medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Quantifying and Realizing the Benefits of Targeting for Pandemic Response

Sergio Camelo
Institute for Computational and Mathematical Engineering, Stanford University camelo@stanford.edu,

Florin Ciocan
Technology and Operations Management, INSEAD, florin.ciocan@insead.edu,

Dan A. Iancu
Operations, Information & Technology, Graduate School of Business, Stanford University, daniancu@stanford.edu,

Xavier S. Warnes
Woods Institute, Stanford University, xwarnes@stanford.edu,

Spyros I. Zoumpoulis
Decision Sciences, INSEAD, spyros.zoumpoulis@insead.edu.

November 29, 2022

Problem definition: To respond to pandemics such as COVID-19, policy makers have relied on interventions that target specific population groups or activities. Since targeting is operationally challenging and contentious, rigorously quantifying its benefits and designing practically implementable policies that achieve some of these benefits is critical for effective and equitable pandemic control.

Methodology/results: We propose a flexible modeling framework that allows computing optimized interventions that target *two* dimensions of heterogeneity: age groups *and* the specific activities that individuals normally engage in. We showcase a complete implementation focused on the Île-de-France region of France, based on commonly available public data. We use this case study to quantify the potential benefits of targeting in an idealized setting where dual targeting is implementable. We find that dual-targeted policies can improve both health and economic outcomes, while maintaining higher activity levels for most age groups (and importantly, for those most confined). We fit decision trees that explain the decisions and the gains afforded by dual targeting. Gains are generated by leveraging complementarities, imposing less confinement on group-activity pairs with high marginal economic value prorated by social contacts. We also derive additional practical insights for identifying settings where gains are high. Because dual targeting can face significant implementation challenges, we introduce two practical proposals inspired by real-world interventions — based on curfews and recommendations — that achieve a significant portion of the benefits without explicitly discriminating based on age.

Implications: Our framework suggests that significant benefits could be generated by targeting confinements based on activities and age groups, and that some of these benefits can be materialized through realistic interventions that avoid contentious age targeting.

Key words: Pandemic management; Targeted interventions; Curfews; Optimization; COVID-19
History:

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

1. 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 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 in pandemic management, and applying it to the COVID-19 pandemic as a real-world case study.

Targeting has been implemented in several different ways during the COVID-19 pandemic. We include a detailed survey of COVID-19 targeted restrictions in Appendix Table EC.14. One real-world contentious example 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 have been implemented in several settings – e.g., with stricter confinements applied to older groups in Finland (Tiirinki et al. 2020), Ireland (Harrison 2020), Israel (Magid 2020) and Moscow (Foy 2020), or curfews applied to children and youth in Bosnia and Herzegovina (Reuters Staff 2020) and Turkey (Kanbur and Ankgül 2020) – but some of the measures were deemed ageist and unconstitutional and were eventually overturned (Magid 2020, Reuters Staff 2020).

A different example of targeting extensively employed in practice has been to tailor confinements to specific *activities* conducted during a typical day. This has been driven by the recognition that different activities (or more specifically, population interactions in locations of certain activities) such as work, schooling, transport or leisure can result in significantly different patterns of social contacts and new infections. This heterogeneity has been recognized in numerous implementations that differentially confine activities through restrictions of varying degrees on schools, workplaces, recreation venues, retail spaces, etc. Additionally, some practical implementations even differentiated based on *both* age groups and activities, e.g., by setting aside dedicated hours when only the senior population was allowed to shop at supermarkets (Aguilera 2020), or by restricting only higher age groups from in-person work activities (Magid 2020).

As these examples suggest, targeted interventions have merits but also pose potentially significant downsides. On the one hand, targeting can generate improvements in both health and economic

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

outcomes, giving policy makers an improved lever when navigating difficult trade-offs. Additionally, explicitly considering *multiple* dimensions of targeting simultaneously – such as 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 is already in place in existing real-world policy implementations, it is critical to transparently model and quantify its benefits and downsides, as well as examining practical policies that can realize some of these benefits. This gives rise to several natural research questions: How large are the health and economic benefits of interventions that can engage in progressively finer targeting? Does finer targeting lead to significant synergies, and is there an interpretable mechanism through which this happens? Because finer targeting may be impractical, are there ways to capture some of its benefits with easier-to-implement policies?

1.1. Contributions

At a high level, the main contributions of this work are in providing a rigorous framework to quantify the potential benefits from targeted confinements and proposing practical implementations that can achieve some of these benefits.

The framework. To develop robust insights on the benefits of targeting, we anchor our analysis on a modeling framework that leverages optimization and publicly available, real data. We embed an optimization framework within a traditional SEIR epidemiological model that differentiates policies based on both population groups and activities, and balances the lost economic value with the cost of deaths. Since the resulting optimization problem is highly non-convex, we design an algorithm — referred to as ROLD (Re-Optimization with Linearized Dynamics) — that can tractably produce high-quality approximate solutions through a linearization and optimization procedure inspired by model predictive control (Bemporad 2006, Camacho and Alba 2013) and trust region methods (Yuan 2015). We then propose a real-life implementation leveraging publicly available data on (i) real-time hospitalization, (ii) community mobility, (iii) social contacts, and (iv) socio-economic measures collected during the COVID-19 pandemic in Île-de-France, a region of France encompassing Paris with a population of approximately 12 million.

Insights. We leverage this framework to address the core research questions posed earlier.

1. **An idealized benchmark for the benefits of targeting.** We first quantify the *potential* benefits that could be achieved in an idealized setting where dual targeting of both activities and age-groups is possible. We find that dual targeting can lead to significant Pareto improvements. Specifically, the optimal ROLD policy that targets both age groups and activities Pareto-dominates the optimal ROLD policies that only target age groups, activities, or neither,

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

i.e., it leads to lower economic costs without increasing the number of pandemic deaths. The gains from targeting are also super-additive: targeting on both dimensions gains more than the sum of gains from unilateral targeting. In addition, the optimized dual-targeted policy also Pareto-dominates a number of benchmarks resembling policies implemented in practice. Finally, although not an explicit objective of the optimization, the dual-targeted ROLD policy also generally reduces the time in confinement for most population groups (and particularly for those groups that are most heavily confined) relative to less fine-grained policies.

2. **Interpretability.** We shed light on the mechanism and the cases when targeting leads to benefits. In term of mechanism, we find that optimized dual-targeted policies impose less confinement on group-activity pairs that generate a relatively high economic value prorated by activity-specific social contacts. This leads to complementary confinement schedules for different groups that reduce both the number of deaths and economic losses, with the important added benefit of not completely confining any group. We confirm these insights by running ROLD on a wide range of problem instances and generating a large dataset, which we use to train decision trees to predict the ROLD decisions based on simple, transparent features. The key feature that explains the targeting is the "econ-to-contacts-ratio," i.e., the *ratio of marginal economic value to total contacts* generated for a given (group, activity) pair.

   To understand *when* targeting leads to benefits, we train separate decision trees to predict the gains from dual targeting over less or no targeting. We find that R_0, the cost of death, the heterogeneity in the econ-to-contacts-ratio values, and the severity of the disease are the most important features, and that some of these can have a non-monotonic impact on the magnitude of the gains. In particular, our analysis leads to practical insights about the value of refined targeting depending on the pandemic's reproduction number and the policy maker's weighting of health over economic outcomes.

3. **Implementability.** Finally, we propose two new policies that do not explicitly target age and are thus simpler to implement and less contentious, but that nonetheless retain some of the superior health and economic outcomes of dual targeting. Both policies are inspired by real-world pandemic response measures. We first consider curfew policies that restrict activities at targeted times of day, *uniformly* for all age groups. Because different age groups perform distinct activities throughout the day, curfews have the potential to implicitly differentiate based on age. In the second proposal, we consider *recommending*, as opposed to enforcing, targeted activity levels separately for each age group, while accounting for imperfect population compliance. The results show that both proposals can significantly improve over activity-based targeting. Specifically, considering the gap between a dual-targeted and an activity-based targeted policy, we find that a curfew policy can close 24.3% (median value) of that gap and a

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

recommendation-based policy can lead to significant improvements even at relatively modest compliance levels, closing 10.4% (median value) of the gap at 40% compliance and 19.2% at 60% compliance.

2. Literature Review

The literature on pandemic response, particularly in the COVID-19 context, is already vast, so we focus our literature review on three key dimensions that our work most closely relates to.

Targeting. Paralleling our aforementioned real-life examples, several papers have studied targeted interventions. Kucharski et al. (2020), Prem et al. (2020), Di Domenico et al. (2020) recognize the importance of heterogeneity in the social contacts generated through activities and examine several interventions limiting them. Though some of the models here are age differentiated, targeting only happens through activities. Population group targeting, either through confinements, testing or vaccinations, has been investigated in Bastani et al. (2021), Acemoglu et al. (2020), Matrajt et al. (2020), Goldstein et al. (2021), Bertsimas et al. (2020), Favero et al. (2020), Birge et al. (2020), Chang et al. (2020), Evgeniou et al. (2020), Giordano et al. (2021). 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 (Acemoglu et al. 2020). A potential benefit of our approach is that, by exploiting complementarities between group and activity targeting, these higher risk groups may experience less confinement for the same level of aggregate deaths and economic losses.

Optimization of interventions. Our proposed optimization algorithm differs in several ways from existing approaches to assessing interventions via SEIR epidemiological models. A number of papers simulate a small number of candidate interventions, e.g., full lockdown versus school-only lockdown (Kucharski et al. 2020, Prem et al. 2020, Di Domenico et al. 2020, El Housni et al. 2020, Favero et al. 2020, Boloori and Saghafian 2020), restrict the candidates to a simple parametric class for which exhaustive search is computationally feasible (e.g., trigger policies based on hospital admissions as in Duque et al. 2020 or confirmed cases as Ahn et al. 2021), or use global optimization methods such as simulated annealing (Dutta et al. 2021). When considering a more complex policy space like in our targeting model, such approaches can lead to significantly sub-optimal results and misleading conclusions. Another stream borrows from the optimal control literature to design non-targeted control policies for single-group and single-activity SEIR dynamics (Bose et al. 2021, Pataro et al. 2021, Morris et al. 2021). Although optimization-based, the models there are simpler than our own and do not capture targeting. Birge et al. (2020) use formal optimization for location-based targeting, but in a one-shot model that does not differentiate age groups or activities and

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

does not account for time in the calculation of health or economic impact. The paper that is most related to ours is Bertsimas et al. (2020), which also proposes a multi-group SEIR formulation, in conjunction with an iterative coordinate descent algorithm to optimize vaccine allocations in a differentiated fashion. Although taking different approaches to doing so, both the algorithm proposed there and the ROLD heuristic crucially depend on solving linearized versions of the true SEIR dynamics which are tractable via commercial solvers. However, the model of Bertsimas et al. (2020) focuses on vaccine allocation decisions, whereas ours captures the dynamics of differential confinements and also allows activity-based targeting.

SEIR model calibration. Our work is also related to several other papers that have estimated SEIR parameters, particularly in the COVID-19 context. A number of papers estimate SEIR epidemiological parameters stratified by age groups – we use the estimates from the Île-de-France study in Salje et al. (2020). Another stream of papers, focusing on forecasting COVID-19 spread, estimate when and how the underlying SEIR parameters evolve in response to government interventions and changes in individual behavior, such as Perakis et al. (2021), Li et al. (2021). Lastly, we relate to the literature that has used Google and other mobility data to inform COVID-19 response strategies or to estimate the realized reductions in social contacts during COVID-19 (Dutta et al. 2021, Ilin et al. 2021, Wellenius et al. 2020, Cot et al. 2021, Xiong et al. 2020), as well as the literature on social contacts estimation (Béraud et al. 2015, Prem et al. 2017).

3. Model and Optimization Problem

We develop a controlled, multi-group SEIR model that includes time-dependent confinements that can be targeted based on *age groups* and types of *activities* that individuals engage in. The framework is flexible, captures resource constraints such as hospital or ICU capacity, and can be extended to capture other targeted interventions such as testing or vaccinations, as well as additional restrictions that make targeting more fair or practical such as integrality requirements, fairness constraints, etc. We discuss several of the potential extensions in Section 9.

3.1. 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 $\boldsymbol{v}_{i:j} := [\boldsymbol{v}_i, \ldots, \boldsymbol{v}_j]$ 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 $\boldsymbol{v}$.

3.2. Epidemiological Model and Controls

We rely on a modified version of the discretized SEIR (Susceptible-Exposed-Infectious-Recovered) epidemiological model (Anderson and May 1992, Prem et al. 2020, Salje et al. 2020) with multiple

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

population groups that interact with each other. In our case study we use nine groups $g \in \mathcal{G}$ determined by age and split in 10-year buckets, 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.

Compartmental Model and States. Figure 1 represents the compartmental model and the SEIR transitions for a specific group g. 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). $I_g(t)$ is further subdivided into $I_{j,g}(t)$ for $j \in \{a, ps, ms, ss\}$ to model different degrees of severity of symptoms: asymptomatic, paucisymptomatic, with mild symptoms, or with severe symptoms; thus, $I_g(t) = \sum_{j \in \{a,ps,ms,ss\}} I_{j,g}(t)$. The model also has states $R_g(t)$ (recovered but not confirmed as having had the virus), $R_g^{\mathrm{q}}(t)$ (recovered and confirmed as having had the virus), and $D_g(t)$ (deceased). Individuals with severe symptoms will need hospitalization, either in general hospital wards ($H_g(t)$) or in intensive care units ($ICU_g(t)$). All the states represent the number of individuals in a compartment of the model in the *beginning* of the time period.

Susceptible individuals get infected and transition to the exposed state at a rate determined by the number of social contacts and the transmission rate $\beta(t)$. Exposed individuals transition to the infectious state at a rate σ and infectious individuals transition out of the infectious state at a rate μ. We assume that an infectious individual in group g exhibits symptoms of degree j with probability $p_{j,g}$. An infectious individual needs to be hospitalized in general hospital wards or ICU with probability p_g^H and p_g^{ICU}, respectively, where $p_{ss,g} = p_g^H + p_g^{ICU}$. On average, patients treated in general hospital wards (ICU) spend λ_H^{-1} (λ_{ICU}^{-1}) days in the hospital (ICU). An infectious individual with severe symptoms in group g deceases (recovers) with probability p_g^D ($p_g^R = 1 - p_g^D$).

We keep track of all living individuals in group g who are not confirmed to have had the disease $N_g(t) := S_g(t) + E_g(t) + I_g(t) + R_g(t)$, and 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 (across groups) at time $0 \le t \le T$. We denote the number of compartments by $|\mathcal{X}|$, so the dimension of $\boldsymbol{X}_t$ is $|\mathcal{G}||\mathcal{X}| \times 1$.

Controls. Individuals interact in activities $\mathcal{A} = \{\mathsf{work}, \mathsf{transport}, \mathsf{leisure}, \mathsf{school}, \mathsf{home}, \mathsf{other}\}$. These interactions generate *social contacts* that 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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Figure 1 Compartmental SEIR model for a specific group g with transition rates.

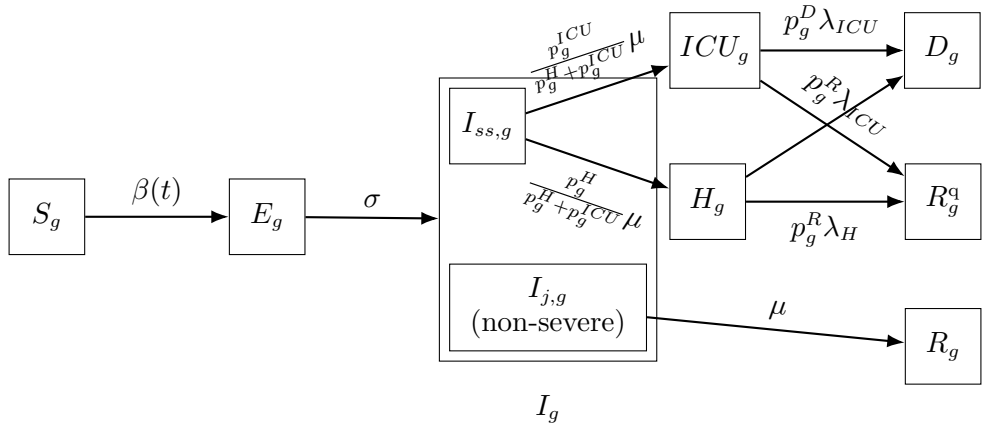

irrespective of confinement policy.[1] We denote the vector of all activity levels for group g at t by $\boldsymbol{\ell}_g(t) = [\ell_g^a(t)]_{a\in\mathcal{A}}$, and we also refer to $\boldsymbol{\ell}_g(t)$ as confinement decisions when no confusion can arise.

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

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

where $C_{g,h}^a$ denote the mean number of daily contacts in activity a under normal course (i.e., without confinement) and $\alpha \in \mathbb{R}$ is a social mixing parameter that captures the elasticity of social contacts to activity levels. This parametrization is similar to a Cobb-Douglas production function (Mas-Colell et al. 1995), using the activity levels as inputs and the number of social contacts as output. We retrieve values for $C_{g,h}^a$ from the data tool of Wille et al. (2020), which is based on the French social contact survey data in Béraud et al. (2015), and we estimate α from health outcome data (French Government 2020) and Google mobility data (Google 2020), as described in Section EC.4.2.

Let $\boldsymbol{u}_t = \big[\boldsymbol{\ell}_g(t)\big]_{g\in\mathcal{G}}$ denote the vector of all decisions at time $t \in \{0, 1, \ldots, T-1\}$, i.e., the confinement decisions for all the groups. 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 $|\mathcal{G}||\mathcal{U}| \times 1$.

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

3.3. Resources and Constraints

We use $K^{\mathrm{H}}(t)$ ($K^{\mathrm{ICU}}(t)$) to denote the capacity of beds in general hospital wards (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 optimizing over such decisions, we choose to not consider this dimension of targeting because it can be extremely contentious in practice. Instead, we implement a proportional rule that allocates any remaining hospital and ICU capacity among patients from all age groups proportionally to the number of cases requiring admission from each group. More formally, with $B_g^H(t)$ denoting the number of patients from age group g who are denied admission to general hospital wards in period t, the proportional rule is:

$$B_g^H(t) = \underbrace{\frac{\text{H-flow}_g(t)}{\sum_h \text{H-flow}_h(t)}}_{\text{proportion of inflow into H from } g} \cdot \left(\underbrace{\sum_h \text{H-flow}_h(t)}_{\text{total inflow into H}} - \underbrace{\left(K^{\mathrm{H}}(t) - \sum_h \left(1 - \lambda_h^H\right) H_h(t) \right)}_{\text{available beds in H}} \right)^+, \tag{2}$$

where $\text{H-flow}_g(t) := \mu \cdot p_g^H \cdot I_g(t)$, and a similar rule holds for patients requiring admission to the ICU. We assume that all patients that are denied admission die immediately.

We can now write a complete set of discrete dynamical equations for the controlled SEIR model ((EC.1)-(EC.9) in Appendix EC.1) and summarize these using the function

$$\boldsymbol{F}_t(\boldsymbol{X}_t, \boldsymbol{u}_t) := \frac{\Delta \boldsymbol{X}_t}{\Delta t}, \tag{3}$$

where $\Delta \boldsymbol{X}_t := \boldsymbol{X}_{t+1} - \boldsymbol{X}_t$. Additionally, we also include the following constraints:

$$\text{(“Hospital capacity”)} \quad \sum_g H_g(t) \le K^{\mathrm{H}}(t), \quad \forall t \tag{4}$$

$$\text{(“ICU capacity”)} \quad \sum_g ICU_g(t) \le K^{\mathrm{ICU}}(t), \quad \forall t \tag{5}$$

$$\text{(“Fractional, non-negative activity levels”)} \quad 0 \le \ell_g^a(t) \le 1, \quad \forall g, a, t \tag{6}$$

We denote by $\mathcal{C}(\boldsymbol{X}_t)$ the feasible set described by (4)-(6)for the vector of decisions $\boldsymbol{u}_t$ at time t.

3.4. Objective

Our objective captures two criteria. The first quantifies the total deaths directly attributable to the pandemic, which we denote by $\mathsf{Total\ Deaths}(\boldsymbol{u}_{0:T-1}) := \sum_{g \in \mathcal{G}} D_g(T)$ to reflect the dependency on the specific policy $\boldsymbol{u}_{0:T-1}$ followed. The second criterion captures the economic losses due to the pandemic, denoted by $\mathsf{Economic\ Loss}(\boldsymbol{u}_{0:T-1})$. These stem from three sources: (a) lost productivity due to confinement, (b) lost productivity during the pandemic due to individuals being quarantined, hospitalized, or deceased, and (c) lost value after the pandemic due to deaths (as deceased individuals no longer produce economic output even after the pandemic ends).

