavier S. Warnes**[c,1], **and Spyros I. Zoumpoulis**[d,1]

[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

This manuscript was compiled on March 22, 2021

To respond to pandemics such as COVID-19, policy makers have relied on interventions that target specific population 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 a set of associated algorithms that compute optimally targeted, time-dependent interventions that coordinate across *two* dimensions of heterogeneity: population group characteristics *and* the specific activities that individuals engage in during the normal course of a day. We showcase a complete implementation in a case study focused on the Île-de-France region of France, based on commonly available hospitalization, community mobility, social contacts and economic data. We find that optimized dual-targeted policies have a simple and explainable structure, imposing less confinement on group-activity pairs that generate a relatively high economic value prorated by activity-specific social contacts. When compared to confinements based on uniform or less granular targeting, 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 population group. 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 intermediate trade-offs may be achievable through limited targeting.

Pandemic management | Confinement | Targeted interventions | Optimization | COVID-19 |

The COVID-19 pandemic has forced policy makers worldwide to deploy a suite of measures aimed at curtailing the spread, including testing, mask wearing, vaccination, and large-scale confinement and social distancing. In determining such measures, a key recognition has been that individuals are heterogeneous in several important dimensions, such as their health outcomes when infected or their daily activities leading to new infections but also creating economic value. Targeting interventions to explicitly account for such heterogeneity could be an important lever to mitigate a pandemic's health and economic 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 targeting mechanism studied in the literature and also implemented by policy makers has been to tailor interventions to *population groups* defined either by age (1–4), age and economic sector (5), geography (6, 7) or clinical risk (8). By enforcing stricter confinements for higher risk groups (e.g., older populations when considering mortality risk or younger populations when considering the risk of new infections), such targeted policies have been shown to generate potentially significant improvements in health outcomes, and even in economic value if optimally tailored (1). However, they could also lead to potentially contentious measures that discriminate based on age or other protected features, giving rise to important ethical and even legal challenges. Such targeted interventions have been implemented during the COVID-19 pandemic in several settings – e.g., with stricter confinements applied to older population groups in Finland (9), Ireland (10), Israel (11) and Moscow (12) or curfews applied to children and youth in Bosnia and Herzegovina (13) and Turkey (14) – but some of the measures were deemed ageist and unconstitutional and were eventually overturned (11, 13).

A different targeting dimension extensively employed in practice has been to tailor confinements to specific *activities* conducted during a typical day. Restrictions of varying degrees on workplace presence, schools, recreation venues, retail activity or outdoor leisure have become commonplace measures during the COVID-19 pandemic. This has been driven by the recognition that certain activities may be responsible for a much larger number of new infections, as social contacts vary significantly depending on whether individuals are engaged in work, schooling, leisure, transport, or other activities (15). In the academic literature, existing studies have only compared a discrete set of candidate scenarios that confine certain activities (16, 17), but without any rigorous quantification of the best activity-targeted policy.

Explicitly considering this dimension of targeting *in addition* to population groups has the potential to generate

Significance Statement

In the fight against pandemics such as COVID-19, policy makers rely on interventions that target distinct groups of individuals or activities for confinement. Such targeting can however lead to contentious policies that excessively confine certain groups like the elderly, so it is critical to understand its relative merits. We propose and implement a rigorous framework to quantify these merits and demonstrate it in a case study of COVID-19 interventions in Île-de-France. We find that optimally designed interventions that differentiate based on both population groups and activities achieve significantly better health and economic outcomes overall and also reduce confinement time for each group, compared to less targeted interventions. The implementation, publicly available at http://insead.arnia.ro, is flexible and portable to other geographies.

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

significant improvements, both by reducing the economic and health burden, but also by potentially overturning the prevailing insight that specific age groups should uniformly face stricter confinements. For instance, high risk groups could remain unrestricted in certain activities, as long as they are protected by confining other groups away from these same activities. Examples of such interventions that target *both* age groups and activities have been seen in practice, e.g., by setting aside dedicated hours when only the senior population is allowed to shop at supermarkets (18) or by restricting only higher age groups from in-person work activities (11).

Given the potential but also the contentiousness of targeted interventions, 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. Several natural research questions emerge: How large are the health and economic benefits of dual targeting? Would dual targeting lead to significant synergies, and if so why? Could dual targeting reduce the time in confinement for *every* population group systematically? More broadly, what is the relationship between the effectiveness and the level of targeting that an intervention enacts across distinct population groups?

Answers to these questions carry important policy implications, but to the best of our knowledge no study to date has provided a rigorous framework that generates quantifiable answers as important analytical and data-related challenges stand in the way. First, embedding complex controls into the highly non-linear and time-dependent dynamics of epidemic models makes optimization extremely challenging, which is why existing proposals rely on either simulating a small number of candidate interventions (16, 17) or picking the best option from a restricted class of policies for which exhaustive search is computationally feasible (such as trigger policies based on hospital admissions, 19). In contrast, we develop a structured optimization framework for the underlying control problem, together with a specialized algorithm to efficiently solve it. Second, calibrating this model requires aligning disparate epidemiological and socio-economic data, at the level of granularity required by the targeting. 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 is publicly available at http://insead.arnia.ro.

Materials and Methods

Our framework relies on a flexible model that captures several important real-world considerations. We extend a multi-group SEIR model to capture controls that target based on (i) *age groups*, and (ii) types of *activities* that individuals engage in. Different policy interventions can be embedded: we include time-dependent confinements as well as testing and quarantining in our study, but vaccinations can also be easily accommodated. The controls modulate the rate of social contacts and the economic value generated, and the objective of the control problem is to minimize a combination of emotional and economic losses caused by deaths, illness, and activity restrictions. The model captures important resource constraints (such as hospital, ICU, and testing capacities), and allows explicitly controlling the amount of targeting through "limited disparity" constraints that limit the difference in the extent of confinement imposed on distinct population groups.

Epidemiological Model and Controls. We rely on a modified version of the discretized SEIR (Susceptible-Exposed-Infectious-Recovered) epidemiological model with multiple population groups that interact with each other. In our case study we use nine groups determined by age, with the youngest group capturing individuals with age 0-9 and the oldest capturing individuals with age 80 or above. Time is discrete, indexed by $t = 0, 1, \ldots, T$ and measured in days. We assume that no infections are possible beyond time T.

For a population group g in time period t, the compartmental model includes states $S_g(t)$ (susceptible to be infected), $E_g(t)$ (exposed but not yet infectious), $I_g(t)$ (infectious but not confirmed through testing and thus not quarantined), $I_g^q(t)$ (infectious and confirmed through testing and thus quarantined; this state is further sub-divided based on the severity of symptoms), $R_g(t)$ (recovered but not confirmed as having had the virus), $R_g^q(t)$ (recovered and confirmed as having had the virus), and $D_g(t)$ (deceased). 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 denote the entire vector of states at time $0 \leq t \leq T$ by $\boldsymbol{X}_t$. See SI Section 2 for details.

Individuals interact in activities belonging to the set $\mathcal{A} = \{\mathsf{work}, \mathsf{transport}, \mathsf{leisure}, \mathsf{school}, \mathsf{home}, \mathsf{other}\}$. 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 a *normal course* of life (i.e., no confinement). In our study, we take $\ell_g^{\mathsf{home}}(t) = 1$, meaning that the number of social contacts at home is unchanged irrespective of the confinement policy followed.* We denote the vector of all activity levels for group g at t by $\boldsymbol{\ell}_g(t)$, and we also refer to $\boldsymbol{\ell}_g(t)$ as confinement decisions when no confusion can arise.

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 $c_{g,h}(\boldsymbol{\ell}_g, \boldsymbol{\ell}_h) = \sum_{a \in \mathcal{A}} C_{g,h}^a \cdot (\ell_g^a)^{\alpha_1} \cdot (\ell_h^a)^{\alpha_2}$, 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 (20), which is based on the French social contact survey data in (15), and we estimate α_1, α_2 from health outcome data (21) and Google mobility data (22).