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

To model (a), we assign a daily economic value $v_g(\boldsymbol{\ell})$ to each individual in group g that depends on the activity levels $\boldsymbol{\ell} := [\boldsymbol{\ell}_g]_{g\in\mathcal{G}}$ 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 assume that an individual who is in quarantine, hospitalized, or deceased, generates no economic value. At the same time, we assume that individuals in R_g^{q} generate economic value as they would under no confinement. For (c), we determine the wages that a deceased 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 economic loss is the difference between the economic value that would have been generated during the pandemic under a "no pandemic" scenario (V) and the value generated during the pandemic, plus the future economic output lost due to deaths.

$$\text{Economic Loss}(\boldsymbol{u}_{0:T-1}) := V - \sum_{t=0}^{T-1}\sum_{g\in\mathcal{G}}\Big(v_g\big(\boldsymbol{\ell}(t)\big)\cdot N_g(t) + v_g(\mathbf{1})\cdot R_g^{\text{q}}(t)\Big) + \sum_{g\in\mathcal{G}} v_g^{\text{life}}\cdot D_g(T), \tag{7}$$

All the details of the economic modelling are deferred to Appendix EC.2.

To allow policy makers to weigh the importance of the two criteria, we associate a cost χ to each death, which we express in multiples of 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 ($\chi \to \infty$).

3.5. Optimization Problem

The optimization problem we solve is to find control policies for confinement that minimize the sum of mortality and economic losses[2] subject to the constraints that (i) the state trajectory follows the SEIR dynamics, and (ii) the controls and states respect the capacity and feasibility constraints discussed above. Formally, we solve:

$$\min_{\boldsymbol{u}_{0:T-1}} \text{Economic Loss}(\boldsymbol{u}_{0:T-1}) + \chi\cdot\text{Total Deaths}(\boldsymbol{u}_{0:T-1}) \tag{8}$$

$$\text{s.t.}\quad \boldsymbol{X}_{t+1} = \boldsymbol{X}_t + \boldsymbol{F}_t\left(\boldsymbol{X}_t, \boldsymbol{u}_t\right), \quad \forall 0 \le t \le T-1 \tag{9}$$

$$\boldsymbol{u}_t \in \mathcal{C}(\boldsymbol{X}_t), \quad \forall 0 \le t \le T-1. \tag{10}$$

[2] As in Acemoglu et al. (2020), we focus on characterizing the frontier between deaths and economic losses, which can be obtained by varying the cost of death χ and minimizing the objective function for each value of χ; this makes our results informative for policy makers who may have differing views on the right value for such a key parameter.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

4. Algorithm: Re-Optimization with Linearized Dynamics

Solving problem (8)-(10) to optimality unfortunately requires solving intractable optimization problems. To see this, note that the key term in the dynamics of any SEIR-type model is the rate of new infections, which involves multiplying the current susceptible population with the infected population. This introduces non-linearity in the state trajectory; for instance, our dynamic for the evolution of the susceptible population in group g from (EC.2) reads:

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

Expanding out $S_g(t)$ produces a complex, non-convex dependency on the past decisions $\ell(\tau)$ for $0 \le \tau \le t-1$, which makes the resulting problem intractable via convex optimization.

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

4.1. Linearization and Optimization

The key idea 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 decisions for all $k, \ldots, T$, and only implement the decisions for the current time step k.

We first describe the linearization procedure. Recall that the true evolution of our dynamical system is given by (3). 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 sequence $\hat{\boldsymbol{u}}_{k:T-1}$ and let $\hat{\boldsymbol{X}}_{k:T}$ denote the resulting nominal system trajectory under the true dynamic (3) and under $\hat{\boldsymbol{u}}_{k:T-1}$. We approximate the original dynamics through a Taylor expansion around $(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)$:

$$\frac{\Delta \boldsymbol{X}_t}{\Delta t} \approx \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) + \nabla_X \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t) + \nabla_u \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t)(\boldsymbol{u}_t - \hat{\boldsymbol{u}}_t), \tag{11}$$

where $\nabla_X \boldsymbol{F}_t$ and $\nabla_u \boldsymbol{F}_t$ 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 (11) is indeed a linear expression of $\boldsymbol{X}_t$ and $\boldsymbol{u}_t$. By induction, every state $\boldsymbol{X}_t$ under dynamic (11) will be a linear function of $\boldsymbol{u}_\tau$ for $\tau < t$, and all the constraints will also depend linearly on the decisions.

In a similar fashion, we also linearize the objective (8). Since $v_g(\boldsymbol{\ell}(t))$ is linear in $\boldsymbol{u}_t$ for all $t = 0, \ldots, T-1$, the objective contains *bilinear* terms and can be written compactly as:

$$\sum_{t=0}^{T-1} \left(\boldsymbol{u}_t^\top \boldsymbol{M} \boldsymbol{X}_t + \boldsymbol{\gamma}^\top \boldsymbol{X}_t \right) + \boldsymbol{\eta}^\top \boldsymbol{X}_T, \tag{12}$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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 $|\mathcal{G}||\mathcal{X}| \times 1$ (detailed expressions are available in Appendix EC.3). By linearizing this using a Taylor approximation, we consider the following objective instead:

$$\sum_{t=0}^{T-1} \left(\hat{\boldsymbol{u}}_t^\top \boldsymbol{M} \hat{\boldsymbol{X}}_t + \hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top (\boldsymbol{u}_t - \hat{\boldsymbol{u}}_t) + \hat{\boldsymbol{u}}_t^\top \boldsymbol{M} (\boldsymbol{X}_t - \hat{\boldsymbol{X}}_t) + \boldsymbol{\gamma}^\top \boldsymbol{X}_t \right) + \boldsymbol{\eta}^\top \boldsymbol{X}_T, \tag{13}$$

which depends linearly on all the decisions $\boldsymbol{u}_0, \ldots, \boldsymbol{u}_{T-1}$.

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

1. Given the current state $\boldsymbol{X}_k$ and a nominal control sequence $\hat{\boldsymbol{u}}_{k:T-1}^{(k)}$ for all remaining periods, calculate a nominal system trajectory $\hat{\boldsymbol{X}}_{k:T}$ under the true dynamic in (3). (The nominal control sequence is set to a solution obtained by a gradient descent method at $k = 0$, and to the algorithm's own output from period $k-1$ for periods $k > 0$, per Step 4 below.)
2. Use (11) to approximate the state dynamic around the nominal trajectory $\hat{\boldsymbol{X}}_t$ and use (13) to approximate the objective-to-go function over the remaining periods $t \in \{k, \ldots, T\}$.
3. Solve the linear program to obtain decision variables $\boldsymbol{u}_{k:T-1}^{*,(k)}$ that maximize the linearized objective-to-go subject to all the relevant linearized constraints.
4. Set the nominal control sequence for the next time step as $\hat{\boldsymbol{u}}_{k+1:T-1}^{(k+1)} := \boldsymbol{u}_{k+1:T-1}^{*,(k)}$.
5. Update the states using the optimal control $\boldsymbol{u}_k^* := \boldsymbol{u}_k^{*,(k)}$ and the true dynamic in (3), i.e. $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k + \boldsymbol{F}_k(\boldsymbol{X}_k, \boldsymbol{u}_k^*)$.

The linearization-optimization procedure described above is run for all periods $k = 0, \ldots, T-1$ sequentially to output a full control policy $\boldsymbol{u}_{0:T-1}^*$.

Trust region implementation. In our experiments, we have found that the linearized model described in (11) may diverge significantly from the real dynamical system when the optimized controls $\boldsymbol{u}_{k:T-1}^{*,(k)}$ determined in Step 3 diverge sufficiently from the nominal controls $\hat{\boldsymbol{u}}_{k:T-1}^{(k)}$ considered in the linearization in Step 2. This can lead to a large sensitivity in performance to the initialization used in the very first step; for example, if the Taylor approximation were constructed around a policy of full confinement, the linearized model could systematically underestimate the number of infections and deaths created when considering more relaxed confinements.

We overcome this by employing an iterative procedure inspired by a trust region optimization method. The key idea is to avoid the large approximation errors by running the linearization-optimization procedure iteratively within each time step k, with each iteration only being allowed to take a small step towards the optimum within a trust region of an ϵ-ball around the nominal control sequence $\hat{\boldsymbol{u}}_{k:T-1}^{(k)}$, and the updated optimized control sequence of each iteration being used as a nominal sequence for the next iteration. 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.

Further algorithmic details for ROLD are provided in Appendix EC.3.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

5. Île-de-France Calibration and Experimental Setup

To complete the description of our framework, we now summarize our approach for calibrating our model using real-world publicly available data on (i) community mobility, (ii) social contacts, (iii) health outcomes, and (iv) economic output for the Île-de-France region of France.

5.1. Parametrization and Calibration

We adopt values for disease progression parameters for Île-de-France directly from the study by Salje et al. (2020); we adopt values for the number of social contacts across different age groups and activities from the study of Béraud et al. (2015) focused on France, using the tool by Wille et al. (2020); and we use Google mobility data (Google 2020) for the Île-de-France region to approximate the mean effective lockdowns for all activities during the horizon of interest. We then estimate any remaining parameters by minimizing the sample-average approximation of an error metric derived by comparing our model's predictions (on several sample paths) with time-series data from the French Public Health Agency (French Government 2020) on hospital and ICU utilization and deaths. We calibrate our economic model using data on full time equivalent wages and employment rates from the French National Institute of Statistics and Economic Studies, and sentiment surveys on business activity levels during confinement from the Bank of France. We provide all the details for calibration and parameter specification in Appendix EC.4. We report experimental results from sensitivity and robustness analyses on the fitted parameters in Appendix EC.6.

5.2. Model Validation

Table EC.5 in the Appendix summarizes the parameter values obtained from our calibration procedure, and Figure 2 compares the fitted model's predictions with the reported values for hospital beds utilization, ICU beds utilization, and cumulative deaths. To further validate our model, we

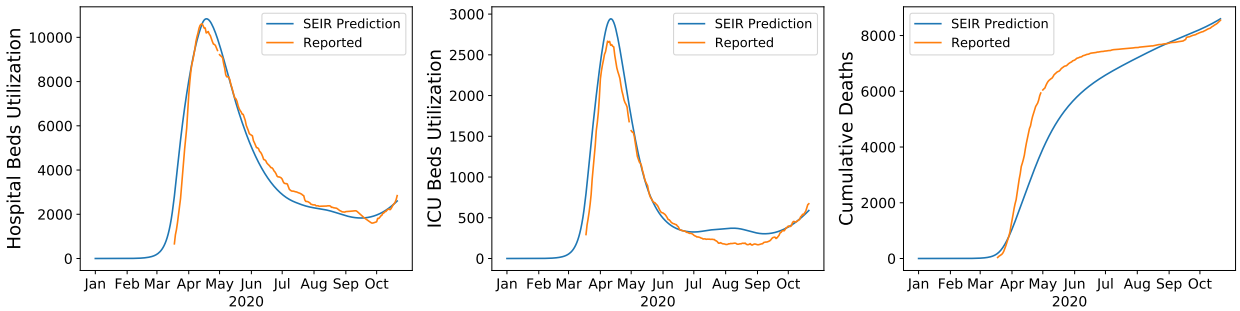

Figure 2 **Predictions of the fitted SEIR model vs. reported values by the French Public Health Agency for Hospital Beds Utilization, ICU Beds Utilization and Cumulative Deaths.**

also assess the goodness of fit *out of sample*. In particular, for each date t from a set of four dates in 2020, we calibrate the SEIR model using data up until day $t-14$, we simulate the calibrated

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

model up until day t, and we compare the model predictions with the reported data in the interval $[t-13, t]$. Figure EC.1 shows that the predictions of the calibrated SEIR model for all three metrics of interest stay close to the reported values in the out-of-sample validation set. As expected, the fit worsens when the model is estimated with less data (e.g., it is worst for the earliest of the four dates shown).

5.3. Experimental and Optimization Setup

We run experiments over a range of values for our model parameters, summarized in Table 1. For parameters for which multiple values are used in our experiments, the "Baseline Value" column reports the values used in results in the main paper and the Appendix, unless specified otherwise. In particular, we use a baseline capacity of 2900 ICU beds in Île-de-France, and experiment with ICU capacities that range from 2000 to 3200 beds.[3] We use an infinite capacity for general hospital wards. We optimize decisions starting on October 21 2020. We allow confinement decisions to change every two weeks.

Optimization Horizon. We use an optimization horizon of $T' = 90$ days in the experiments reported in the main paper (and allow up to 360 days in additional experiments). To optimize confinements using ROLD and other policies on days $0, 1, \ldots, T'-1$, we set the total time horizon to be $T = T' + 14$ and further constrain the policies to be fully open on days $T', \ldots, T'+13$, while assuming no infections are possible starting on day T'. We do this to mitigate possible end-of-horizon effects: any deaths and loss of economic value between day $T'+1$ and $T'+14$ will still count towards the objective, so ROLD cannot quite allow for too many infections towards the end of the T'-day optimization horizon. We discuss further techniques for mitigating end-of-horizon effects in Section 9.

ROLD Variants. 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.[4] 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. The four versions of ROLD we consider are no targeting whatsoever (NO-TARGET),

[3] ICU capacity in Île-de-France was 1,200 beds at the beginning of the pandemic (Godeluck 2020) and was enhanced during the crisis, with a reported peak ICU utilization during the spring of 2020 of 2,668 beds (Sterlé et al. 2020).

[4] 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. 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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table 1 Parameter values for experimental and optimization setup. The parameters $\nu^{\text{other activities}}, r, f_g$ and θ related to our economic model are defined in Appendix EC.2

Parameter Description	Notation	Values in Experiments	Baseline Value
Cost of death	χ	30 values in $[0, 1000]\times$ GDP per capita of France	
GDP per capita of France		€37199.03	
ICU capacity	K^{ICU}	2000, 2300, 2600, 2900, 3200	2900
Hospital capacity	K^{H}	∞	
Sensitivity of econ. value on confinement	$\nu^{\text{other activities}}$	0, 0.1, 0.2	0.1
Discount rate (used to calculate $v_g(\boldsymbol{\ell})$ and v_g^{life})	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	

targeting age groups only (AGE), targeting activities only (ACT), or targeting both (AGE-ACT, or simply ROLD when no confusion can arise). To obtain each policy, we run suitably constrained versions of the ROLD optimization problem initialized using the solution of a gradient descent algorithm (details in Appendix EC.3.5).

6. How Large Are the Potential Gains from Dual Targeting?

We next apply our framework to the Île-de-France context to quantify the magnitude of the gains from targeting in an idealized setting where dual targeting can be implemented. This will then serve as a useful benchmark against which to assess more implementable targeted policies.

To isolate the benefits of each type of targeting, we compare the four versions of ROLD that differ in the level of targeting allowed, as described in Section 5.3. Figure 3a records each policy's performance in several problem instances, where each problem instance corresponds to a different value for 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%-36.4%) in AGE, by EUR 0.6B-2.2B (6.5%-52.3%) in ACT, and by EUR 2.3B-5.4B (24.4%-80.6%) in AGE-ACT. This Pareto-dominance is unexpected since it is not explicitly required 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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

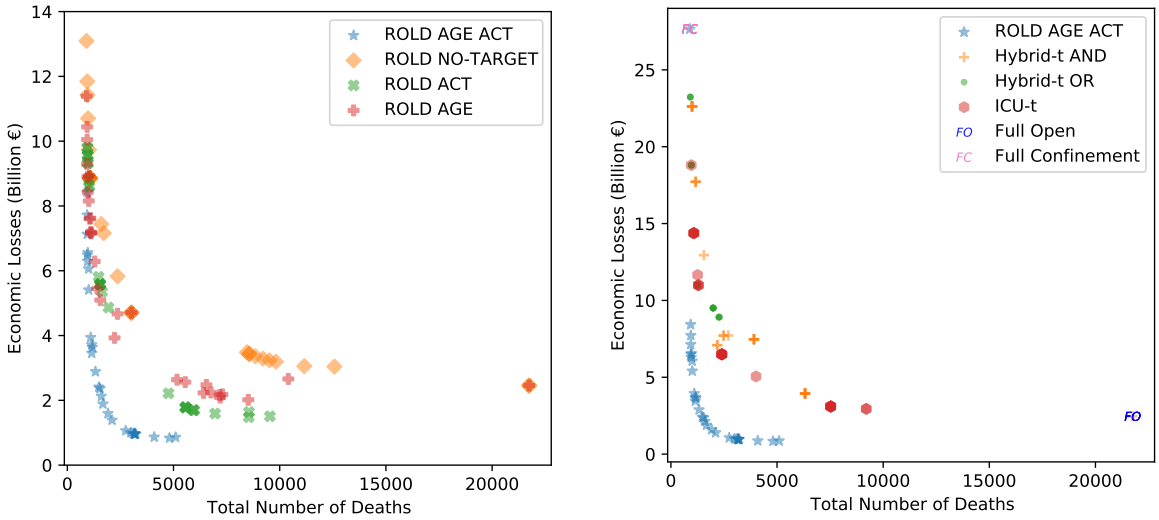

(a) (b)

Figure 3 **The total number of deaths and the economic losses generated by targeted ROLD policies and 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 parameterized by the cost of death χ. We include 30 distinct values of χ from 0 to 1000$\times$, and panel (b) also includes a very large value ($\chi = 10^{16}\times$, for which ROLD AGE-ACT recovers the full confinement policy).**

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 (Figure 4a). In contrast and crucially, AGE-ACT Pareto-dominates *all* other policies (and also dominates in terms of the total loss objective). Moreover, targeting both age groups and activities 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* (Figure 5). This suggests that substantial complementarities may be unlocked through *dual* targeting that may not be available under less granular targeting. These results are robust under more problem instances (Appendix Section EC.6.1).

To gain additional perspective on these idealized gains, we also compare ROLD AGE-ACT with various practical benchmark policies in Figures 3b, 4b. Benchmarks ICU-t and Hybrid-t AND/Hybrid-t OR mimic implementations in the U.S. Austin area (Duque et al. 2020) and, respectively, France (Lehot and Borgne 2020). These policies switch between a stricter and a relaxed

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

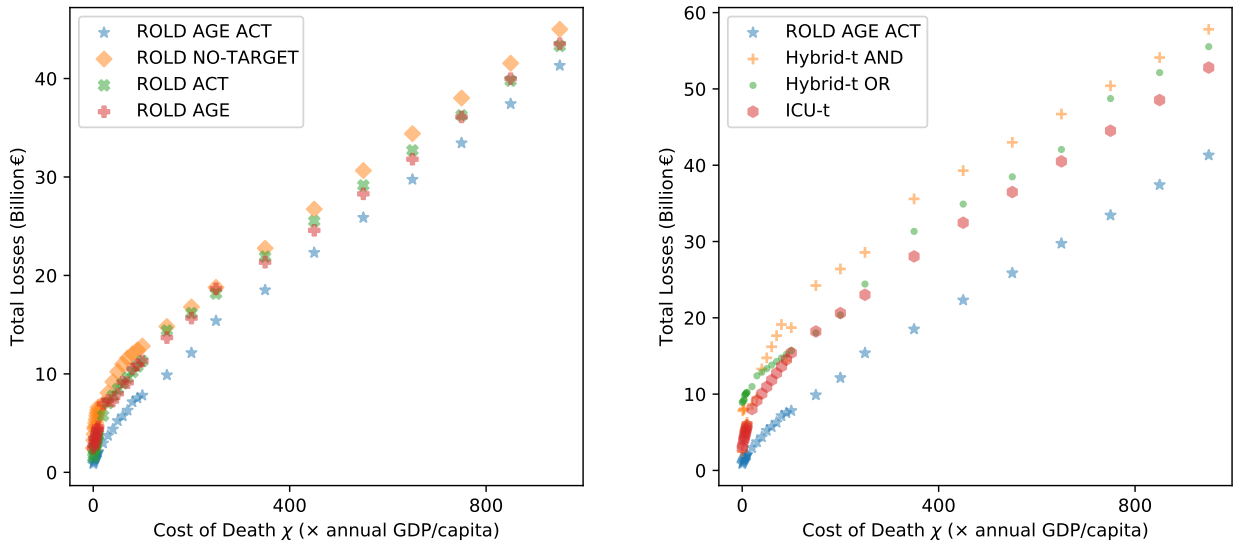

(a) (b)

Figure 4 **Total losses generated by targeted ROLD policies and by the benchmark policies, at different values of the cost of death χ. 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 parameterized by the cost of death χ. We include 30 distinct values of χ from 0 to 1000×.**

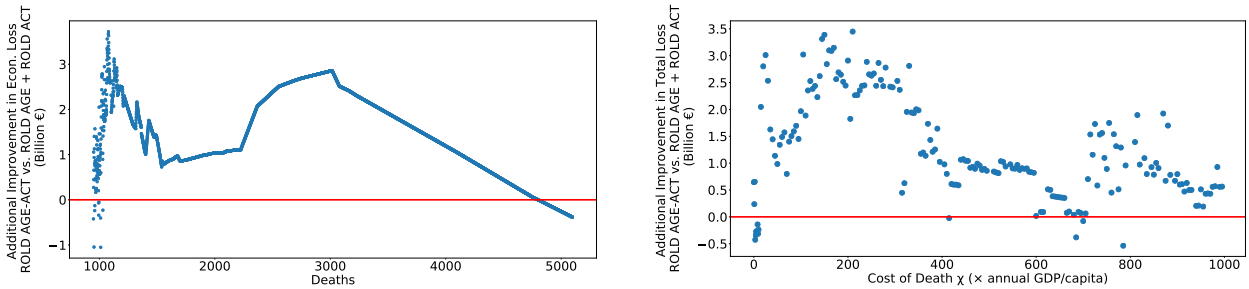

Figure 5 **The super-additivity of ROLD AGE-ACT. The figures compare the improvement from AGE-ACT with the sum of the improvements from AGE and ACT. All improvements are with respect to NO-TARGET. The left panel compares the improvements in economic losses for the same number of deaths. The right panel compares the improvements in total losses for the same cost of death χ.**

confinement level based on conditions related to hospital admissions and occupancy and the rate of new infections (Appendix EC.5). We optimize the parameters of these benchmark policies and

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

compare with the best-performing model in each class. We also consider two extreme benchmarks corresponding to enforcing "full confinement" (FC) or remaining "fully open" (FO), which can be expected to perform well when completely prioritizing the number of deaths (i.e., $\chi \to \infty$) or the economic losses (i.e., $\chi = 0$), respectively.

ROLD Pareto-dominates all these benchmarks, decreasing economic losses by EUR 5.8B-16.9B (66.6%-84.0%) relative to Hybrid-t AND, by EUR 7.3B-11.8B (55.5%-85.7%) relative to Hybrid-t OR, and by EUR 5.7B-11.8B (55.5%-81.9%) 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 27.7B; for a sufficiently low χ, ROLD actually Pareto-dominates the FO policy, reducing the number of deaths by 16,642 (76.6%) and reducing economic losses by EUR 1.6B (64.9%). 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.

Another possible benefit of finer targeting is that it could reduce the time in confinement for all groups. To check this, we calculate the time-average activity level for each age group under each ROLD policy, averaged over the activities relevant to that age group and weighting each activity by the participation rate of that age group in that activity (details provided in Appendix EC.6.2). The results are visualized in Figure 6. The dual-targeted AGE-ACT policy has higher population-wide activity levels relative to all other policies and maintains higher activity levels for most age groups, with systematically higher levels for the age groups that are *most confined*. In particular, for the AGE-ACT policy the group with the lowest median activity level (of 81.7%) is the 40-49 y.o., whereas under the AGE (ACT) policy, that group is the 80+ (60-69) y.o., with median activity level of 34.2% (75.3%). This is a reasonable requirement when considering the "fairness"

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

of confinement policies.[5] Moreover, even when AGE-ACT confines certain age groups more than other policies, it never does so by much: its activity levels are within 5% (in absolute terms) from the levels achieved by ACT (respectively, AGE and NO-TARGET) for *every* age group in 96% (57%, 98%) of all instances, and within 10% from ACT (respectively, AGE and NO-TARGET) for every age group in 98% (88%, 98%) of all instances.

These results suggest that dual targeting has the potential to improve outcomes along multiple dimensions: in terms of *both* health and economic outcomes, and by generating more socially acceptable confinements, through a reduced and more "fair" distribution of confinement time.

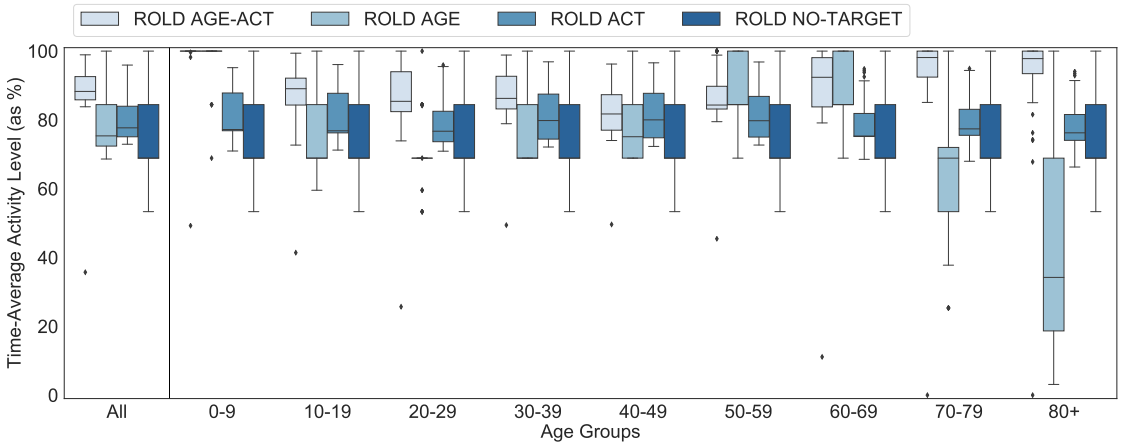

Figure 6 **Time-average activity level for the targeted ROLD policies. Each boxplot depicts the time-average activity level under the respective policy averaged over the activities relevant to that age group and weighting activities by participation rates, for different problem instances parameterized by χ. The first boxplot on the left depicts activity levels aggregated over all the age groups.**

7. How Do Gains Arise from Dual Targeting?

That targeting can generate such improvements along multiple metrics is quite unexpected because this is not something that the ROLD framework explicitly optimizes for (recall that the ROLD objective only involves a total loss metric). To better understand how gains could arise from

[5] Maximizing the time-average activity levels for the group(s) that are most confined can be seen as a reasonable fairness requirement, e.g., consistent with max-min fairness and Rawlsian justice (Young 1994).

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

targeting, we now examine the structure of the optimal dual-targeted ROLD decisions and explain these as well as their associated gains in terms of a set of interpretable features.

7.1. Dual Targeting Generates Complementarities

By examining the optimal ROLD AGE-ACT decisions, we first note that the policy maintains high activity levels for (age group, activity) pairs with high economic value and few social contacts. This can be visualized in Figure 7, which plots the activity levels for different (age group, activity) pairs divided into three buckets (high, medium, low) depending on their econ-to-contacts-ratio. This is defined as the marginal economic value divided by the total social contacts generated by a group in an activity, i.e.,

$$\text{econ-to-contacts-ratio}(g,a,t) := \frac{\partial \mathsf{Economic\ Value}(t)/\partial \ell_g^a(t)}{N_g(t)\sum_{h\in\mathcal{G}} C_{h,g}^a},$$

where $\mathsf{Economic\ Value}(t) = \sum_{g\in\mathcal{G}} \left(v_g\left(\boldsymbol{\ell}(t)\right)\cdot N_g(t) + v_g(\mathbf{1})\cdot R_g^{\mathrm{q}}(t)\right)$. Similarly, Figure 8 shows that ROLD prioritizes activity in transport, school, work, other, and then leisure, in that order, in accordance with the relative econ-to-contacts-ratio of these activities.[6]

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, ROLD coordinates confinements to account for this (see Figure 9): groups 20-69 y.o. remain more open in work but face confinement in leisure for up to the first ten weeks, whereas groups 0-19 and 70+ y.o. remain quite open in leisure but have reduced activity in work. These complementary confinement schedules allow ROLD to reduce both the number of deaths and economic losses, with the important added benefit that no age group is completely confined. (For an additional visualization, Figure EC.4 in the Appendix depicts the ROLD activity levels for a specific value of the cost of death χ.)

[6] These figures also show that the policy generally imposes stricter confinements for larger cost of death χ. We also find that the policy tends to relax confinements through time — see Figure EC.3 in the Appendix.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

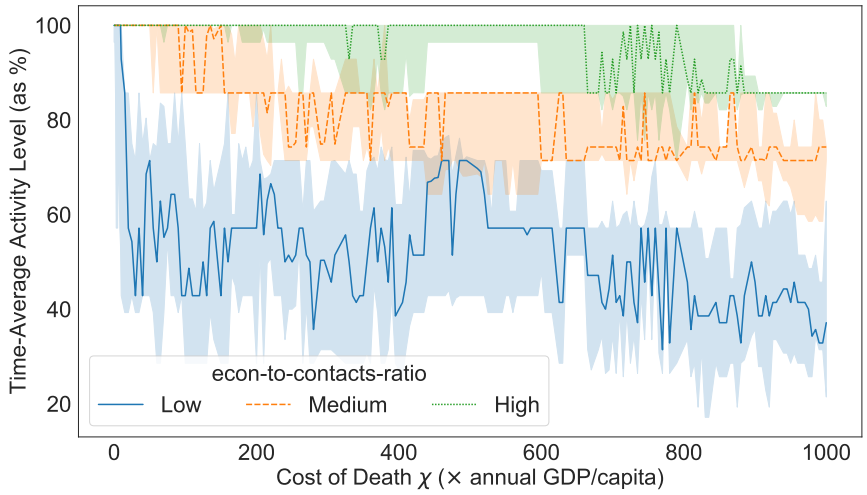

Figure 7 **The time-average activity levels of optimized ROLD AGE-ACT policies for different (age group, activity) pairs, bucketed in three equally-sized buckets according to their econ-to-contacts-ratio. The lines indicate the median for each bucket while the bands indicate interquartile ranges, where for each value of the cost of death χ, the statistics are taken over the relevant (age group, activity) pairs.**

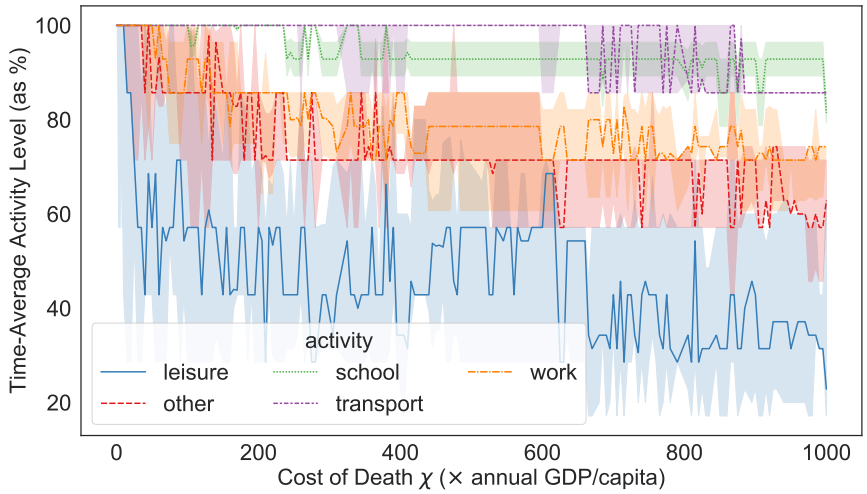

Figure 8 **The time-average activity levels of optimized ROLD AGE-ACT policies for different activities. The lines indicate the median for each activity, while the bands indicate interquartile ranges, where for each value of the cost of death χ the statistics are taken over the relevant age groups for that activity.**

7.2. Understanding Dual-Targeted Policies and Their Gains

To gain a better understanding of the ROLD policy and the potential benefits of dual targeting, we also take a different approach and train interpretable machine learning models (regression trees)

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

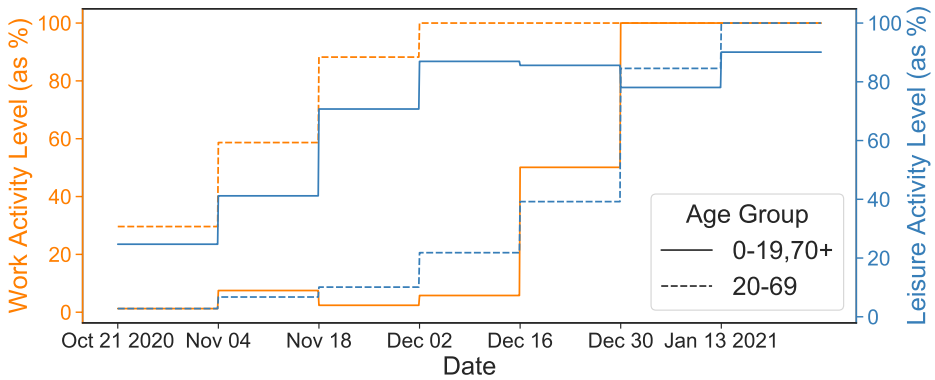

Figure 9 **The activity levels of optimized ROLD AGE-ACT policies for work and leisure, and for the 20-69 y.o., vs. the 0-19 and 70+ y.o. age groups. The lines indicate mean activity levels, where for each time t the mean is taken across different values for the cost of death χ and across the respective relevant age groups. (The figure is best viewed in color.)**

to predict (i) the optimal ROLD decisions and (ii) the gains from dual targeting, as a function of several salient epidemiological and economic features.

More specifically, we generate problem instances with randomly drawn values from a wide range for the transmission rate β (equivalently, for the basic reproduction number R_0), the disease severity ("severity-scaling" — a scaling factor for the probabilities of an infectious individual having severe symptoms $\{p_{ss,g}\}_{g\in\mathcal{G}}$), the social mixing parameter α (or equivalently, the elasticity of social contacts to activity levels[7]), the cost of death χ, the ICU capacity K^{ICU}, and the economic model parameters $\nu^{\text{work}}, \nu^{\text{other activities}}, \nu^{\text{fixed}}$. We create 14,039 problem instances by drawing each parameter independently and uniformly at random from a specified range. Details for the parameter variation are provided in Appendix EC.6.4 and in particular in Table EC.9.

Interpreting the dual-targeted policy. We create a training set of 5,524,785 samples — with one sample for every 14-day period, each age group g, and each activity a relevant to that age group. We fit a decision tree to predict the ROLD AGE-ACT-optimized activity levels $\ell_g^a(t)$ based on 20 features. We include as features several important problem parameters (e.g., the basic reproduction

[7] This is defined as the percentage drop in total contacts when activity levels drop by 50%, and is equal to $1-4^{-\alpha}$, where α is the social mixing parameter.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

number R_0, the cost of death, the elasticity of social contacts to activity levels, the ICU capacity) as well as derived features based on parameters and SEIR state values (e.g., econ-to-contacts-ratio, ICU utilization and admission rates, infection rates). Some of these features are targeted, meaning they take different values for different age groups or activities (e.g., econ-to-contacts-ratio), whereas others are non-targeted, such as R_0, the time t or the cost of death χ. The details of the fitting procedure are in Appendix EC.6.4.1, and the full set of features is in Tables EC.10 and EC.11.

Figure 10 displays the depth-four tree obtained by fitting using the entire feature set, together with the resulting root-mean-squared-error (RMSE). This simple tree predicts the optimal ROLD activity levels reasonably well (RMSE of 0.31) and confirms 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 the tree and in several sub-trees, with higher ratios always leading to higher activity levels. The tree also confirms that R_0 and time are relevant, with the optimal ROLD policy enforcing stricter confinements for larger values of R_0 and in earlier time periods. Another relevant feature is R-perc, which quantifies the total number of individuals in a recovered state as a fraction of the overall population; this is a natural measure of the level of herd immunity and the ROLD policy leverages it intuitively, allowing more activity at larger levels of herd immunity.

We also quantify the importance of the various features by calculating their permutation importance scores, a commonly used metric that measures importance as the increase in model prediction error after the values of the respective feature are randomly permuted. For trees of depth four, the scores for econ-to-contacts-ratio, R_0, time, R-perc(t), and $N_g(t)$ are respectively 0.163, 0.065, 0.037, 0.008, 0.007; and all the other features have scores of zero. The econ-to-contacts-ratio feature thus has by far the largest permutation importance score of all features. Figure EC.6 in the Appendix confirms similar results for a tree of depth 10, and Figure EC.7 documents an additional exercise that further shows econ-to-contacts-ratio as the most salient of all targeted features. Appendix EC.7 contains additional results to complement this section, including a theoretical justification for the salience of the econ-to-contacts-ratio derived in a simplified model.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

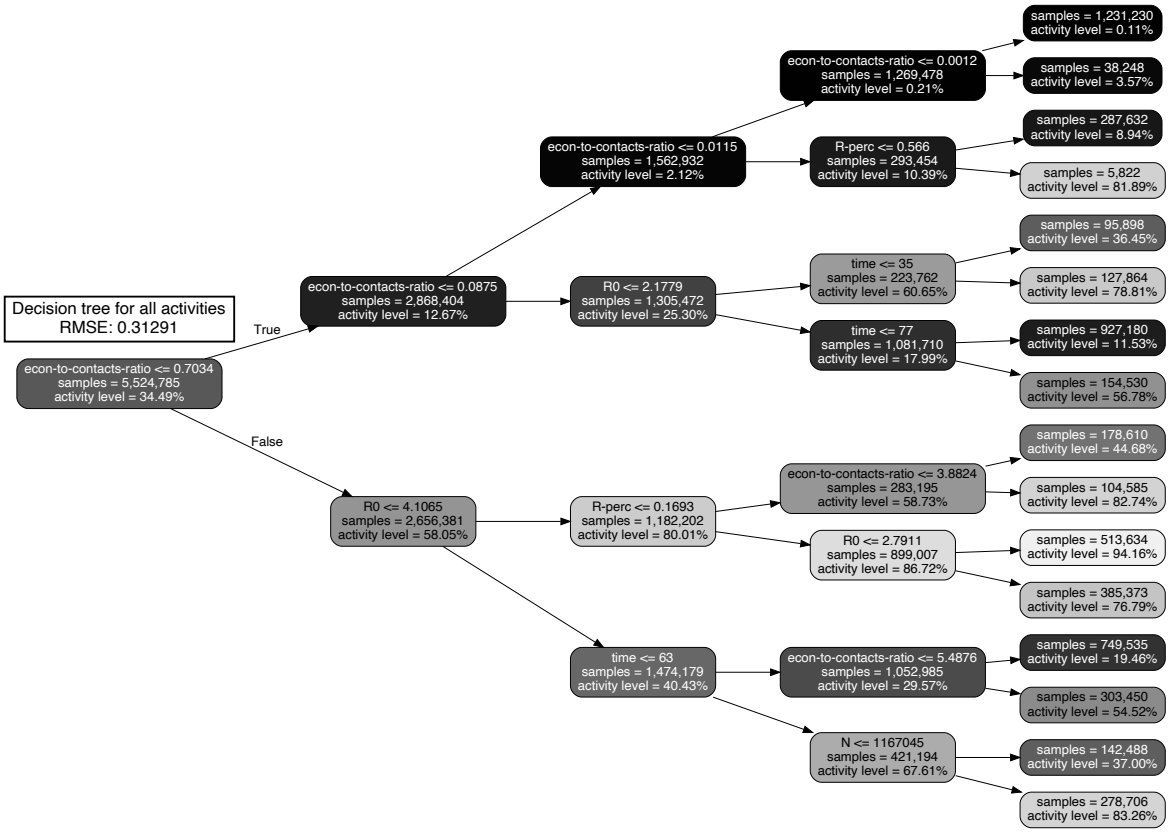

Figure 10 **Decision tree of depth four approximating the optimized ROLD AGE-ACT decisions trained on 14,039 problem instances with an optimization horizon of $T' = 90$ days (a total of 5,524,785 samples). Nodes are color-coded based on activity level, with darker colors corresponding to stricter confinement.**

Interpreting the gains from dual targeting. For each of the 14,039 problem instances, we retrieve the optimal policies under ROLD AGE-ACT, ACT, and NO-TARGET, and calculate the relative gain of ROLD AGE-ACT over the less targeted policies ROLD NO-TARGET and ROLD ACT in terms of the total loss objective. We then fit decision trees to predict these relative gains. We include as features several problem parameters (e.g., the basic reproduction number R_0, the cost of death, the disease severity, the elasticity of social contacts to activity levels, the ICU capacity) and derived features based on statistics of the econ-to-contacts-ratio across age groups and activities. Specifically, because each sample in the training set corresponds to an entire problem instance (including all time periods and all age groups and activities), we first calculate

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

econ-to-contacts-ratio values at the start of the pandemic, denoted as time t_0:

$$\text{econ-to-contacts-ratio}(g, a, t_0) := \frac{\left(\partial \text{Economic Value}(t) / \partial \ell_g^a(t)\right)_{t=t_0}}{N_g(t_0) \sum_{h \in \mathcal{G}} C_{h,g}^a},$$

and use these to calculate two kinds of economic features: (i) for each activity other than home, we take the standard deviation of econ-to-contacts-ratio across all age groups that are active in that activity, which we label econ-to-contacts-ratio-activity-stdev; (ii) for each age group, we take the standard deviation of econ-to-contacts-ratio across all activities (other than home) in which that age group is active, which we label econ-to-contacts-ratio-agegroup-stdev. (Relevant activities for each age group are listed in Table EC.8.) The full set of features is summarized in Table EC.12 and the details of the fitting procedure can be found in Appendix EC.6.4.2.