Social mixing leads to new infections, proportional to

$$\propto S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \frac{I_h(t)}{N_h(t) + R_h^q(t)} \right), \qquad [1]$$

where $N_g(t) := S_g(t) + E_g(t) + I_g(t) + R_g(t)$ captures all living individuals in group g who are not confirmed to have had the disease.

Besides confinements, we also model targeted viral testing decisions, which capture how much of a finite capacity of tests to allocate to each age group. We model random mass testing, where a test detects infectious individuals with probability equal to the fraction of infectious individuals in the group's population. Infected individuals in group g who are detected are placed in the quarantined SEIR state $I_g^q(t)$, where they can no longer infect others. Testing is discussed in more detail in SI Section 2.

When the total patient inflow into the hospital or ICU exceeds the remaining available beds, we assume that policy makers turn away patients from each group g according to a proportional rule, admitting patients proportionally to the demand from that group up to available bed capacity, and turning away remaining patients.†

In compact notation, we use $\boldsymbol{u}_t$ to denote the vector of all decisions at time $t \in \{0, \ldots, T-1\}$, and we use the function

$$F_t(\boldsymbol{X}_t, \boldsymbol{u}_t) := \frac{\Delta \boldsymbol{X}_t}{\Delta t} \qquad [2]$$

to capture the dynamics of the SEIR states.

Objective. Our objective captures losses of both economic and emotional nature that are directly attributable to the pandemic. For economic losses, we capture lost value stemming from two sources: (a) confinement and quarantine, and (b) deceased individuals who no longer generate economic output. To model (a), we

*The number of social contacts at home arguably increases when other activities are restricted, but these contacts would likely be with the same individuals and would not constitute independent trials that could result in infections, as in a typical SEIR model. We therefore assume the contacts in the home activity are unchanged, but our model could easily accommodate other assumptions.

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

assign a daily economic value $v_g(\boldsymbol{\ell})$ to each individual in group g that depends on the activity levels $\boldsymbol{\ell}$ across all groups and activities. For the working age groups, $v_g(\boldsymbol{\ell})$ comes from wages from employment and is a linear function of group g's activity level in work (ℓ_g^{work}) and of the average activity levels in leisure, other and transport for the entire population (equally weighted). This reflects that the value generated in some industries, like retail, is impacted by confinements across all these three activities. For the school age groups, $v_g(\boldsymbol{\ell})$ captures future wages from employment due to schooling and depends only on the group's activity level in school (ℓ_g^{school}). For (b), we determine the wages that an individual would have earned based on their current age until retirement age under the prevailing wage curve, and denote the resulting amount of lost wages with v_g^{life}. The overall aggregate economic loss, which we denote by Economic Loss($\boldsymbol{u}_{0:T-1}$), is:[‡]

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

where V captures the economic value that would have been generated without the pandemic (defined precisely in SI Section 2).

We capture emotional losses by associating a (non-pecuniary) cost χ to each death, which results in an aggregate loss of:

$$\text{Emotional Loss}(\boldsymbol{u}_{0:T-1}) := \chi \cdot \sum_{g \in \mathcal{G}} D_g(T). \qquad [3]$$

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 ($\chi \to \infty$).

Optimization Problem. The optimization problem we solve is to find control policies (for confinement and testing) that minimize the sum of economic and emotional losses. Formally, we solve:

$$\min_{\boldsymbol{u}_{0:T-1}} [\text{Economic Loss}(\boldsymbol{u}_{0:T-1}) + \text{Emotional Loss}(\boldsymbol{u}_{0:T-1})] \qquad [4]$$

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, hospital beds, and tests. The model and its dynamics are presented in detail in SI Section 2.

Re-Optimization with Linearized Dynamics — ROLD. The problem above is computationally challenging: the term in Eq. (1) that captures new infections by multiplying $S_g(t)$ and $I_h(t)$ (as is typical in all SEIR models) implies that states depend on the control policy in a highly non-convex fashion, which makes the problem difficult to solve to optimality even for small time horizons. To tackle this problem, we rely on ideas from model predictive control to design an approximate algorithm that remains tractable for realistic time horizons and problem sizes. The algorithm is based on linearizing the true SEIR dynamic, solving a tractable convex optimization problem to determine a control policy, and applying this policy in the current period. In particular, at each time step k:

1. Given the current state $\boldsymbol{X}_k$ and a nominal control sequence $\hat{\boldsymbol{u}}_{k:T-1}$ for all remaining periods, we calculate a nominal system trajectory $\hat{\boldsymbol{X}}_{k:T}$ under the true dynamic in Eq. (2). (The nominal control is set to a solution obtained by gradient descent at $k = 0$, and to the algorithm's own output for periods $k > 0$.)

2. We linearize the state dynamics around $(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)$ at all future periods $t \in \{k, \dots, T-1\}$:

$$F_t(\boldsymbol{X}_t, \boldsymbol{u}_t) \approx F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) + \nabla_X F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t) + \nabla_u F_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{u}_t - \hat{\boldsymbol{u}}_t),$$

   where $\nabla_X F_t$ and $\nabla_u F_t$ denote the Jacobians with respect to $\boldsymbol{X}_t$ and $\boldsymbol{u}_t$, respectively, and we similarly linearize the objective in Eq. (4). This leads to a linear optimization problem over decisions $\boldsymbol{u}_k, \dots, \boldsymbol{u}_{T-1}$.

3. We solve the linear program to obtain $\boldsymbol{u}^*_{k:T-1}$.

4. The true states are updated using the decision $\boldsymbol{u}^*_k$ at period k.

Resolving at each time k is aimed to prevent the linearization errors from growing too large. To further ensure this, we also iteratively update the linearization *within* each period k: each iteration is 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 the updated control is used as a seed for the next iteration. A detailed description of ROLD can be found in SI Section 3.

Additional modeling features. The framework we described is very flexible, and allows easily embedding constraints to capture several considerations, such as ensuring that the policies are only allowed a specific limited level of targeting in either age-groups or activities, that they can only be re-evaluated with a certain frequency (e.g., every two weeks), that confinements in a certain activity reduce the maximum level possible in another activity (e.g., school closures reducing the maximum activity level in work), etc. We refer to SI Section 2 for details.

Parametrization and Model Calibration. We adopt the confidence regions for SEIR parameters reported in (23) for the Île-de-France region, which we complement with Google mobility data to approximate the mean effective lockdowns for all activities during the horizon of interest. Based on these, we simulate our SEIR model to generate several potential sample paths, which we use in conjunction with real data on hospital and ICU utilization and deaths to generate a fitting error metric. Lastly, we estimate values of all parameters of interest by minimizing the sample-average-approximation of the error metric. 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. We provide all the details for calibration and parameter specification in the SI, where we also report results from sensitivity and robustness analyses on the fitted parameters.

Experimental Setup. All results reported in the main paper are obtained under a testing capacity of zero, so only confinement decisions are optimized and compared (SI Section 7 discusses the additional benefits of targeting the viral testing strategy.) We use a horizon of $T = 90$ days, and 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. The details of the experimental setup are provided in SI Section 6.

To quantify the benefits of targeting, 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× the annual GDP per capita in France[§]. For each χ value, we calculate all the ROLD policies of interest, and we record separately the economic losses and the number of deaths generated by each policy.

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 emotional 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: i.e., 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.

[‡] For a time series of vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_n$, we use $\boldsymbol{v}_{i:j} := [\boldsymbol{v}_i, \dots, \boldsymbol{v}_j]$ to denote the concatenation of vectors $\boldsymbol{v}_i$ through $\boldsymbol{v}_j$.

[§] We quantify the emotional 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.

When comparing the different types of targeting, note that neither AGE nor ACT Pareto-dominate each other, and in fact neither policy dominates in terms of the total loss objective (see SI Figure S3). In contrast and quite 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 the ability to target *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, Hybrid-t AND, and Hybrid-t OR mimic implementations in the U.S. Austin area (19) and, respectively, France (24). These policies switch between a stricter and a relaxed confinement level based on conditions related to hospital occupancy and the rate of new infections (details are provided in SI Section 5). Additionally, we also consider two extreme benchmarks corresponding to enforcing a "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 as FO ensuring low economic losses.