To appreciate how the various features influence the gains, consider Figure 11, which displays permutation importance scores corresponding to depth-10 trees for explaining the gains of the ROLD AGE-ACT policy over ROLD NO-TARGET and ROLD ACT; and Figures EC.8 and EC.9 in the Appendix, which depict trees of depth five explaining the same gains.

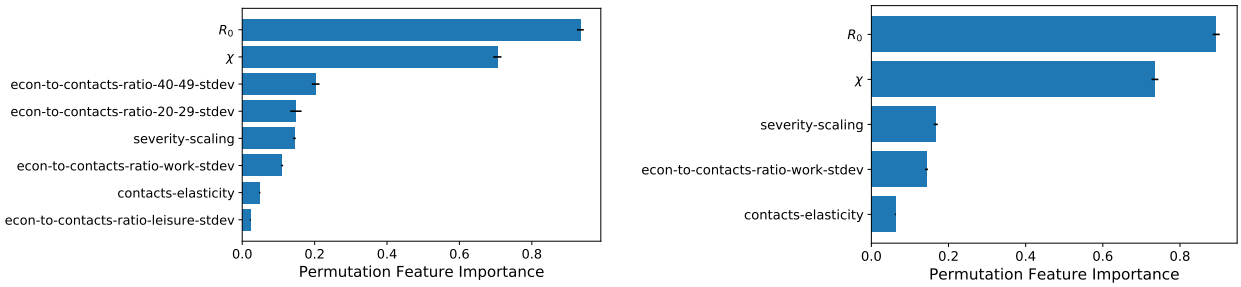

Figure 11 Permutation feature importance for the depth-10 decision trees explaining the relative gains of ROLD AGE-ACT over ROLD NO-TARGET (left) and ROLD ACT (right).

The most salient features for predicting the gains of dual targeting are R_0, the cost of death χ, the heterogeneity in the econ-to-contacts-ratio values, the disease severity, and the elasticity of social contacts to activity levels. The two most important features, R_0 and χ, impact gains in a similar, *non-monotonic* fashion, with moderate values leading to larger gains than very low or

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

very high values. This can be best viewed in the partial dependence plots shown in Figure 12. This is explainable because all policies tend to enforce similar activity levels at very low or very high values of R_0 or χ (that is, high activity values at sufficiently low values of R_0 or χ, and vice-versa), which limits the benefits that could be generated through more refined targeting.

Figure 12 also leads to practical insights. Policy makers faced with pandemic variants with very high basic reproduction number R_0 and who heavily prioritize health over economic outcomes should not concern themselves with targeting. The gains are also generally smaller in regimes with low R_0 — although then, dual targeting could lead to more substantial gains for policy makers who weigh more heavily (albeit not too heavily) health over economic outcomes. Lastly, in regimes with moderate values of R_0, all policy makers would gain substantially from targeting, and especially those who weigh economic outcomes more heavily.

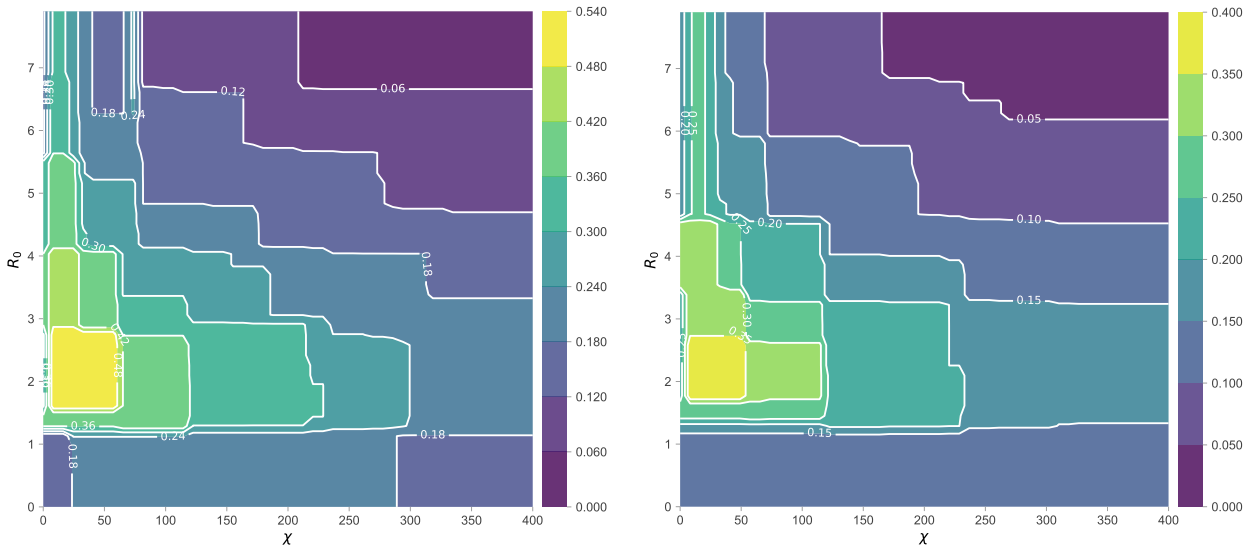

Figure 12 **Partial Dependence Plots for the depth-10 decision trees showing the relative gains of ROLD AGE-ACT over ROLD NO-TARGET (left) and ROLD ACT (right) as a function of R_0 and χ. The color scale depicts the magnitude of the relative gains. (The figure is best viewed in color.)**

The gains from dual targeting increase as the heterogeneity in the econ-to-contacts-ratio values (captured by the standard deviation) increases. This aligns well with the intuition developed in

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Section 7.1 that dual targeting exploits such heterogeneity to prioritize which (age group, activity) pairs to keep open. Intuitively, the gains relative to no-targeting are explained by heterogeneity across *both* age groups and activities, with the 40-49 and 20-29 y.o. emerging as the most relevant groups and work, leisure as the most relevant activities; in contrast, the gains relative to activity-based targeting are best explained by heterogeneity *across age groups* alone (in the work activity).

Generally, the gains decrease as the disease severity increases, reflecting that a higher probability of severe symptoms limits the improvements available from finer targeting. In the regime of small R_0, where the gains from targeting are limited, lower severity predicts even smaller gains. This is consistent with all policies recommending high activity levels, which leaves little room for improvements from targeting.

8. Towards Practically Implementable Policies

The previous sections established dual targeting as an idealized benchmark. However, discriminating activity levels based on age can be challenging in practice, from both an operational and an ethical/legal standpoint. We thus study two policy variations inspired by real-world interventions that are simpler to implement than dual targeting and that retain some of its superior health and economic outcomes.

In the first proposal, we consider *curfews* for different activities, implemented *uniformly* over all age groups. These draw inspiration from real-life policies during COVID-19: as an example, France implemented population-wide curfews during the first half of 2021 that started at a time varying from 6 p.m. to 11 p.m. and lasted until 6 a.m., while maintaining school and work activities largely de-confined (Reuters 2021). This is effectively implementing restrictions similar to ROLD AGE-ACT: under a 6 p.m. curfew, a typical member of the 20-65 y.o. group would be engaged in work until the start of the curfew, so their leisure and other activities would be quite restricted; in contrast, since most individuals aged above 65 are not in active employment, they would not face such restrictions in these last two activities. More generally, since separate age groups spend different parts of the day in different activities, curfews are a potential way to *implicitly* target age groups, but are operationally simpler and less contentious than the ROLD AGE-ACT policy.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

In the second proposal, we consider *recommending*, as opposed to enforcing, targeted activity levels separately for each age group, while accounting for imperfect population compliance. Such recommendations have been deployed in the pandemic response of many countries, with governments encouraging specific at-risk groups such as the elderly to limit their activities outside of home (see the examples in Table EC.15). Such policies are thus more readily implementable than fully targeted confinements, and our goal with the exercise is to determine the compliance levels needed to ensure a sufficiently good performance relative to activity-based targeting.

8.1. Optimized Curfews

We focus on *practical* curfew policies: we do not allow targeting by age groups, we consider a single start and end-time within a 24-hour day (further constrained to occur at the hour or half-hour mark), we assume that activity levels are completely restricted during the curfew and unrestricted otherwise, and we allow curfews to change at most once every two weeks.

To formalize our model, we break the 24-hour day into 48 slots of 30 minutes each, starting at the hour and half-hour mark. For day t, we denote the τ-th time slot of the day by t_τ, for $\tau = 1, \ldots, 48$, with $\tau = 1$ denoting the 12 a.m.-12.30 a.m. time window. We encode the curfew for activity a as starting at the beginning of slot $s^a_{\text{start}}(t)$ and ending at the end of slot $s^a_{\text{end}}(t)$. For each activity $a \in \mathcal{A} \setminus \{\mathsf{home}\}$, we let $\ell^a(t_\tau)$ denote the activity level allowed during time slot τ of day t, which we take to be 0 during the curfew and 1 outside of the curfew; we set $\ell^{\mathsf{home}}(t) = 1$ throughout.

To model curfews, we refine our social mixing model to take into account the participation rates of different age groups in distinct activities throughout the day. To that end, we requested and obtained time-use survey data from the French National Archive of Data from Official Statistics (ADISP, INSEE 2010) that include detailed information of how survey participants of different ages spent every 10-minute interval of their day, from which we estimate participation rates in the different activities (details of our procedure are discussed in Appendix EC.8.1). Let $\mathcal{R}^a_g(\tau)$ denote the participation rate of age group g in activity a during time slot τ of a generic day, so that $\sum_{a \in \mathcal{A}} \mathcal{R}^a_g(\tau) = 1, \forall \tau, g$. As in our base model, let $c_{g,h}(\boldsymbol{\ell}(t))$ denote the mean number of total daily

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

contacts between an individual in group g and individuals in group h across all activities during day t, when the activity levels are $\boldsymbol{\ell}(t)$. We then set

$$c_{g,h}(\boldsymbol{\ell}(t)) = \sum_{a\in\mathcal{A}} C^a_{g,h} \frac{\sum_\tau \mathcal{R}^a_g(\tau)\mathcal{R}^a_h(\tau)\cdot(\ell^a(t_\tau))^{2\alpha}}{\sum_\tau \mathcal{R}^a_g(\tau)\mathcal{R}^a_h(\tau)}. \tag{14}$$

This expression extends our base-case social mixing model from Section 3 by capturing the distinct participation rates of the age groups through a weighted-average combination of contacts during the day. Note that if we removed the differentiation between time slots, i.e., $\ell^a_g(t_\tau) = \ell^a_g(t_{\tau'}) = \ell^a_g(t), \forall \tau, \tau', g, a$, then (14) would collapse to our familiar social mixing model in (1) (for the case when all age groups have the same activity levels).

To determine the optimized curfew policy, which we henceforth refer to as CURFEW, we solve an optimization problem analogous to the one in (8)-(10), with the following modifications. First, the decision variables are the "start" and "end" time of the curfew in each activity on each day t, $\{s^a_{\text{start}}(t), s^a_{\text{end}}(t)\}_{a\in\mathcal{A}}$, allowed to change once every 14 days. Second, the calculation of Economic Loss in the objective replaces $\boldsymbol{\ell}(t)$ in (7) with $\boldsymbol{\ell}^{\text{eff}}(t)$, where $\boldsymbol{\ell}^{\text{eff}}(t) := [\ell^{a,\text{eff}}_g(t)]_{a\in\mathcal{A},g\in\mathcal{G}}$ and

$$\ell^{a,\text{eff}}_g(t) := \frac{\sum_\tau \mathcal{R}^a_g(\tau)\ell^a(t_\tau)}{\sum_\tau \mathcal{R}^a_g(\tau)}, \quad \forall g, a \tag{15}$$

denotes the (weighted-average) effective activity level of group g in activity a during day t, accounting for the group's participation rates. Third, the SEIR dynamics in constraint (9) are modified to reflect the revised social mixing in (14). Lastly, we solve the optimization problem through a gradient descent procedure, and round its optimal solution to the closest hour or half-hour mark (details are provided in Appendix EC.8.2).

Figure 13 shows that optimized curfews can significantly outperform ROLD ACT in terms of the total loss objective: CURFEW closes 24.3% (median value) of the performance gap between ROLD AGE-ACT and ROLD ACT, and the improvement can reach up to 65.1% of the gap for moderate values of the cost of death χ, namely $\chi = 145$. As a side remark, we conjecture that these improvements are conservative estimates of what is achievable because the optimization procedure used for CURFEW is significantly less sophisticated than the one used for ROLD AGE-ACT.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

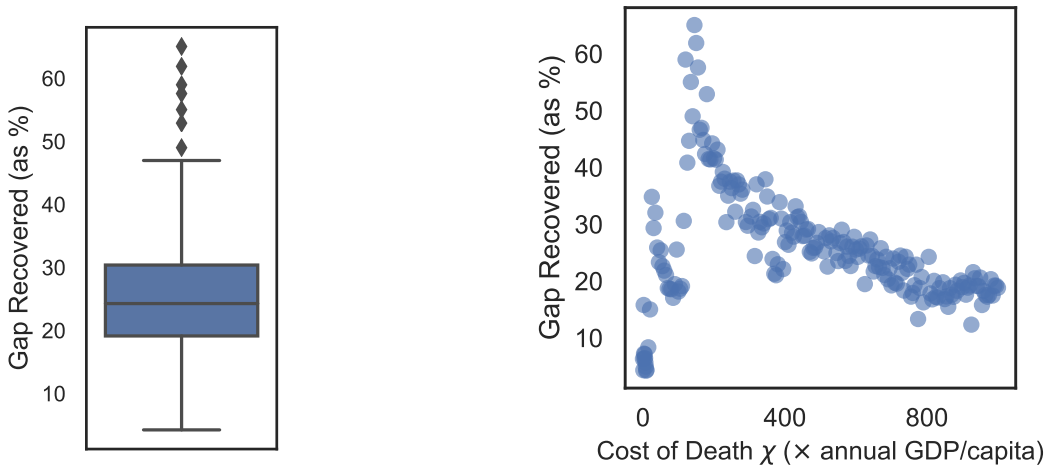

Figure 13 **The performance of the CURFEW policy expressed as a fraction of the gap in total loss between ROLD ACT and AGE-ACT that it recovers. We include 199 distinct values of χ, from 0 to 1000×.**

That CURFEW can recover a substantial portion of AGE-ACT's performance despite not explicitly targeting age groups suggests that time-of-day targeting implicitly allows some differentiation by age. To examine this, we calculate the time-average activity levels of all the relevant age groups engaged in a given activity (per Table EC.8), based on which we determine the range of activity levels encountered across groups. Figure 14 plots the magnitude of these ranges for the CURFEW policy. When comparing these with ROLD ACT, which by definition has zero range in each activity, it can be seen that CURFEW effectively differentiates restrictions by age in each activity.

Overall, these results highlight the benefits of the CURFEW policy, which can significantly improve performance relative to ROLD ACT without incurring the operational or ethical implementation challenges of ROLD AGE-ACT.

8.2. Dual-Targeted Recommendations with Limited Compliance

We next examine a practically-appealing implementation of dual targeting that phrases ROLD's output as *recommendations* instead of mandatory restrictions and that also accounts for imperfect compliance, whereby only a subset of the population would follow the recommendations. More formally, we capture compliance through a parameter p, with $0 \leq p \leq 1$, such that a p fraction of the population of each age group adheres to the recommended age group-targeted activity levels, and the remaining $1-p$ fraction does not adhere, but must still abide by a population-wide restriction

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

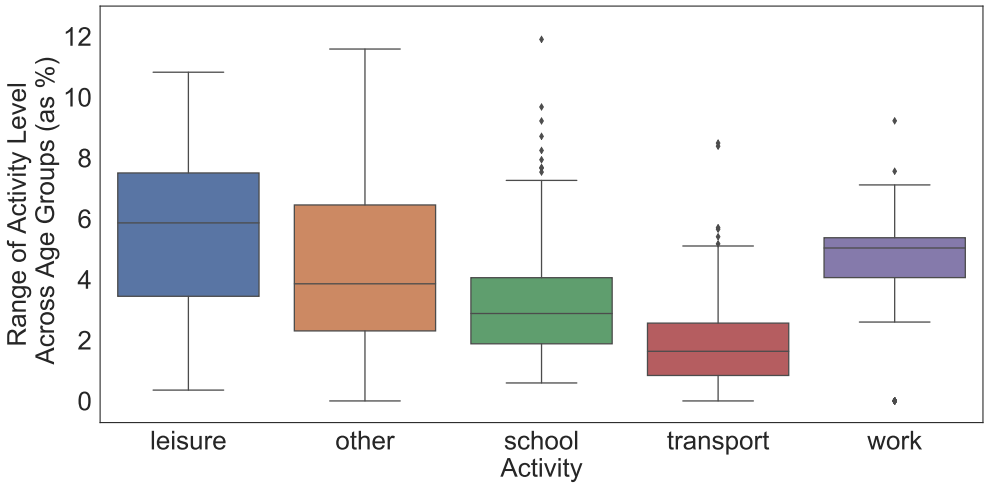

Figure 14 **The range of time-average activity levels (defined as maximum minus minimum level over all relevant age groups) for the CURFEW policy, in several problem instances parametrized by χ. In comparison, the ROLD ACT policy has equal activity levels for all age groups, and thus zero range for each activity.**

in each activity (which is assumed enforceable, as in ROLD ACT). We refer to the resulting policy as the ROLD COMP(p) policy.

More precisely, ROLD COMP(p) still optimizes for $\boldsymbol{\ell}$ via the optimization problem (8)-(10), but builds in limited compliance into this optimization by adding the following constraints:

$$\ell_g^a(t) = p \cdot x_g^a(t) + (1-p) \cdot y^a(t), \quad \forall g, a, t \tag{16}$$

$$y^a(t) \geq x_g^a(t), \quad \forall g, a, t \tag{17}$$

$$0 \leq x_g^a(t) \leq 1, \quad \forall g, a, t \tag{18}$$

$$0 \leq y^a(t) \leq 1, \quad \forall a, t. \tag{19}$$

This confinement policy can be interpreted as a linear combination of two policies: an age-and-activity targeted policy, $x_g^a(t)$, which the policy maker can only recommend; and an activity-only targeted policy, $y^a(t)$, which the policy maker enforces. The linear weights reflect the percentage of the population that would be compliant with the recommendation $x_g^a(t)$, versus those who would only abide by the enforceable activity levels $y^a(t)$. Constraint (17) then expresses that the enforceable activity level cannot be more restrictive than the recommended one.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Figure 15, which shows ROLD COMP(p)'s performance for different compliance levels p, confirms the practical promise of the approach. The left panel, which depicts the fraction of the performance gap (in terms of the total cost objective) between ROLD ACT and AGE-ACT recovered by ROLD COMP(p), suggests that significant improvements are already available at relatively modest compliance levels. Specifically, the median gap recovered is 10.4% at 40% compliance and 19.2% at 60% compliance. As p increases, the performance naturally improves, and the relationship between performance and compliance level is overall convex, with larger marginal gains occurring at higher compliance levels.[8] The right panel suggests that ROLD COMP(p) Pareto-dominates ROLD ACT even for small compliance levels when restricting to moderate values of χ, and Pareto-dominates over all tested χ's for sufficiently large levels of compliance. We also consider a variation of limited compliance in which we allow each age group to have a group-specific compliance fraction calibrated from empirical data (Ganslmeier et al. 2022) and the resuls are very similar (Appendix EC.8.4).