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 also 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, which confirms that dual targeting is critical and powerful.

The Pareto-dominance of ROLD AGE-ACT implies its dominance in terms of the total loss objective, as confirmed in SI Figure S3. Figure S5 in SI confirms the robustness of these results under more problem instances.

How do gains arise from dual targeting? To address this question, 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 (25) and is representative of the overall behavior we observe across all experiments (SI Section 7 has a more detailed discussion of how the policies vary with problem parameters). 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., $\frac{dv_g(\boldsymbol{\ell})}{d\ell_g^a} / \sum_{h\in\mathcal{G}} C_{hg}^a$, henceforth referred to as *"econ-to-contacts-ratio"*. 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 with a horizon of $T = 90$ days, and then train regression decision trees to predict the optimal ROLD activity levels as a function of several features (see SI Section 7 for details). The results are captured in Figure 2b and Figure S10 in the SI. 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 seems to be the most salient feature when targeting confinements, with higher ratios leading to higher activity levels in all activities considered. The trees also confirm that time and the emotional 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.

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. Accordingly, we see that 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; and the 70+ y.o. groups are completely open in leisure. These complementary confinement schedules allow ROLD to reduce both the number of deaths and economic losses, with the important added benefit that they do not require completely confining any age group.

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 (see SI Section 7 for a formal definition). 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

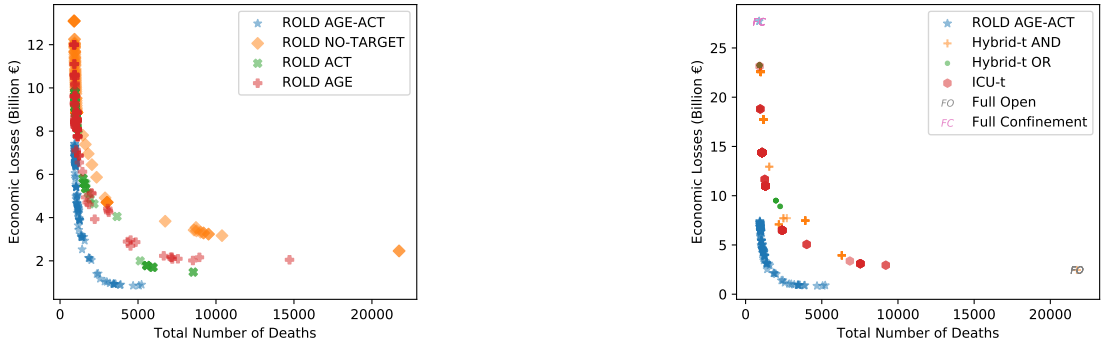

(a) Comparison between four ROLD policies with different levels of targeting

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

Fig. 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 emotional cost of death χ. We include 128 distinct values of χ from 0 to 990$\times$, and panel (b) also includes a very large value ($\chi = 10^{16}\times$).

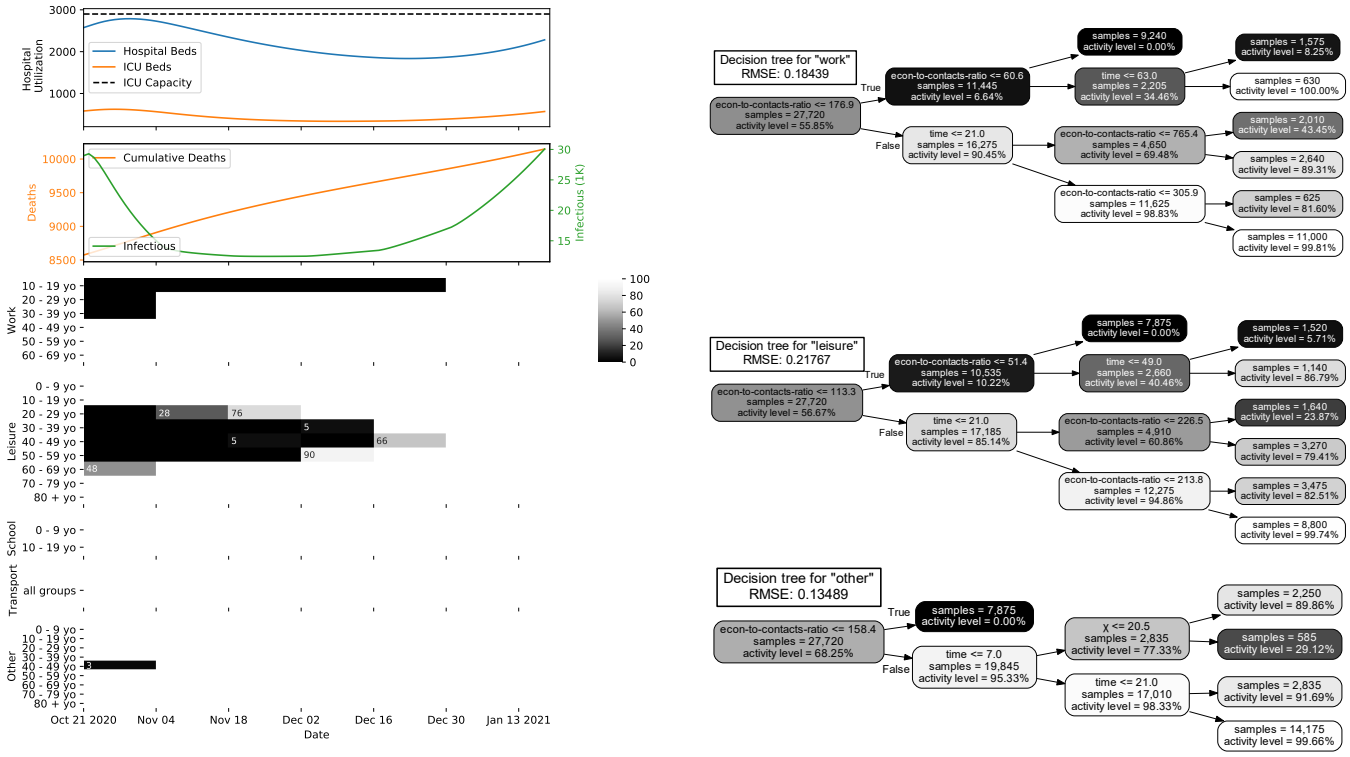

(a) Optimized ROLD confinement policy and its estimated impact

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

Fig. 2. The optimized ROLD AGE-ACT policies for problem instances with a 90-day optimization horizon starting on October 21, 2020. Figure **(a)** corresponds to a problem instance where the emotional cost of death χ is $50\times$. 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). 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 ≤ 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.

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 economic and emotional loss objective in Eq. (4).[¶]

It is worth noting that although ROLD AGE-ACT reduces confinements for every population group compared to less targeted policies, it does not do so uniformly, and it can sometimes lead to a larger discrepancy in the confinements

[¶] That each group should spend less time in confinement can be seen as a reasonable fairness requirement, e.g., consistent with Rawlsian justice (26). In this sense, allowing ROLD increased levels of explicit discrimination has the potential to improve both efficiency and fairness — something that has been noted before in the algorithmic fairness literature (27).

faced by different age groups: those aged 10-59 are generally more confined than those aged 0-9 or 60+.