These results suggest that *recommendations* for reducing activity levels according to the dual-targeted ROLD AGE-ACT policy, accompanied by reasonable levels of compliance in the population, can lead to significant health and economic benefits compared to less targeted policies.

9. Discussion

Our study established that dual targeting could in principle lead to significant benefits and that certain simple, practical policies can achieve some of these benefits. Dual targeting can improve both health *and* economic outcomes and it can maintain higher activity levels for most age groups (and particularly for those most confined), thus making confinements more socially acceptable. Such benefits are generated by heavily leveraging complementarities through a simple "bang-for-the-buck" rule: impose less confinement on group-activity pairs with a high econ-to-contacts-ratio,

[8] Note that at $p = 1$, the fraction of the ROLD ACT/AGE-ACT performance gap recovered can be less than 100%. This is due to our optimization approach, which puts ROLD COMP at a relative disadvantage: ROLD COMP is initialized with the ROLD ACT solution, but ROLD AGE-ACT is initialized with the best of two (separate) gradient descent runs initialized at the ROLD ACT and the ROLD AGE solutions. One can easily improve ROLD COMP by improving its initialization, so the current results should be viewed as a conservative estimate of its performance.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

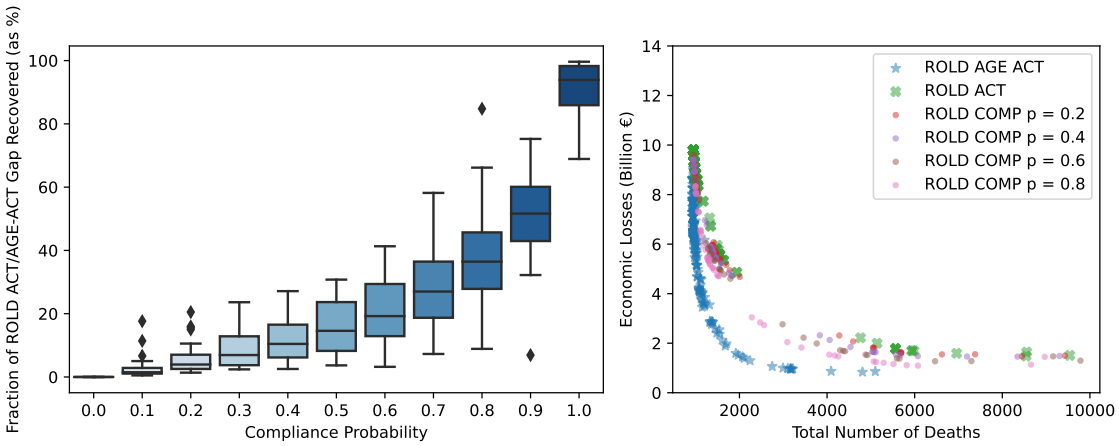

Figure 15 **The performance of limited compliance policies for different compliance levels. In the left panel, each boxplot summarizes the fraction of the ROLD ACT/AGE-ACT performance gap (in terms of the total cost objective) that ROLD COMP(p) recovers, in 201 problem instances parametrized by the cost of death χ (with values from 0 to 1000$\times$). The right panel illustrates the total number of deaths and the economic losses generated by ROLD COMP(p) for different compliance levels p, compared to ROLD AGE-ACT and ROLD ACT; each marker corresponds to a different value of χ.**

i.e., high marginal economic value prorated by activity-specific social contacts. The gains afforded by targeting are generally largest for moderate values of the basic reproduction number R_0 and the cost of death, high heterogeneity in the econ-to-contacts-ratio, and low disease severity. Because dual targeting can face significant implementation challenges, we introduced two practical proposals — based on curfews and recommendations — that achieve some of the benefits without explicitly discriminating based on age.

Future work. The algorithmic developments in our proposed framework for quantifying the possible benefits of dual targeting could be improved in several ways. One is end-of-horizon effects: ROLD may allow for increased infections as it nears the end of the horizon, since these may not progress to deaths. A policy maker may want to adjust ROLD's behavior to prevent this, for example through a thorough sensitivity analysis of the optimal policy with respect to the problem horizon, or by imposing penalties on the remaining infections, or by running ROLD in a rolling (non-shrinking) horizon fashion. Secondly, as seen with emerging COVID-19 variants, the

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

underlying epidemic environment could change during a typical ROLD time horizon. To track and respond to such changes, a potential approach could be to re-estimate the model parameters on the fly as ROLD advances, or to adopt a robust model that explicitly allows for mis-specification. One could also consider second-order approximations of the dynamics and objective and leverage (non-convex) quadratic programming to deal with the resulting optimization problems. Lastly, work could also be devoted to deriving a theoretical justification for the importance of the econ-to-contacts-ratio by generalizing our simple analysis in Appendix EC.7, or by deriving tractable upper bounds for the performance of controlled SEIR models that could be used to benchmark interventions.

Our modeling choices were constrained by available data, but further data collection may allow for more precise modeling. Our social mixing model relied on social contact matrices by age group and activity, which are available from surveys conducted in a number of countries; additional data may allow refining the population group or activity definitions. Additionally, the available social contacts data only record interactions in the same activity, whereas contacts occur as individuals are engaged in different activities (e.g., a services industry professional interacts during work with individuals who are engaged in leisure activities). A more refined mixing model that captures such interactions would be desirable, provided that data are available to calibrate it. Similarly, we worked with economic data reported by industry activities and for coarse population groups, but a dataset that splits economic value into fine-grained (group, activity) contributions would be more suitable. Lastly, a more detailed economic model would include cross effects, e.g., children in school producing value in conjunction with educational staff engaged in work.

Although we focus on confinements, a direction for future research is to see how to combine these with other types of targeted interventions. The framework is sufficiently flexible to accommodate interventions such as contact tracing, mass testing and also vaccinations, although a careful implementation is beyond the scope of this article.

Finally, to address fairness concerns related to targeting in a principled way, one could impose fairness requirements based on the interventions' actions and/or outcomes, e.g., requiring that the

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

health or economic losses or the time in confinement faced by different groups should satisfy certain axiomatic fairness properties (Young 1994).

Acknowledgments

References

Acemoglu D, Chernozhukov V, Werning I, Whinston MD (2020) Optimal targeted lockdowns in a multi-group SIR model. Working Paper 27102, NBER, URL `http://dx.doi.org/10.3386/w27102`.

Aguilera J (2020) Some supermarkets are launching senior-only hours during the coronavirus pandemic. Not all retailers think that's a good idea. URL `https://time.com/5804574/grocery-store-chains-seniors-coronavirus/`, accessed January 12, 2021.

Ahn HS, Silberholz J, Song X, Wu X (2021) Optimal covid-19 containment strategies: Evidence across multiple mathematical models. *Available at SSRN 3834668* .

Alvarez FE, Argente D, Lippi F (2020) A simple planning problem for COVID-19 lockdown. Working Paper 26981, National Bureau of Economic Research, URL `http://dx.doi.org/10.3386/w26981`.

Anderson RM, May RM (1992) *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press).

Bastani H, Drakopoulos K, Gupta V, Vlachogiannis I, Hadjicristodoulou C, Lagiou P, Magiorkinis G, Paraskevis D, Tsiodras S (2021) Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature* 599(7883):108–113, URL `http://dx.doi.org/10.1038/s41586-021-04014-z`.

Bemporad A (2006) Model predictive control design: New trends and tools. *Proceedings of the 45th IEEE Conference on Decision and Control*, 6678–6683 (IEEE).

Béraud G, et al. (2015) The French connection: The first large population-based contact survey in France relevant for the spread of infectious diseases. *PLOS ONE* 10(7):1–22.

Bertsimas D, Ivanhoe JK, Jacquillat A, Li ML, Previero A, Lami OS, Bouardi HT (2020) Optimizing vaccine allocation to combat the covid-19 pandemic. *medRxiv* .

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

Boloori A, Saghafian S (2020) Health and economic impacts of societal intervention policies in the early stage of covid-19 in the u.s. *Available at SSRN* .

Bose S, Souyris S, Mukherjee UK, Ivanov A, Seshadri S, III ACE, Xu Y (2021) Control of epidemic spreads via testing and lock-down. Technical report.

Camacho EF, Alba CB (2013) *Model predictive control* (Springer science & business media).

Chang S, Pierson E, Koh PW, Gerardin J, Redbird B, Grusky D, Leskovec J (2020) Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* 589:1–6.

Cot C, Cacciapaglia G, Sannino F (2021) Mining Google and Apple mobility data: Temporal anatomy for COVID-19 social distancing. *Scientific reports* 11(1):1–8.

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

Du Z, Xu X, Wu Y, Wang L, Cowling BJ, Meyers L (2020) Serial interval of COVID-19 among publicly reported confirmed cases. *Emerging Infectious Diseases* 26(6).

Duque D, Morton DP, Singh B, Du Z, Pasco R, Meyers LA (2020) Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proceedings of the National Academy of Sciences* 117(33):19873–19878.

Dutta R, Gomes SN, Kalise D, Pacchiardi L (2021) Using mobility data in the design of optimal lockdown strategies for the COVID-19 pandemic. *PLoS Computational Biology* 17(8):e1009236.

El Housni O, Sumida M, Rusmevichientong P, Topaloglu H, Ziya S (2020) Can testing ease social distancing measures? future evolution of covid-19 in nyc. *arXiv preprint arXiv:2005.14700* .

Evgeniou T, Fekom M, Ovchinnikov A, Porcher R, Pouchol C, Vayatis N (2020) Epidemic models for personalised covid-19 isolation and exit policies using clinical risk predictions. *medRxiv* .

Favero CA, Ichino A, Rustichini A (2020) Restarting the economy while saving lives under COVID-19. C.E.P.R Discussion Paper 14664, CEPR.

Foy H (2020) Russian businesses prepare for fresh lockdowns as Covid-19 cases soar. *The Financial Times* URL `https://www.ft.com/content/f0696f19-fe0b-425c-8e73-3703bd44dec3`, accessed January 12, 2021.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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

Ganslmeier M, Van Parys J, Vlandas T (2022) Compliance with the first uk covid-19 lockdown and the compounding effects of weather. *Scientific Reports* 12(1):3821, URL http://dx.doi.org/10.1038/s41598-022-07857-2.

Giordano G, Colaneri M, Di Filippo A, Blanchini F, Bolzern P, De Nicolao G, Sacchi P, Colaneri P, Bruno R (2021) Modeling vaccination rollouts, SARS-CoV-2 variants and the requirement for non-pharmaceutical interventions in Italy. *Nature Medicine* 27(6):993–998.

Godeluck S (2020) Coronavirus : plus de 2.000 patients en réanimation en ˆIle-de-France. Les Echos, URL https://www.lesechos.fr/economie-france/social/coronavirus-plus-de-2000-patients-en-reanimation-en-ile-de-france-1191003, accessed June 30, 2020.

Goldstein JR, Cassidy T, Wachter KW (2021) Vaccinating the oldest against COVID-19 saves both the most lives and most years of life. *Proceedings of the National Academy of Sciences* 118(11), ISSN 0027-8424.

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

Harrison S (2020) Coronavirus: Ireland's restrictions eased for over 70s. *BBC* URL https://www.bbc.com/news/uk-northern-ireland-52547169, accessed January 12, 2021.

Ilin C, Annan-Phan S, Tai XH, Mehra S, Hsiang S, Blumenstock JE (2021) Public mobility data enables COVID-19 forecasting and management at local and global scales. *Scientific reports* 11(1):1–11.

INSEE (2010) Emploi du temps (version pour Eurostat). 2009-2010, INSEE [producer], ADISP [distributor].

Kanbur N, Ankgül S (2020) Quaranteenagers: A single country pandemic curfew targeting adolescents in Turkey. *Journal of Adolescent Health* 67(2).

Kucharski AJ, et al. (2020) 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 Infectious Diseases* 20(10):1151–1160.

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

Li ML, Bouardi HT, Lami OS, Trikalinos TA, Trichakis NK, Bertsimas D (2021) Forecasting covid-19 and analyzing the effect of government interventions. *MedRxiv* 2020–06.

Magid J (2020) Cabinet approves removal of age restriction for those returning to work. *The Times of Israel* URL https://www.timesofisrael.com/liveblog-may-2-2020/, accessed January 12, 2021.

Mas-Colell A, Whinston MD, Green JR, et al. (1995) *Microeconomic theory*, volume 1 (Oxford University Press New York).

Matrajt L, Eaton J, Leung T, Brown ER (2020) Vaccine optimization for COVID-19, who to vaccinate first?

Morris DH, Rossine FW, Plotkin JB, Levin SA (2021) Optimal, near-optimal, and robust epidemic control. *Communications Physics* 4(1):1–8.

Pataro IML, et al. (2021) A control framework to optimize public health policies in the course of the COVID-19 pandemic. *Scientific Reports* 11(1):13403.

Perakis G, Singhvi D, Skali Lami O, Thayaparan L (2021) COVID-19: A multipeak SIR based model for learning waves and optimizing testing. *Available at SSRN 3817680* .

Prem K, Cook AR, Jit M (2017) Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLoS computational biology* 13(9):e1005697.

Prem K, et al. (2020) 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 Health* 5(5):e261–e270.

Reuters (2021) France goes under nationwide 6pm curfew as Covid-19 death toll surpasses 70,000. *France24* URL https://www.france24.com/en/europe/20210116-france-set-for-nationwide-6pm-curfew-in-effort-to-stem-covid-19, accessed January 12, 2021.

Reuters Staff (2020) Bosnian region eases lockdown on seniors, children after court ruling. *Reuters* URL https://www.reuters.com/article/us-health-coronavirus-bosnia-idUSKCN2261RA, accessed January 12, 2021.

Salje H, et al. (2020) Estimating the burden of SARS-CoV-2 in France. *Science* 38.

Sterlé C, Alexandre V, de Livonnière S (2020) Coronavirus en île-de-France : 13 000 malades encore hospitalisés, déjà 40 000 tests positifs. Le Parisien, URL https://www.leparisien.fr/paris-75/coronavirus-en-ile-de-france-13-000-malades-encore-hospitalises-deja-40-000-tests-positifs-29-04-2020-8307801.php, accessed June 30, 2020.

Tiirinki H, et al. (2020) COVID-19 pandemic in Finland–preliminary analysis on health system response and economic consequences. *Health policy and technology* 9(4):649–662.

Wellenius GA, et al. (2020) Impacts of state-level policies on social distancing in the United States using aggregated mobility data during the COVID-19 pandemic. *arXiv e-prints* arXiv–2004.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Wille L, Van Hoang T, Funk S, Coletti P, Beutels P, Hens N (2020) SOCRATES: an online tool leveraging a social contact data sharing initiative to assess mitigation strategies for COVID-19. *BMC Research Notes* 13(293).

Xiong C, Hu S, Yang M, Luo W, Zhang L (2020) Mobile device data reveal the dynamics in a positive relationship between human mobility and COVID-19 infections. *Proceedings of the National Academy of Sciences* 117(44):27087–27089.

Young HP (1994) *Equity* (Princeton University Press).

Yuan Yx (2015) Recent advances in trust region algorithms. *Mathematical Programming* 171:249–281.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

E-companion to Quantifying and Realizing the Benefits of Targeting for Pandemic Response

EC.1. Dynamics of the Controlled SEIR Epidemic Model

We write down a set of discrete time dynamics for the controlled SEIR model. We use notation $\Delta Z(t)$ to denote $Z(t+1)-Z(t)$. For all groups $g \in \mathcal{G}$, we have:

$$\Delta N_g(t) = -\underbrace{\mu \cdot (p_g^H + p_g^{ICU}) \cdot I_g(t)}_{\text{infected going to ICU, H}} \tag{EC.1}$$

$$\Delta S_g(t) = -\beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)} \right) \tag{EC.2}$$

$$\Delta E_g(t) = \beta(t) S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)} \right) - \sigma E_g(t) \tag{EC.3}$$

$$\Delta I_g(t) = \sigma E_g(t) - \underbrace{\mu \cdot (p_g^H + p_g^{ICU}) \cdot I_g(t)}_{\text{go to ICU, H}} - \underbrace{\mu \cdot (1 - p_g^H - p_g^{ICU}) \cdot I_g(t)}_{\text{recover}} \tag{EC.4}$$

$$\Delta R_g(t) = \mu \cdot (1 - p_g^H - p_g^{ICU}) \cdot I_g(t) \tag{EC.5}$$

$$\Delta R_g^{\mathrm{q}}(t) = \underbrace{\lambda_g^H \cdot p_g^R \cdot H_g(t) + \lambda_g^{ICU} \cdot p_g^R \cdot ICU_g(t)}_{\text{recovery of known infected cases}} \tag{EC.6}$$

$$\Delta H_g(t) = -\lambda_g^H H_g(t) + \mu \cdot p_g^H \cdot I_g(t) - B_g^H(t) \tag{EC.7}$$

$$\Delta ICU_g(t) = -\lambda_g^{ICU} ICU_g(t) + \mu \cdot p_g^{ICU} \cdot I_g(t) - B_g^{ICU}(t) \tag{EC.8}$$

$$\Delta D_g(t) = \lambda_g^H \cdot p_g^D \cdot H_g(t) + \lambda_g^{ICU} \cdot p_g^D \cdot ICU_g(t) + B_g^H(t) + B_g^{ICU}(t). \tag{EC.9}$$

In (EC.5) and (EC.6), 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 (EC.9), 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 now provide justification for how we account for social contacts and, in particular, for the expressions in (EC.2) and (EC.3). We note that individuals in R_g^{q} can interact with members of S_g, E_g, I_g and R_g. Fix a person i in age group $g \in \mathcal{G}$, in state S_g. Then:

$$Pr\left(\bigcup_{h \in \mathcal{G}} \{i \text{ got infected through socializing with age group } h\} \right) \tag{EC.10}$$

$$= Pr\left(\bigcup_{h \in \mathcal{G}} \{i \text{ got infected through socializing with individuals in } I_h\} \right) \tag{EC.11}$$

$$= 1 - Pr\left(\{i \text{ did not get infected through socializing with individuals in any } I_h \text{ for any } h \in \mathcal{G}\}\right) \tag{EC.12}$$

$$= 1 - \prod_{h \in \mathcal{G}} Pr\left(\{i \text{ did not get infected through socializing with individuals in } I_h\}\right) \tag{EC.13}$$

$$= 1 - \prod_{h \in \mathcal{G}} \left(1 - \beta(t) \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)} \right)^{c_{g,h}} \tag{EC.14}$$

$$\approx 1 - \prod_{h \in \mathcal{G}} \left(1 - \beta(t) c_{g,h} \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)} \right) \tag{EC.15}$$

$$\approx \beta(t) \sum_{h \in \mathcal{G}} c_{g,h} \frac{I_h(t)}{N_h(t) + R_h^{\mathrm{q}}(t)}. \tag{EC.16}$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

In (EC.14) we use the following reasoning. Having fixed person i in age group g, (a) a contact with a randomly chosen individual in group h will result in person i getting infected with probability $\beta(t)\frac{I_h(t)}{N_h(t)+R_h^{\mathrm{q}}(t)}$, and (b) the number of person i's contacts with individuals in age group h is given by $c_{g,h} = c_{g,h}(\ell_g(t), \ell_h(t))$. Finally, person i getting infected as the result of a contact with someone from group h is considered to be an independent event across different contacts, therefore we raise the probability of no infection from a contact to the power of the number of contacts. (EC.15) and (EC.16) follow from linear approximations.

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 (EC.2) and (EC.3).

EC.2. Details of the Economic Model

As discussed in Section 3.4, economic losses come from three separate sources:

Effect of confinement. To account for confinement in the non-quarantined population, we make the economic value generated per day 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

$$v_g(\boldsymbol{\ell}) := \begin{cases} v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{0-9 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) + v_g^{\text{schooling}}(\boldsymbol{\ell}) & \text{if } g = \text{10-19 y.o.} \\ v_g^{\text{employment}}(\boldsymbol{\ell}) & \text{if } g = \text{20-29, 30-39, 40-49, 50-59, 60-69 y.o.} \\ 0 & \text{otherwise.} \end{cases} \tag{EC.17}$$

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 generated by 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 weighted average of $\ell^{\text{transport}}, \ell^{\text{leisure}}$ and ℓ^{other} across these three activities and all groups $g \in \mathcal{G}$:

$$v_g^{\text{employment}}(\boldsymbol{\ell}) := w_g \cdot \left(\nu^{\text{work}} \cdot \ell_g^{\text{work}} + \nu^{\text{other activities}} \cdot \left(\frac{1}{|\mathcal{G}|} \sum_{h\in\mathcal{G}} w^{\text{transport}} \ell_h^{\text{transport}} + w^{\text{leisure}} \ell_h^{\text{leisure}} + w^{\text{other}} \ell_h^{\text{other}} \right) + \nu^{\text{fixed}} \right). \tag{EC.18}$$

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 (EC.18). Then w_g measures the overall daily employment value of a typical member of group g under no confinement, and is equal to $v_g^{\text{employment}}(\mathbf{1})$. The weights $w^{\text{transport}}, w^{\text{leisure}}, w^{\text{other}}$ capture the relative importance of each of these

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

three activities for employment value. For our baseline setting, we take $w^{\text{transport}} = w^{\text{leisure}} = w^{\text{other}} = 1/3$. We let these parameters vary in our analyses in Section 7.2.