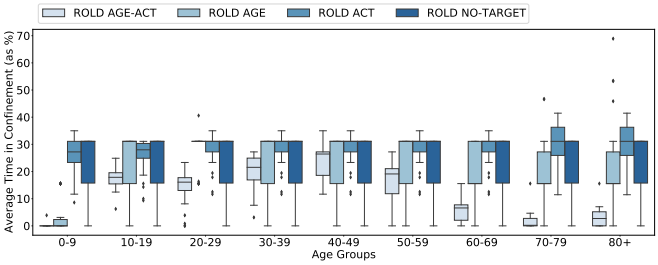

Fig. 3. Average time in confinement for the ROLD policies with different targeting types. Each boxplot depicts the fraction of time the age group spends in confinement under the respective policy averaged over the activities relevant to that age-group, for different problem instances parameterized by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

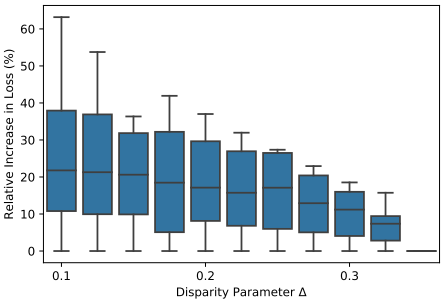

Fig. 4. The impact of limited disparity requirements. The plot shows the relative (%) increase in total loss generated by a ROLD policy, compared to a fully targeted policy, as a function of the disparity parameter Δ that measures the maximum allowed difference in activity levels for distinct age groups. 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.

The impact of limited disparity requirements. That targeted policies confine some age groups more than others could be perceived as evidence of disparate or unfair treatment, so it is important to quantify how an intervention's effectiveness is impacted when requiring less differentiation across age groups. To examine this, 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 (see definition in SI Section 7). A value $\Delta = 0$ thus 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 could also be achieved with some intermediate disparity.

Discussion

Our case study suggests that an optimized intervention that targets both population groups *and* activities carries significant promise for alleviating a pandemic's health, economic and even psychological burden, but also points to certain challenges in designing and implementing such finely targeted interventions 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, targeted confinements based on age groups and activities can reduce economic losses more than any of the simpler interventions that uniformly confine population 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 largely 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 increased transparency into how targeted confinement decisions could be made. This is valuable in a real-world implementation even if policy makers do not directly rely on such metrics to justify their actions.[‖]

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 population groups. As different age groups may be responsible for generating a larger economic value prorated by social contacts in distinct activities, a dual-targeted confinement policy may enable all age groups to remain more active and engage in activity levels that more closely resemble normal life 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 discriminating based on age and can result in differences in confinements across population 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. Apart from the examples discussed in the introduction, such dual-targeting can also arise *implicitly* in interventions that only seem to target activities. As an example, France is currently implementing a population-wide 6 p.m. to 6 a.m. curfew (28), while maintaining school and work activities largely de-confined. This is effectively implementing activity level restrictions similar to ROLD AGE-ACT: since a typical member of the 20-65 y.o. group is engaged in work

[‖] Policy makers may prefer to tie confinement decisions to certain observed outcomes, such as the confirmed infections or hospitalizations for specific age group-activity pairs. Provided that social contacts data are accurate, such observables will be strongly correlated with the raw social contact data, so our insights regarding the econ-to-contacts ratios would also approximately transfer to, e.g., ratios of the economic value prorated by the infections in specific age group-activity pairs.

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 by the curfew.

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 based on population groups.

Challenges and limitations. An immediate practical challenge is data availability. Estimating our model requires access to data on (i) hospitalizations and deaths, (ii) community activity levels over time, (iii) social contacts between different population groups in each relevant activity under normal (no confinement) conditions, and (iv) economic value generated by each population group in each activity. 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 (e.g., 15) and Google mobility reports use different activity categories. Responding to these difficulties requires non-trivial fitting, as we explain in SI Section 4, and highlights the value of collecting epidemiological and socio-economic data in a coordinated and structured fashion that allows minimizing fitting errors.

Availability of data also constrains our model's structure in several important ways. The first limitation is that social contacts between age groups only depend on confinements in the same activity, since the available social contacts dataset (15) only reports interactions in the same activity. However, contacts occur as individuals from different groups are engaged in different activities (e.g., a professional in the services industry interacts during work with individuals who may be engaged in leisure activities). A more refined contact mixing model that captures such interactions would be more appropriate for this study, provided that relevant social contact data are available.

Another limitation is that our current economic model assumes additive value across activities and age groups. This choice ignores that complementary activities conducted by different groups may be needed to generate output: for example, the young age groups engaged in school produce value in conjunction with educational staff engaged in work. Relatedly, the available data is insufficient to precisely estimate the contribution of leisure, transport and other activities in a group's economic output, which is why our model assigns additive and equal economic value to all three. Our framework can be amended to more accurately model the economic objectives, should more detailed economic data be available.

Another challenge with targeted interventions is the perception that they may lead to unfair outcomes, as certain population groups face more confinement than others. Such discrimination does arise in the optimized dual-targeted policies, and as discussed, our framework can partially address the concerns through explicit constraints that limit the disparities across groups. Our requirement that limited disparity should 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, it may be more meaningful to impose fairness requirements based on the intervention's *outcomes*, e.g., requiring that the health or economic losses faced by different population groups satisfy certain axiomatic fairness properties (26).

Lastly, 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 Section 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.

1. D Acemoglu, V Chernozhukov, I Werning, MD Whinston, Optimal targeted lockdowns in a multi-group SIR model, (NBER), Working Paper 27102 (2020).
2. L Matrajt, J Eaton, T Leung, ER Brown, Vaccine optimization for covid-19, who to vaccinate first? (2020).
3. JR Goldstein, T Cassidy, KW Wachter, Vaccinating the oldest against covid-19 saves both the most lives and most years of life. *Proc. Natl. Acad. Sci.* **118** (2021).
4. D Bertsimas, et al., Optimizing vaccine allocation to combat the covid-19 pandemic (2020).
5. CA Favero et al., Restarting the economy while saving lives under covid-19 (2020).
6. JR Birge, O Candogan, Y Feng, Reducing economic losses with targeted closures, (University of Chicago, Becker Friedman Institute for Economics), Working Paper 2020-57 (2020).
7. S Chang, et al., Mobility network models of covid-19 explain inequities and inform reopening. *Nature* **589**, 1–6 (2020).
8. T Evgeniou, et al., Epidemic models for personalised covid-19 isolation and exit policies using clinical risk predictions (2020).
9. H Tiirinki, et al., Covid-19 pandemic in finland–preliminary analysis on health system response and economic consequences. *Heal. policy technology* **9**, 649–662 (2020).
10. S Harrison, Coronavirus: Ireland's restrictions eased for over 70s (2020) Access. 01/12/2021.
11. J Magid, Cabinet approves removal of age restriction for those returning to work (2020) Accessed January 12, 2021.
12. H Foy, Russian businesses prepare for fresh lockdowns as Covid-19 cases soar (2020) Accessed January 12, 2021.
13. Reuters Staff, Bosnian region eases lockdown on seniors, children after court ruling (2020) Accessed January 12, 2021.
14. N Kanbur, S Ankgül, Quaranteenagers: A single country pandemic curfew targeting adolescents in Turkey. *J. Adolesc. Heal.* **67** (2020).
15. 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).
16. K Prem, 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 (2020).
17. L Di Domenico, G Pullano, CE Sabbatini, PY Boëlle, V Colizza, Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC medicine* **18**, 1–13 (2020).
18. J Aguilera, 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.
19. D Duque, et al., Timing social distancing to avert unmanageable covid-19 hospital surges. *Proc. Natl. Acad. Sci.* **117**, 19873–19878 (2020).
20. 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).
21. 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.
22. Google, COVID-19 community mobility report (https://www.google.com/covid19/mobility/) (2020) Accessed October 21, 2020.
23. H Salje, et al., Estimating the burden of SARS-CoV-2 in France. *Science* **38** (2020).
24. 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.
25. FE Alvarez, D Argente, F Lippi, A simple planning problem for COVID-19 lockdown, (National Bureau of Economic Research), Working Paper 26981 (2020).
26. HP Young, *Equity.* (Princeton University Press), (1994).
27. J Kleinberg, J Ludwig, S Mullainathan, A Rambachan, Algorithmic fairness. *AEA Pap. Proc.* **108**, 22–27 (2018).
28. Reuters, France goes under nationwide 6pm curfew as Covid-19 death toll surpasses 70,000 (2021) Accessed January 12, 2021.