We estimate the coefficients of this model from data, as we describe in detail in Section EC.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 discount over 15 years, and for 10-19 year olds we discount over five years. The discounting factor we apply is thus

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

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.[9] Lastly, we also use a multiplicative factor θ 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 5.3.

Thus, the definition for value of school days is

$$v_g^{\text{schooling}}(\boldsymbol{\ell}) := \theta \cdot f_g \cdot \delta_g \cdot v_{\text{20-29 y.o}}^{\text{employment}}(\mathbf{1}) \cdot \ell_g^{\text{school}}. \tag{EC.19}$$

Effect of quarantine, illness and deaths during the pandemic. We capture the economic effect of quarantine, hospitalization, and death during the pandemic by assuming that if at some time period an individual in group g is in one of the SEIR chambers H_g, ICU_g, D_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 lost future wages due to death. We account for a deceased individual's lost wages which they would have earned from their current age until retirement age, given the prevailing wage curve under normal circumstances, $\{v_g(\mathbf{1})\}_{g\in\mathcal{G}}$. For group g, we set the current age to the midpoint 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[10]

$$\begin{aligned} v_{30-39\text{ y.o.}}^{\text{life}} := \sum_{\tau=35}^{69} \frac{1}{(1+r)^{\tau-35}} \\ \cdot \Big(\mathbb{1}\{35 \le \tau \le 39\} \cdot v_{30-39\text{ y.o.}}^{\text{employment}}(\mathbf{1}) + \mathbb{1}\{40 \le \tau \le 49\} \cdot v_{40-49\text{ y.o.}}^{\text{employment}}(\mathbf{1}) \\ + \mathbb{1}\{50 \le \tau \le 59\} \cdot v_{50-59\text{ y.o.}}^{\text{employment}}(\mathbf{1}) + \mathbb{1}\{60 \le \tau \le 69\} \cdot v_{60-69\text{ y.o.}}^{\text{employment}}(\mathbf{1}) \Big). \end{aligned} \tag{EC.20}$$

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

$$V := \sum_{t=0}^{T-1} \sum_{g\in\mathcal{G}} v_g(\mathbf{1}) \cdot \big(N_g(0) + H_g(0) + ICU_g(0) + R_g^{\text{q}}(0) \big). \tag{EC.21}$$

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

[9] 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.

[10] 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 EC.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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

EC.3. Algorithmic Details for ROLD

In this section, we clarify the algorithmic details of the linearization-optimization procedure described in Section 4. We first focus on how we build a linear model given $k, \boldsymbol{X}_k$ and $\hat{\boldsymbol{u}}_{k:T-1}$.

EC.3.1. Linearized Dynamics

In Step 2, the algorithm builds an approximation of the state dynamics that is linear in the controls $\boldsymbol{u}_k, \ldots, \boldsymbol{u}_{T-1}$. Here, we compute the coefficients for each $\boldsymbol{u}_t$ explicitly. We introduce the notation:

$$\boldsymbol{A}_t := \mathbb{I} + \Delta t \cdot \nabla_X \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \tag{EC.22}$$

$$\boldsymbol{B}_t := \Delta t \cdot \nabla_u \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \tag{EC.23}$$

$$\boldsymbol{c}_t := \Delta t \cdot \left[\boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) - \nabla_X \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \cdot \hat{\boldsymbol{X}}_t - \nabla_u \boldsymbol{F}_t(\hat{\boldsymbol{X}}_t, \hat{\boldsymbol{u}}_t) \cdot \hat{\boldsymbol{u}}_t\right], \tag{EC.24}$$

where matrix $\boldsymbol{A}_t$ has dimensions $|\mathcal{G}||\mathcal{X}| \times |\mathcal{G}||\mathcal{X}|$, matrix $\boldsymbol{B}_t$ has dimensions $|\mathcal{G}||\mathcal{X}| \times |\mathcal{G}||\mathcal{U}|$, and vector $\boldsymbol{c}_t$ has dimensions $|\mathcal{G}||\mathcal{X}| \times 1$. With this, we have

$$\boldsymbol{X}_{t+1} = \boldsymbol{A}_t \boldsymbol{X}_t + \boldsymbol{B}_t \boldsymbol{u}_t + \boldsymbol{c}_t, \quad t = 0, \ldots T-1.$$

We can then express the state $\boldsymbol{X}_t$ as[11]

$$\boldsymbol{X}_t = \left(\prod_{\tau=t-1,t-2,\ldots,k} \boldsymbol{A}_\tau\right) \boldsymbol{X}_k + \sum_{\tau=k}^{t-1} \left(\prod_{i=t-1,t-2,\ldots,\tau+1} \boldsymbol{A}_i\right) \boldsymbol{B}_\tau \boldsymbol{u}_\tau + \sum_{\tau=k}^{t-1} \left(\prod_{i=t-1,t-2,\ldots,\tau+1} \boldsymbol{A}_i\right) \boldsymbol{c}_\tau. \tag{EC.25}$$

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

EC.3.2. Constraint Coefficients

We can write each of the constraints (4) and (5) in the form

$$L_t \leq K(t), \text{ where } L_t := \boldsymbol{\gamma_x}^\top \cdot \boldsymbol{X}_t + \boldsymbol{\gamma_u}^\top \cdot \boldsymbol{u}_t, \forall t \in \{k, k+1, \ldots, T-1\}, \tag{EC.26}$$

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

$$\begin{cases} \boldsymbol{\gamma_u}^\top, & \text{for } \tau = t \\ \boldsymbol{\gamma_x}^\top \left(\prod_{i=t-1,t-2,\ldots,\tau+1} \boldsymbol{A}_i\right) \boldsymbol{B}_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases} \tag{EC.27}$$

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

EC.3.3. Objective Coefficients

Up to constants, the objective in (13) can be written as

$$\sum_{t=k}^{T} \left(\boldsymbol{d}_t^\top \boldsymbol{X}_t + \boldsymbol{e}_t^\top \boldsymbol{u}_t\right) \tag{EC.29}$$

[11] In the expressions in (EC.25), (EC.27), (EC.31), we follow the convention that a product of matrices over an empty set of indices results in the identity matrix.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Algorithm 1 CALCULATECONSTRAINTCOEFFICIENTS

Initialization: $\bar{\boldsymbol{X}}_k := \boldsymbol{X}_k$
for $t = k, k+1, \ldots, T-1$ **do**
  Calculate the constants in L_t as: ${\gamma_x}^\top \bar{\boldsymbol{X}}_t$
  Set $\bar{\boldsymbol{X}}_{t+1} := \boldsymbol{A}_t \bar{\boldsymbol{X}}_t + \boldsymbol{c}_t$
  Calculate coefficients for $\boldsymbol{u}_\tau$ in L_t as:

$$\begin{cases} {\gamma_u}^\top, & \text{for } \tau = t \\ {\gamma_x}^\top \bar{\boldsymbol{A}}_{t,\tau} \boldsymbol{B}_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases}$$

where

$$\bar{\boldsymbol{A}}_{t,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = t-1 \\ \bar{\boldsymbol{A}}_{t,\tau+1} \boldsymbol{A}_{\tau+1}, & \text{for } \tau = t-2, t-3, \ldots, k. \end{cases} \tag{EC.28}$$

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

with

$$\boldsymbol{d}_t{}^\top = \begin{cases} \hat{\boldsymbol{u}}_t^\top \boldsymbol{M} + \boldsymbol{\gamma}^\top, & \text{if } t < T \\ \boldsymbol{\eta}^\top, & \text{if } t = T, \end{cases} \qquad \boldsymbol{e}_t{}^\top = \begin{cases} \hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top, & \text{if } t < T \\ 0, & \text{if } t = T. \end{cases} \tag{EC.30}$$

In (13) the decisions $\boldsymbol{u}_t$, for $k \le t \le T-1$, will have objective coefficients:

$$\hat{\boldsymbol{X}}_t^\top \boldsymbol{M}^\top + \sum_{\tau=t+1}^{T-1} \left(\hat{\boldsymbol{u}}_\tau^\top \boldsymbol{M} + \boldsymbol{\gamma}^\top\right) \left(\prod_{i=\tau-1,\tau-2,\ldots,t+1} \boldsymbol{A}_i\right) \boldsymbol{B}_t + \boldsymbol{\eta}^\top \left(\prod_{i=T-1,T-2,\ldots,t+1} \boldsymbol{A}_i\right) \boldsymbol{B}_t. \tag{EC.31}$$

This allows calculating the coefficients recursively, just as we did for the constraints. The detailed function CALCULATEOBJECTIVECOEFFICIENTS is defined in Algorithm 2.

Calculation of $\boldsymbol{M}, \boldsymbol{\gamma}$ and $\boldsymbol{\eta}$. Expanding the objective (8), we have:

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

From this equation and the definitions of $v_g(\cdot)$ in Appendix (EC.2) above, we can write $\boldsymbol{M}$ (where the rows are indexed by the controls and the columns indexed by the compartments) as

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

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

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

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Algorithm 2 CALCULATEOBJECTIVECOEFFICIENTS

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

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

Calculate coefficients for $\boldsymbol{u}_\tau$ in the period t summand of (EC.29) as:

$$\begin{cases} \boldsymbol{e}_t{}^\top, & \text{for } \tau = t \\ \boldsymbol{d}_t{}^\top \bar{\boldsymbol{A}}_{t,\tau} \boldsymbol{B}_\tau, & \text{for } \tau = t-1, t-2, \ldots, k \end{cases} \tag{EC.32}$$

where

$$\bar{\boldsymbol{A}}_{t,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = t-1 \\ \bar{\boldsymbol{A}}_{t,\tau+1} \boldsymbol{A}_{\tau+1}, & \text{for } \tau = t-2, t-3, \ldots, k. \end{cases} \tag{EC.33}$$

end for

for $t = T$ **do**

Calculate coefficients for $\boldsymbol{u}_\tau$ in the period T summand of (EC.29) as:

$$\boldsymbol{d}_T{}^\top \bar{\boldsymbol{A}}_{T,\tau} \boldsymbol{B}_\tau, \quad \text{for } \tau = T-1, T-2, \ldots, k \tag{EC.34}$$

where

$$\bar{\boldsymbol{A}}_{T,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = T-1 \\ \bar{\boldsymbol{A}}_{T,\tau+1} \boldsymbol{A}_{\tau+1}, & \text{for } \tau = T-2, T-3, \ldots, k. \end{cases} \tag{EC.35}$$

end for

Sum up the terms attributed to a common $\boldsymbol{u}_\tau$:

$$\boldsymbol{e}_\tau{}^\top + \sum_{t=\tau+1}^{T} \boldsymbol{d}_t{}^\top \bar{\boldsymbol{A}}_{t,\tau} \boldsymbol{B}_\tau. \tag{EC.36}$$

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

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

$$\boldsymbol{\gamma}[S_g] = \boldsymbol{\gamma}[E_g] = \boldsymbol{\gamma}[I_g] = \boldsymbol{\gamma}[R_g] = \begin{cases} -w_g \nu^{\text{fixed}}, & \text{if } g = \text{10-19, ..., 60-69 y.o.} \\ 0, & \text{otherwise.} \end{cases} \tag{EC.37}$$

$$\boldsymbol{\gamma}[R_g^q] = -v_g(\mathbf{1}), \quad \text{for } g \in \mathcal{G} \tag{EC.38}$$

$$\boldsymbol{\gamma}[\cdot] = 0, \quad \text{otherwise.} \tag{EC.39}$$

Finally, $\boldsymbol{\eta}$ (indexed by the compartments for each group) is

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

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

EC.3.4. Specifics of the Iterative Linearization-Optimization Procedure

Having defined functions CALCULATECONSTRAINTCOEFFICIENTS and CALCULATEOBJECTIVECOEFFICIENTS, the LINEARIZATION-OPTIMIZATION function which is the main subroutine of ROLD

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

is described in Algorithm 3. This function builds the linear approximation for the remaining trajectory of the system, and optimizes it via an LP in a trust region of an infinity-norm ϵ-ball around the initial nominal control $\hat{\boldsymbol{u}}^{(k)}_{k:T-1}$. We denote that ϵ-ball by $B^{\infty}_{\epsilon}(\hat{\boldsymbol{u}}^{(k)}_{k:T-1})$.

Algorithm 3 LINEARIZATION-OPTIMIZATION

Require: time period k, starting state $\boldsymbol{X}_k$, nominal control initialization $\hat{\boldsymbol{u}}^{(k)}_{k:T-1}$
Calculate nominal trajectory $\hat{\boldsymbol{X}}_{k:T}$ from $\boldsymbol{X}_k$ and $\hat{\boldsymbol{u}}^{(k)}_{k:T-1}$
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 $\boldsymbol{u}_{k:T-1} \in B^{\infty}_{\epsilon}(\hat{\boldsymbol{u}}^{(k)}_{k:T-1})$ to obtain controls $\boldsymbol{\zeta}_{k:T-1}$
return $\boldsymbol{\zeta}_{k:T-1}$.

The full ROLD procedure is given in Algorithm 4. Within each period k, ROLD calls the LINEARIZATION-OPTIMIZATION 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 to choose an initialization of the $k=0$ nominal control $\hat{\boldsymbol{u}}^{(0)}_{0:T-1}$; in our experiments we initialize this with a heuristic solution generated via a gradient method.

For the termination condition, we combine a fixed upper bound on the number of iterations with a condition that we do not repeat the control sequences produced by LINEARIZATION-OPTIMIZATION, 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 $\boldsymbol{u}_{k:T-1}$ can be changed to any value in $[0,1]$ with ϵ-length steps, i.e., the upper bound is at least $\lceil \frac{1}{\epsilon} \rceil$. We further multiply the allowed number of iterations by a multiple $mult \geq 1$, fixing the upper bound to be $mult \lceil \frac{1}{\epsilon} \rceil$.

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. Higher values of $mult$ improved the heuristic only up to around $mult=2$, after which point the non-cycling termination condition was triggered almost always. On the other hand, reducing ϵ had 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. In particular, for all the runs presented we take $\epsilon=0.05$ and $mult=2$, resulting in an upper bound of 40 runs for the inner loop.

EC.3.5. Initialization for ROLD

In this subsection we flesh out the details for the nominal control initialization $\hat{\boldsymbol{u}}^{(0)}_{0:T-1}$. Our implementation of the initialization varies depending on the ROLD variant (i.e., the level of targeting):

- ROLD NO-TARGET: ROLD is initialized at a solution of a gradient descent method that is constrained to have the same activity levels across all age groups and activities, which can vary through time.
- ROLD AGE: A gradient descent method that is constrained to have the same activity levels across all activities for a given age group (and the activity levels are allowed to vary through time) is initialized at the ROLD NO-TARGET solution. Then ROLD is initialized at the solution of the gradient descent method.
- ROLD ACT: A gradient descent method that is constrained to have the same activity levels across all age groups in a given activity (and the activity levels are allowed to vary through time) is initialized at the ROLD NO-TARGET solution. Then ROLD is initialized at the solution of the gradient descent method.
- ROLD AGE-ACT: We keep the best solution out of:

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Algorithm 4 ROLD

Require: initial state $\boldsymbol{X}_0$ and initial nominal control $\hat{\boldsymbol{u}}_{0:T-1}^{(0)}$
for $k=0,\ldots,T-1$ **do**
$i:=0$ and $\boldsymbol{u}_{k:T-1}^{(k,0)}:=\hat{\boldsymbol{u}}_{k:T-1}^{(k)}$
while $i\leq mult\lceil\frac{1}{\epsilon}\rceil$ and $\boldsymbol{u}_{k:T-1}^{(k,i)}\neq\boldsymbol{u}_{k:T-1}^{(k,j)}$, $\forall 0\leq j<i$ **do**
$\boldsymbol{u}_{k:T-1}^{(k,i+1)}=$ LINEARIZATION-OPTIMIZATION$(k,\boldsymbol{X}_k,\boldsymbol{u}_{k:T-1}^{(k,i)})$
$i=i+1$
end while
Set $\boldsymbol{u}_{k:T-1}^{*,(k)}=\arg\min_{\boldsymbol{u}_{k:T-1}\in\left\{\boldsymbol{u}_{k:T-1}^{(k,0)},\ldots,\boldsymbol{u}_{k:T-1}^{(k,i-1)}\right\}}\mathsf{Total\ Loss}(\boldsymbol{u}_{0:k-1}^{*},\boldsymbol{u}_{k:T-1})$
Set the nominal control sequence to $\hat{\boldsymbol{u}}_{k+1:T-1}^{(k+1)}=\boldsymbol{u}_{k+1:T-1}^{*,(k)}$
Set $\boldsymbol{u}_k^{*}=\boldsymbol{u}_k^{*,(k)}$ and update the system state:

$$\boldsymbol{X}_{k+1}=\boldsymbol{X}_k+\boldsymbol{F}_k\left(\boldsymbol{X}_k,\boldsymbol{u}_k^{*}\right). \tag{EC.40}$$

end for

(i) A gradient descent method that can vary activity levels across activities, age groups, and time is initialized at the ROLD AGE solution. Then ROLD is initialized at the solution of the gradient descent method.

(ii) A gradient descent method that can vary activity levels across activities, age groups, and time is initialized at the ROLD ACT solution. Then ROLD is initialized at the solution of the gradient descent method.

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

EC.4.1. Basic SEIR Model Parameters

The SEIR model parameters that are constant across age groups are summarized in Table EC.1. The age-group specific parameters are reported in Table EC.2. We start with the parameters as reported in Salje et al. (2020)[12], 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 so as to model changes in how hospitals manage COVID-19 patients and changes in mandates for using masks and other measures that reduce transmission (details are presented in Section EC.4.2).

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 $4\pm0.8\cdot0.6$ days, where 0.6 days is half the width of the 95% confidence interval for the incubation period reported in Bi et al. (2020), and 0.8 accounts for the fact that the stay in compartment E is 4/5 of the mean incubation time in Salje et al. (2020). 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 Du et al. (2020).[13] 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 (Salje et al. 2020).

[12] We retrieve the parameter values as reported before Salje et al. (2020) updated them on July 8, 2020.

[13] We note that Du et al. (2020) 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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.1 SEIR model parameters

Variable	Description	Value [Uncertainty Range]	Reference
R_0	Basic reproduction number	2.9 [2.8 - 2.99]	Salje et al. (2020)
σ^{-1}	Latency period	4.0 days [3.52 - 4.48]	Salje et al. (2020), Bi et al. (2020)
μ^{-1}	Infectious period	4.0 days [3.57 - 4.43]	Salje et al. (2020), Du et al. (2020)
$(\lambda^H)^{-1}$	Average time spent in Hospital	14.94 days [14.55 - 15.32]	Salje et al. (2020)
$(\lambda^{ICU})^{-1}$	Average time spent in ICU	21.96 days [21.38 - 22.6]	Salje et al. (2020)

Table EC.2 Age-group specific SEIR model probability parameters

Age group g (y.o.)	$p_{ss,g}$ 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:

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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

$$\boldsymbol{\Phi} = \beta \cdot \left[\begin{array}{c|c} \mathbf{0} & [c_{gh}]_{g,h\in\mathcal{G}} \\ \hline \mathbf{0} & \mathbf{0} \end{array}\right] \tag{EC.41}$$

and

$$\boldsymbol{\Lambda} = \left[\begin{array}{c|c} -\mathbf{diag}(\boldsymbol{\sigma}) & \mathbf{0} \\ \hline \mathbf{diag}(\boldsymbol{\sigma}) & -\mathbf{diag}(\boldsymbol{\mu}) \end{array}\right]. \tag{EC.42}$$

Then R_0 can be identified as the spectral radius (i.e., the largest absolute value of the eigenvalues) of the matrix $-\boldsymbol{\Phi}\boldsymbol{\Lambda}^{-1}$ (Diekmann et al. 2010, Perasso 2018). 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 $(-\boldsymbol{\Phi}/\beta)\,\boldsymbol{\Lambda}^{-1}$.

EC.4.2. Epidemiological Model Parameter Fitting Using Health Outcomes and Mobility Data

We use data on health outcomes from the French Public Health Agency (French Government 2020), as well as Google mobility data (Google 2020), 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 who 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 $\mathcal{P}$.

• Date of patient zero. We assume that the SEIR process starts with an infected individual of the 40-49 y.o. age group (Mohammad 2020). We wish to estimate the date when this infection occurs.

• Epidemiological parameters. We use the epidemiological parameters of Salje et al. (2020) to initialize the SEIR model. We allow these parameters to change in time in order to model changes in the way hospitals manage COVID-19 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 EC.4.1), according to the relationship

$$s_{\text{after_}d} = m_s \cdot s_{\text{before_}d}, \quad s \in \mathcal{S}, \tag{EC.43}$$

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 (Google 2020). 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 EC.3.

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 until October 21 2020. 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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.3 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

- Social mixing parameter. To reduce the number of parameters to be calibrated, we simplify our mixing model in (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 3.2). 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:

$$\tilde{\boldsymbol{\ell}}_g(t) = \bar{\boldsymbol{\ell}}_g(t)(\mathbf{1} + \boldsymbol{Z}(t)), \quad \boldsymbol{Z}(t) \sim N(\mathbf{0}, \sigma^2 \mathbb{I}), \quad \sigma = 10^{-2}. \tag{EC.44}$$

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, respectively. We denote these quantities with $H_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}), ICU_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}), D_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}})$ to emphasize the dependence of the SEIR process on parameters $\mathcal{P}$ and the vector of activity levels $\tilde{\boldsymbol{\ell}} := [\tilde{\boldsymbol{\ell}}_g]_{g \in \mathcal{G}}$. We aggregate these quantities over all age groups:

$$H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \sum_{g \in \mathcal{G}} H_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) \tag{EC.45}$$

$$ICU^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \sum_{g \in \mathcal{G}} ICU_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) \tag{EC.46}$$

$$D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \sum_{g \in \mathcal{G}} D_g^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) \tag{EC.47}$$

Denote by $H^{obs}(t), ICU^{obs}(t), D^{obs}(t)$ the general ward hospital beds utilization, ICU beds utilization, and cumulative deaths, respectively, at time t, as observed in the real data for Île-de-France from the French Public Health Agency (French Government 2020). We calculate the relative fitting error of the SEIR model at time t for each of these three quantities as

$$\mathcal{E}_H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \left| \frac{H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) - H^{obs}(t)}{H^{obs}(t)} \right| \tag{EC.48}$$

$$\mathcal{E}_{ICU}^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \left| \frac{ICU^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) - ICU^{obs}(t)}{ICU^{obs}(t)} \right| \tag{EC.49}$$

$$\mathcal{E}_D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) := \left| \frac{D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) - D^{obs}(t)}{D^{obs}(t)} \right|. \tag{EC.50}$$

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

$$\mathcal{E}^{\mathcal{P}} := \mathbb{E}_{\tilde{\boldsymbol{\ell}}} \Bigg[\sum_{t \in \text{peak_H}} w^{\text{peak}} \cdot \mathcal{E}_H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + \sum_{t \in \text{peak_ICU}} w^{\text{peak}} \cdot \mathcal{E}_{ICU}^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + \sum_{t \in \text{peak_D}} w^{\text{peak}} \cdot \mathcal{E}_D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}})$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

$$+ \sum_{t \in \text{tail}} \left(w_H^{\text{tail}} \cdot \mathcal{E}_H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + w_{ICU}^{\text{tail}} \cdot \mathcal{E}_{ICU}^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + w_D^{\text{tail}} \cdot \mathcal{E}_D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) \right)$$

$$+ \sum_{t \in \text{main}} \left((1 - w_H^{\text{tail}}) \cdot \mathcal{E}_H^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + (1 - w_{ICU}^{\text{tail}}) \cdot \mathcal{E}_{ICU}^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) + (1 - w_D^{\text{tail}}) \cdot \mathcal{E}_D^{\mathcal{P}}(t, \tilde{\boldsymbol{\ell}}) \right) \Bigg], \tag{EC.51}$$

where the expectation is taken with respect to random vector $\tilde{\boldsymbol{\ell}}$, and where the time intervals are defined in Table EC.4 and comprise the entire period between March 17 2020 and October 21 2020. Our approach penalizes the errors at the peak times for hospital beds utilization, ICU 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 $w_H^{\text{tail}} = w_D^{\text{tail}} = 1/3$ and $w_{ICU}^{\text{tail}} = 2/3$, with a higher weight for ICU beds utilization so as to ensure low error in the tail predictions of ICU utilization, as ICU beds utilization towards the tail of the calibration window plays an important role in the dynamics of the model right after the calibration window.

Table EC.4 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 $\mathcal{P}$ that minimize $\mathcal{E}^{\mathcal{P}}$ in (EC.51). We approximate the expectation in (EC.51) through a Monte Carlo sample-average approximation, using 100 samples. The set $\mathcal{P}$ contains both discrete and continuous parameters. To minimize $\mathcal{E}^{\mathcal{P}}$, we first do a grid search over all possible combinations for the discrete parameters, and then for each such combination, we perform a gradient descent procedure over the space of the continuous parameters. Each parameter is optimized within an allowed range, which is informed from the context of Île-de-France in 2020. We used a wide allowed range if we had no information on what could be reasonable values for a given parameter.

Our calibration procedure yields the parameter fitting summarized in Table EC.5. We compare the fitted values of the SEIR model with the values reported by the French Public Health Agency in Figure 2. Figure EC.1 presents the results from the out-of-sample model validation analysis.

EC.4.3. Economic Model Parameter Fitting

We obtain data on population, employment, and wages from the French National Institute of Statistics and Economic Studies (Institut national de la statistique et des études économiques — INSEE). Where relevant, we discount all cash flows at a 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 INSEE (2020).

Estimation of w_g. Recall that w_g in (EC.18) corresponds to the employment value for a member of group g, under normal conditions. To estimate w_g, we use two datasets from INSEE:

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.5 Calibration: fitted epidemiological and economic parameters and allowed estimation ranges

Parameter	Estimated Value	Allowed range
Date of patient zero	Jan 1 2020	Dec 8 2019 - Jan 27 2020
Date d when epidemiological parameters change	Apr 16 2020	Mar 1 2020 - Apr 30 2020
Multiplier m_β	0.60	0.50-2.00
Multiplier $m_{p_g^{ICU}}$	1.66	0.50-2.00
Multiplier $m_{p_g^D}$	1.13	0.50-2.00
Multiplier m_{λ_H}	1.50	0.50-2.00
Multiplier $m_{\lambda_{ICU}}$	2.00	0.50-2.00
Weight for activities α_{other}	0.96	0.00-1.00
School activity ℓ^{school} Mar 17 - May 10 2020	0.00	0.00-0.10
School activity ℓ^{school} May 11 - Jun 30 2020	0.00	0.00-1.00
School activity ℓ^{school} Jul 1 - Aug 31 2020	0.20	0.00-0.20
School activity ℓ^{school} Sep 1 - Oct 21 2020	1.00	0.50-1.00
Social mixing α	0.39	0.10-2.00
Sensitivity of econ. value ν^{work}	0.50	
Sensitivity of econ. value ν^{fixed}	0.47	

• Yearly full time equivalent (FTE)[14] 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" (INSEE 2016b).

• 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" (INSEE 2019).

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 INSEE (2019) 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 midpoint 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 INSEE (2016b); 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.

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

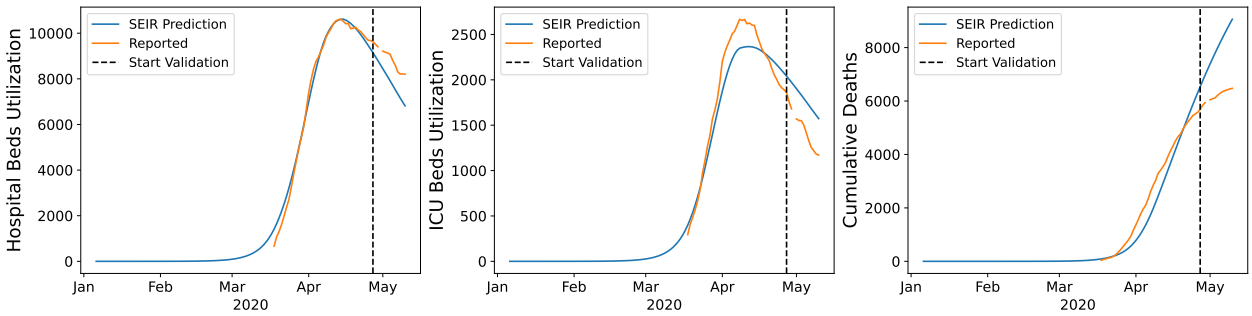

(a) Calibrate until April 26 2020; validate until May 10 2020

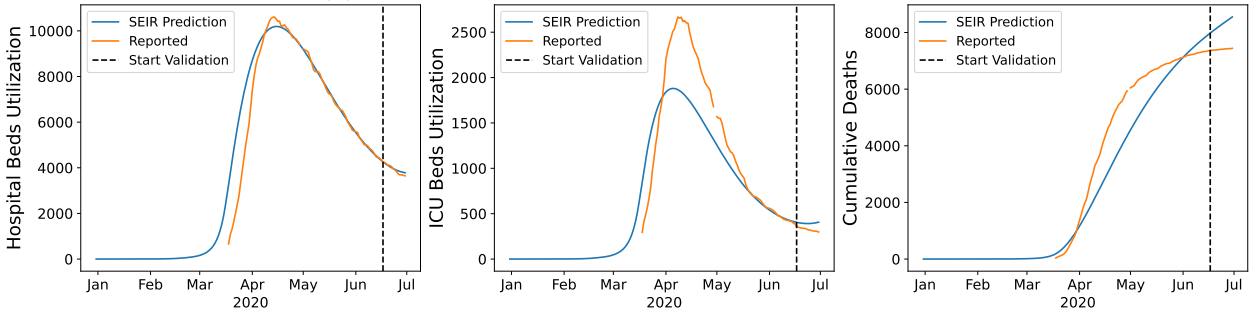

(b) Calibrate until June 16 2020; validate until June 30 2020

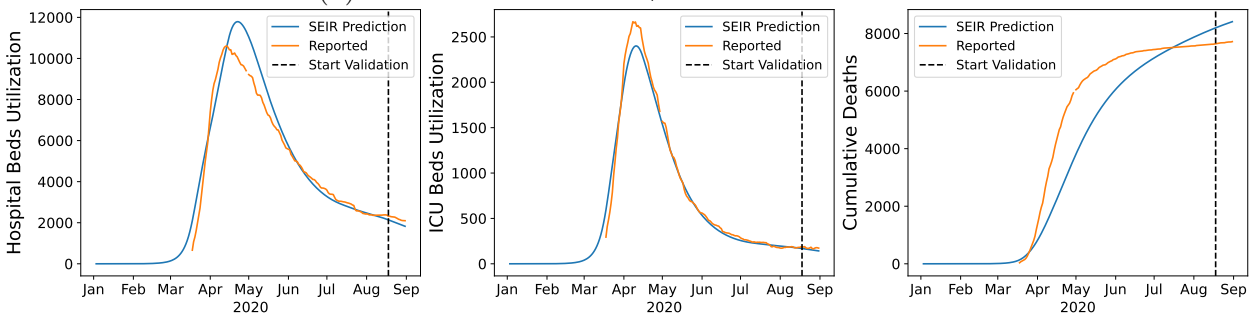

(c) Calibrate until August 17 2020; validate until August 31 2020

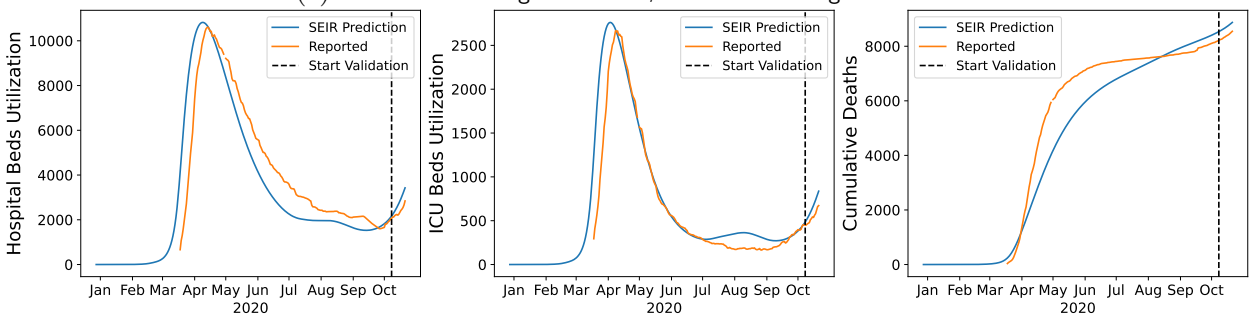

(d) Calibrate until October 7 2020; validate until October 21 2020

Figure EC.1 **Out-of-sample predictions of the fitted SEIR model vs. reported values by the French Public Health Agency for hospital beds utilization (left), ICU beds utilization (middle) and cumulative deaths (right). For each of four dates t ((a): May 10 2020; (b): June 30 2020; (c): August 31 2020; (d): October 21 2020) we calibrate the SEIR model using data up until day $t-14$; and then run the SEIR model up to day t. The dashed line indicates the start of the out-of-sample validation (i.e., $t-13$).**

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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

$$\text{average_wage}_g = \text{employed_average_wage}_g \times \text{employment_rate}_g \times \text{fraction_work_eligible_population}_g. \tag{EC.52}$$

• 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 scale all wages average_wage_g proportionally so as to retrieve the real total wage amount w_g.

Table EC.6 summarizes the year-based employment contribution parameters per age group. We note that when using them in the objective of the optimization problem, we divide these year-based values by 365, in order to capture employment value on a daily basis.

Table EC.6 Year-based economic value parameters under normal activity, per age group (in €)

Age Group g (y.o.)	$w_g \times 365$ (€)
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 (EC.18). 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 from the Bank of France. This is sentiment data where managers are asked to compare current business conditions to normal conditions for the same relevant time period (Banque de France 2020a,b). 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 (INSEE 2016a) 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 (Google 2020), as explained in Section EC.4.2. To simplify the estimation, we set $\nu^{\text{other activities}} = 0$ and then determine parameters $\nu^{\text{work}}, \nu^{\text{fixed}}$ solving the system of equations

$$\nu^{\text{work}} + \nu^{\text{fixed}} = 1 \tag{EC.53}$$

$$\nu^{\text{work}} \ell^{\text{work}}_{\text{April}} + \nu^{\text{fixed}} = 0.5851, \tag{EC.54}$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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 (EC.53)-(EC.54), and adjust $\nu^{\text{work}} = 1 - \nu^{\text{other activities}} - \nu^{\text{fixed}}$. The specific values we test are $\nu^{\text{other activities}} \in \{0, 0.1, 0.2\}$.

EC.5. Benchmark Policies

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 Duque et al. (2020) 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 hospital admissions (i.e., inflow of patients admitted into the hospital due to COVID-19 complications) exceeds a pre-determined threshold, and then changes the confinement to a relaxed level when the average seven-day hospital admissions and the hospital utilization rate drop below pre-determined thresholds.

Since Duque et al. (2020) 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 5.

Algorithm 5 ICU Admissions Trigger Policy — ICU-t

Require: parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{admissions}}$, $\rho_{\text{utilization}}$
  Initialize $\ell^a_g(0) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$
  for $t \in \{0, \ldots, T-1\}$, $g \in \mathcal{G}$ **do**
    if 7 day average ICU admissions $> \rho_{\text{admissions}}$ **then**
      Set $\ell^a_g(t) = \ell_{\text{strict}}$ 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 $\leq \rho_{\text{admissions}}$) AND $\left(\sum_g ICU_g(t)/K^{ICU}(t) \leq \rho_{\text{utilization}}\right)$
  then
      Set $\ell^a_g(t) = \ell_{\text{relaxed}}$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$
    else
      Set $\ell^a_g(t) = \ell^a_g(t-1)$ for all $a \in \{\text{work}, \text{school}, \text{transport}, \text{leisure}, \text{other}\}, g \in \mathcal{G}$
    end if
  end for

We optimize over the parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{admissions}}$ and $\rho_{\text{utilization}}$ via grid search with the goal of minimizing the objective in (8) corresponding to the total economic and death loss due to the pandemic, and we report the performance of the best policy.

• **Hybrid Trigger Policy — Hybrid-t**. This policy resembles the rule used in France for declaring a region in "maximum alert"[15] (Lehot and Borgne 2020). Like the previous policy, this policy also switches between a (uniform) strict and a more relaxed confinement level, but the trigger condition combines ICU utilization with the rate of new infections in the population. Specifically, the policy switches to strict confinement if the average seven-day incidence rate in the population, defined[16] as $\sum_{t-6 \leq \tau \leq t} \sum_g \text{new infections}_g(\tau) / \sum_g N_g(0)$, is greater than a threshold $\rho_{\text{incidence}}$,

[15] 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.

[16] 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.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

and the incidence rate in age groups corresponding to the population that is 60 y.o. and above, $\sum_{t-6\leq\tau\leq t}\sum_{g\geq 60\text{ y.o.}}\text{new infections}_g(\tau)/\sum_{g\geq 60\text{ y.o.}}N_g(0)$, is greater than a threshold $\rho_{\text{incidence_60+}}$, *and* the ICU utilization rate is greater than a threshold $\rho_{\text{utilization}}$. We optimize over all parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{incidence}}$, $\rho_{\text{incidence_60+}}$ and $\rho_{\text{utilization}}$ with the goal of minimizing the total loss objective in (8), and we report the performance of the best policy.

This is the Hybrid-t AND policy, and it is described in Algorithm 6. 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*, as the trigger for setting the strict confinement level.

Algorithm 6 HYBRID TRIGGER POLICY — Hybrid-t AND

Require: parameters ℓ_{strict}, ℓ_{relaxed}, $\rho_{\text{incidence}}$, $\rho_{\text{incidence_60+}}$, $\rho_{\text{utilization}}$
  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 $t\in\{0,\ldots,T-1\}$, $g\in\mathcal{G}$ **do**
    if $\left(\sum_{t-6\leq\tau\leq t}\sum_g \text{new infections}_g(\tau)/\sum_g N_g(0)>\rho_{\text{incidence}}\right)$ AND
  $\left(\sum_{t-6\leq\tau\leq t}\sum_{g\geq 60y.o.}\text{new infections}_g(\tau)/\sum_{g\geq 60y.o.}N_g(0)>\rho_{\text{incidence_60+}}\right)$ AND
  $\left(\sum_g ICU_g(t)/K^{ICU}(t)>\rho_{\text{utilization}}\right)$ **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 $a\in\{\text{work},\text{school},\text{transport},\text{leisure},\text{other}\}, g\in\mathcal{G}$
    end if
  end for

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

EC.6. Additional Results

EC.6.1. 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 EC.7. For each parameter, we sample 40 values uniformly at random from a 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.

Figure EC.2 shows robustness results for seven values of the economic cost of death χ: $[0,10,15,25,50,100,150]\times$ 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 from Section 6 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.

EC.6.2. Can Dual Targeting Increase Average Activity Levels for Each Age Group?

For each ROLD policy, we calculate the time-average activity level of each age group, averaged over the activities relevant to that age group (listed in Table EC.8), and weighting different activities by the respective participation rates. We split the 24-hour day in 48 slots of 30 minutes each, indexed $\tau=1,\ldots,48$, and let $\mathcal{R}_g^a(\tau)$ denote the participation rate of age group g in activity a during time slot τ of a generic day (so that $\sum_{a\in\mathcal{A}}\mathcal{R}_g^a(\tau)=1,\forall\tau,g$). We retrieve the participation rates from time-use survey data we obtained from the French National Archive of Data from Official

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.7 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	0.76-1.00
School activity ℓ^{school} Mar 17 - May 10 2020	0.00	0.00-0.20
School activity ℓ^{school} May 11 - Jun 30 2020	0.00	0.00-0.20
School activity ℓ^{school} Jul 1 - Aug 31 2020	0.20	0.00-0.40
School activity ℓ^{school} Sep 1 - Oct 21 2020	1.00	0.80-1.00
Social mixing α	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%

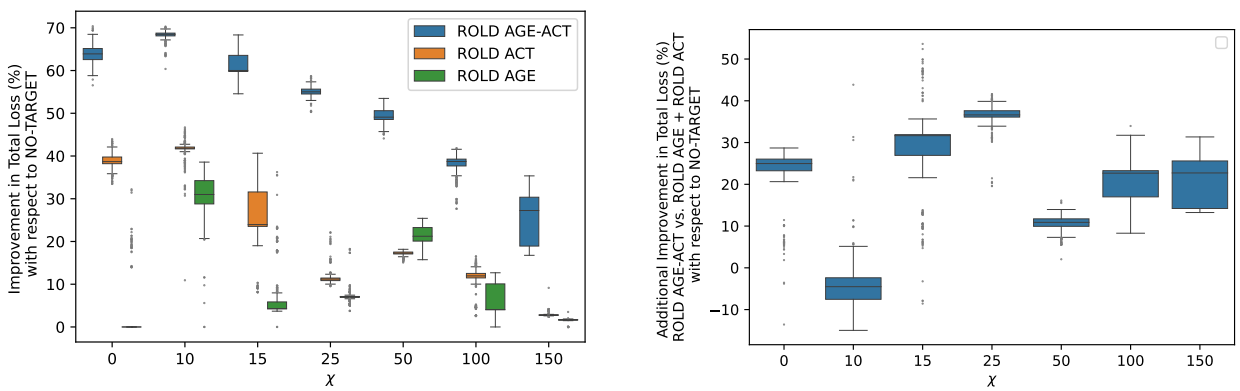

(a) Comparison between ROLD policies with different levels of targeting in terms of total losses

(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

Figure EC.2 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.

Statistics (ADISP, INSEE 2010). We detail our procedure for estimating participation rates in Appendix EC.8.1. Then, for group 0-9 y.o., the time-average activity level is defined as

$$\frac{1}{T'}\sum_{t=1}^{T'}\sum_{a\in\{\text{school,leisure,transport,other}\}}\frac{\sum_\tau \mathcal{R}_g^a(\tau)}{\sum_{a'\in\{\text{school,leisure,transport,other}\}}\sum_\tau \mathcal{R}_g^{a'}(\tau)}\ell_g^a(t), \quad \text{if } g = \text{0-9 y.o.} \tag{EC.55}$$

and similarly for other age groups.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.8 Relevant activities for each age group, excluding home

Age Group g (y.o.)	Relevant Activities
0-9	school, leisure, transport, other
10-19	school, work, leisure, transport, other
20-29	work, leisure, transport, other
30-39	work, leisure, transport, other
40-49	work, leisure, transport, other
50-59	work, leisure, transport, other
60-69	work, leisure, transport, other
70-79	leisure, transport, other
80+	leisure, transport, other

EC.6.3. Additional Details for How Gains Arise from Dual Targeting

Figure EC.3 shows the activity levels of different (age group, activity) pairs, with pairs divided into three buckets (high, medium, low) in terms of their econ-to-contacts-ratio. The ROLD policy tends to enforce stricter confinements in earlier periods and subsequently relaxes these through time.

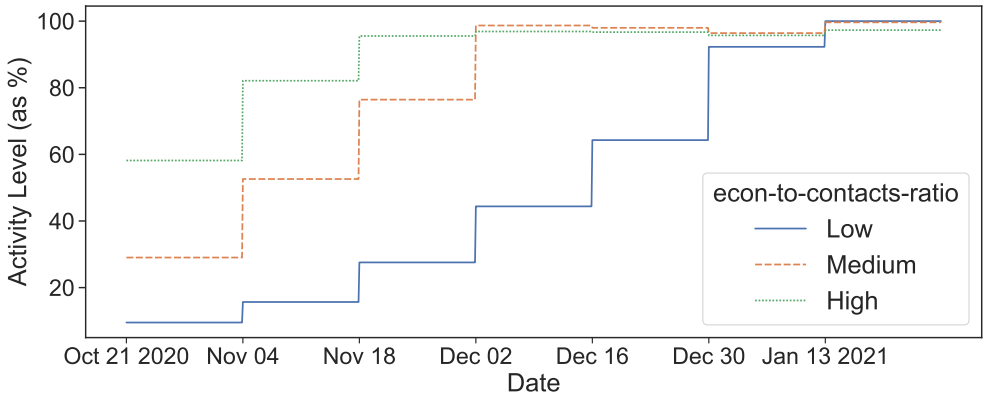

Figure EC.3 The activity levels of optimized ROLD AGE-ACT policies for different (age group, activity) pairs, bucketed in three equally-sized buckets according to their econ-to-contacts-ratio. The lines indicate the mean for each bucket, where for each time t the mean is taken over the different values of the cost of death χ and over the relevant (age group, activity) pairs.

Figure EC.4 visualizes the optimized confinement policy for the value $\chi = 60\times$, which is in the mid-range of estimates used in the economics literature on COVID-19 (Alvarez et al. 2020) and is representative of the overall behavior we observe across all experiments. Groups 20-69 y.o. remain more open in work but face confinement in leisure for up to the first ten weeks; on the contrary, 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 also remain open in leisure.

EC.6.4. Additional Details for the Regression Decision Trees

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

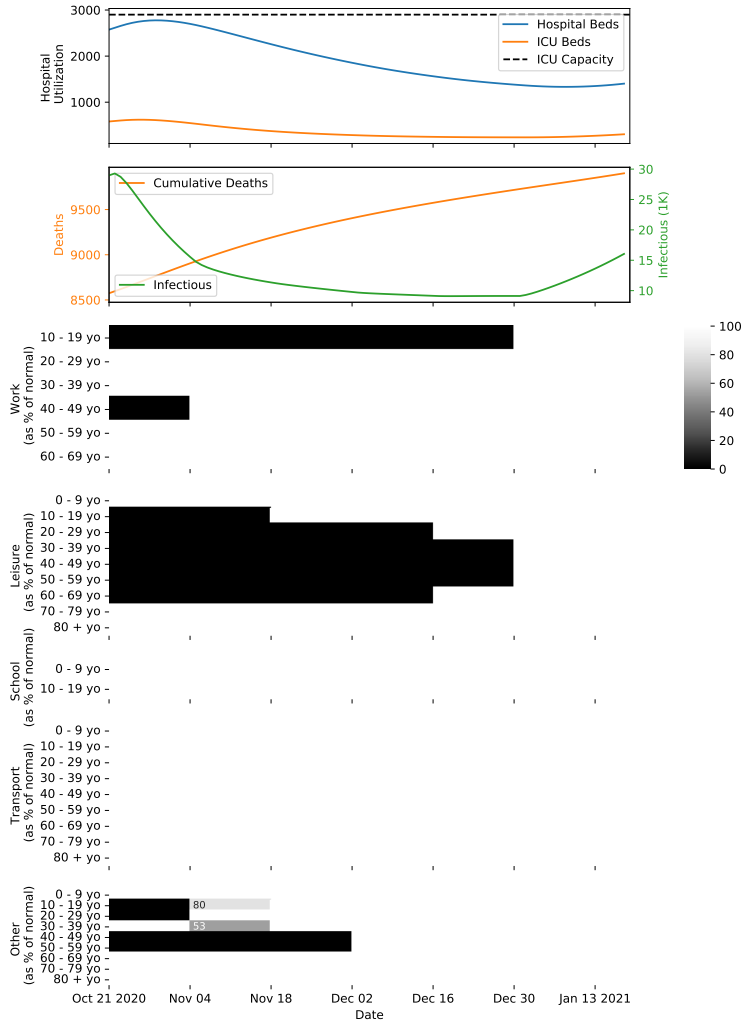

Figure EC.4 **The optimized ROLD AGE-ACT policy for a problem instance with a 90-day optimization horizon starting on October 21, 2020, with 2900 ICU beds, and with cost of death $\chi = 60\times$. From top to bottom, the seven panels depict the time evolution for the occupation of hospital and ICU beds, the number of infectious individuals, the cumulative number of deaths in the population, and the confinement policy imposed by ROLD in each age group and activity. In panels 3-7, the values correspond to the activity levels allowed and are color-coded so that darker shades capture a stricter confinement.**

To build a training set, we first create a larger set of problem instances built using a wide range of problem parameters. We vary the following parameters: basic reproduction number R_0, the scaling of the group-specific probabilities $p_{ss,g}$ of having severe symptoms conditioned on being in state I , the social mixing parameter α, the cost of death χ, the ICU capacity K^{ICU}, the economic model parameter ν^{fixed}, the ratio of economic model parameters $\nu^{\text{other activities}}/\nu^{\text{work}}$, and the activity weights $w^{\text{leisure}}, w^{\text{transport}}, w^{\text{other}}$ used in the model for $v_g^{\text{employment}}$ in (EC.18). We create each instance by drawing each of these parameters independently and uniformly at random from a range that is specified in Table EC.9.

All resulting instances use an optimization horizon of $T' = 90$ days and change the confinement decisions every 14 days. For each problem instance, we compute the optimal ROLD AGE-ACT, ACT, and NO-TARGET activity levels, i.e., the decisions $\ell_g^a(t)$ for $g \in \mathcal{G}, a \in \mathcal{A}, 0 \le t \le T-1$, which we then use to simulate the SEIR dynamics and calculate all the corresponding values of the SEIR states.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.9 Parameter variation used for fitting trees. Each problem instance is defined by drawing each parameter independently from the uniform distribution over the indicated range for that parameter.

Parameter	Description	Min	Max	Notes
R_0	Basic reproduction number	1	8	
severity-scaling	Scales the probability of $p_{ss,g}$ for all groups g	0.1	3.125	Max value ensures $p_{ss,80+} = 1$
α	Social mixing parameter	0.16	1.16	Min/max values correspond resp. to contacts dropping to 80% and 20% of their level
χ	Cost of death	0	400	In multiples of GDP/capita
K^{ICU}	ICU capacity	592	4175	Min(max) corresponds to the ICU filling completely in 10 (60) days from day 0 of the optimization horizon, assuming a fully open policy
ν^{fixed}	Economic model parameter	0.1	0.9	
$\frac{\nu^{\text{other activities}}}{\nu^{\text{work}}}$	Ratio of economic model parameters	0.01	0.36	The estimated ratio of 3/50 was scaled up and down by a factor of 6
$w^{\text{leisure}}, w^{\text{transport}}, w^{\text{other}}$	Weights for econ. model for $v_g^{\text{employment}}$ in (EC.18)	-	-	Sample uniformly at random from the simplex defined by $w^{\text{leisure}} + w^{\text{transport}} + w^{\text{other}} = 1$; $w^{\text{leisure}}, w^{\text{transport}}, w^{\text{other}} \geq 0$

EC.6.4.1. Additional Details for the Regression Decision Trees for the ROLD AGE-ACT Policy We create a training set of 5,524,785 samples — with one sample for every 14-day period (i.e., for $t = 0, 14, 28$, etc.), 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 AGE-ACT decision $\ell_g^a(t)$. We consider both targeted features, by which we mean features that differentiate on either group g or activity a,[17] or on both, and non-targeted features (Tables EC.10, EC.11). Some of the non-targeted features are allowed to depend on t.

Using this data as training set, we employ the CART algorithm (implemented in the SCIKIT-LEARN Python library) to train regression decision trees to predict the optimal ROLD AGE-ACT confinement decisions across all activities and age groups as a function of the considered features, using the traditional mean-squared-error (MSE) criterion as a goodness of fit metric. When using k-fold cross-validation to assess the optimal depth of the tree, we noticed that the optimal depth exceeded 10. As such, we decided to train a tree of depth four for interpretability purposes, and we also trained trees with maximum depth ranging from four to ten to test the robustness of our results.

In addition, we also report in Figure EC.5 the tree of depth five trained using the features described above. Beyond appearing in the first four levels of the tree (which is consistent with Figure 10), it is worth noting that the econ-to-contacts-ratio continues to be used for splits in the fifth level, which further supports the importance of that feature.

[17] Our targeted features include age group and activity themselves. To handle the categorical nature of activities, we create dummy indicator variables for each a.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.10 Targeted features used for fitting trees for ROLD AGE-ACT

Feature Name	Description
g, a	Age group, activity
$p_{ss,g}$	Probability of having severe symptoms conditional on infection for group g
p_g^D	Probability of dying conditional on severe symptoms for group g
v_g^{life}	Lost wages for group g
$N_g(t)$	$S_g(t) + E_g(t) + I_g(t) + R_g(t)$
econ-gradient(g, a, t)	The gradient of the economic value with respect to the activity level $\ell_g^a(t)$ normalized by $N_g(t)$, i.e., $\frac{\partial \text{Economic Value}(t)/\partial \ell_g^a(t)}{N_g(t)}$
contacts-given(g, a)	The number of social contacts given to other groups by group g in activity a, i.e., $\sum_{h\in\mathcal{G}} C_{h,g}^a$
econ-to-contacts-ratio(g, a, t)	The ratio of econ-gradient(g, a, t) and contacts-given(g, a)

Table EC.11 Non-targeted features used for fitting trees for ROLD AGE-ACT

Feature Name	Description
t	Time period
R_0	Basic reproduction number
contacts-elasticity	The percentage drop in total contacts when activity levels drop by 50%. Equal to $1 - 4^{-\alpha}$, where α is the social mixing parameter.
χ	Cost of death
K^{ICU}	ICU capacity
R-perc(t)	The percentage of the population estimated by the model to be in the recovered state at time t, i.e., $\sum_{g\in\mathcal{G}} R_g(t)/[\sum_{g\in\mathcal{G}} N_g(t) + R_g^q(t)]$
ICU-utilization(t)	The percentage of ICU beds occupied at time t
infection-rate-new-seven-day-avg(t)	The rate of new infections averaged over the last seven days, i.e., $\frac{1}{7}\sum_{\tau\in[t-6,t]}\sum_{g\in\mathcal{G}}[\sigma E_g(\tau)]/[\sum_{g\in\mathcal{G}} N_g(\tau) + R_g^q(\tau)]$
infection-rate-new-elderly-seven-day-avg(t)	The rate of new infections in elderly groups averaged over the last seven days, i.e., $\frac{1}{7}\sum_{\tau\in[t-6,t]}[\sum_{g\in\mathcal{G}\,:\,g\geq 60y.o.} \sigma E_g(\tau)]/[\sum_{g\in\mathcal{G}\,:\,g\geq 60y.o.} N_g(\tau) + R_g^q(\tau)]$
ICU-daily-admissions(t)	New daily admissions into the ICU at time t, i.e., the sum of the second and third term in the right-hand-side of (EC.8) over all age groups
ICU-daily-admissions-seven-day-avg(t)	Seven-day average of new daily admissions into the ICU, i.e., ICU-daily-admissions(t) averaged over the last seven days

Lastly, we perform a few other tests of feature importance. Figure EC.6 reports permutation importance for a very high depth (equal to 10) tree. While at this higher depth there are more features with non-zero permutation importance, econ-to-contacts-ratio continues to overwhelmingly dominate the others. We also train another set of depth-four trees to verify this finding through

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

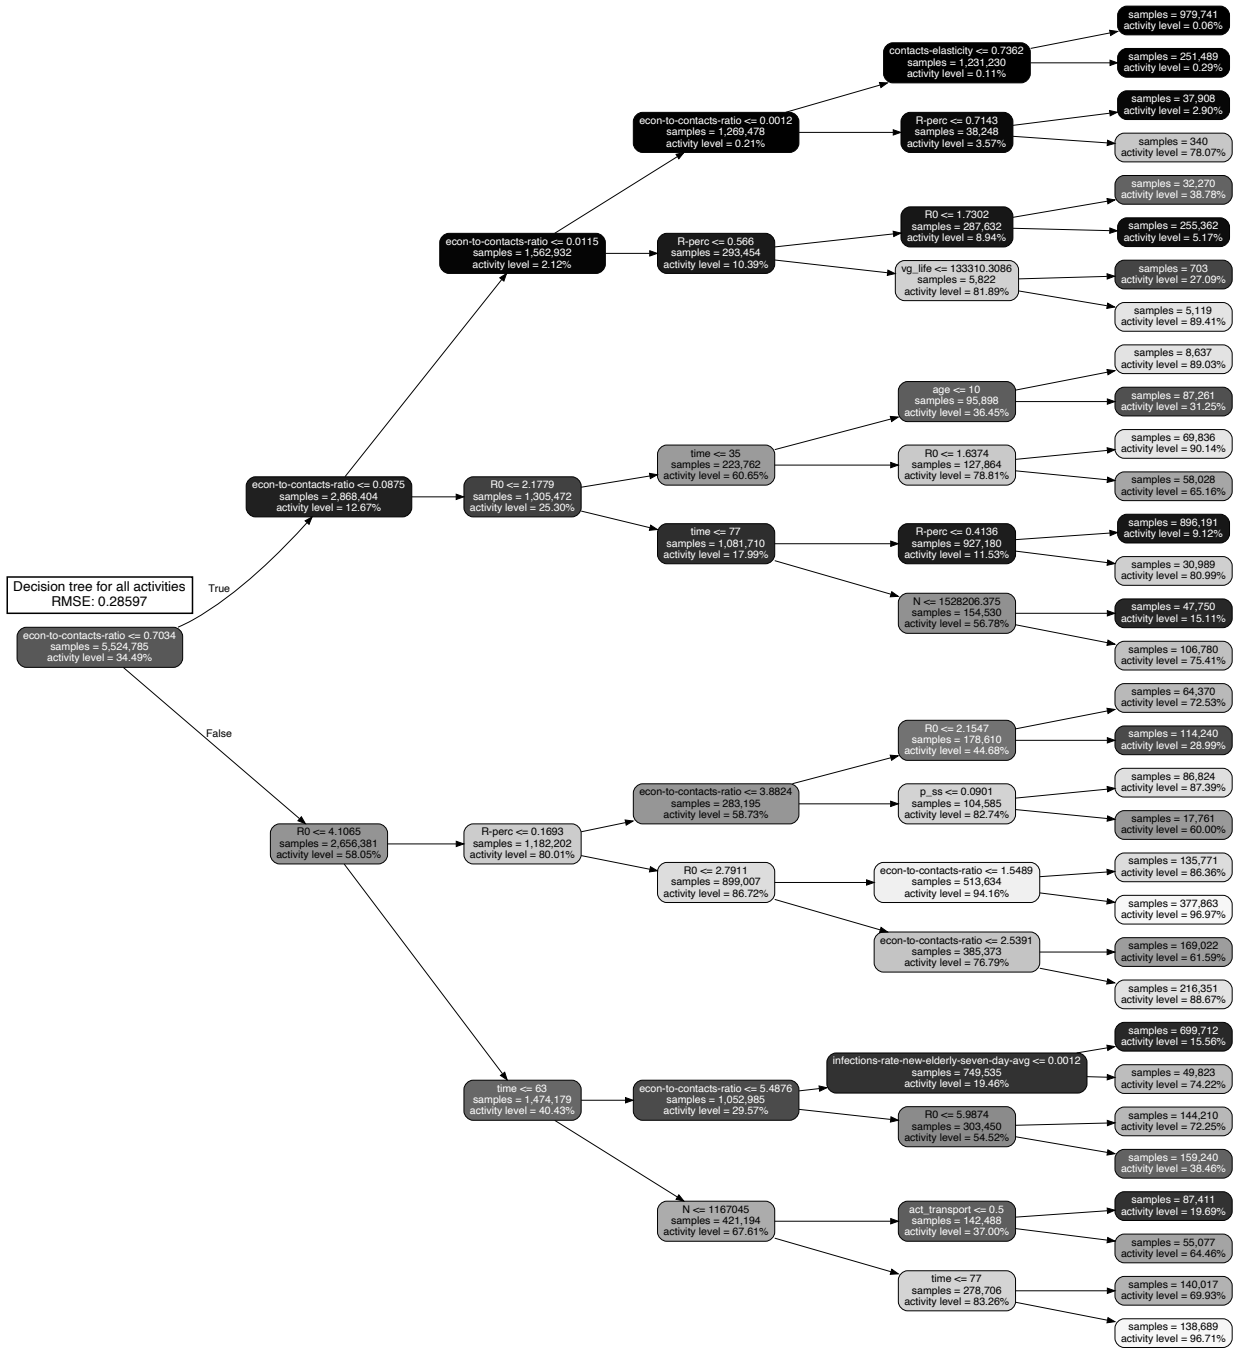

Figure EC.5 Decision tree of depth five approximating the optimized ROLD AGE-ACT confinement decisions (trained with 5,524,785 samples), with an optimization horizon of $T' = 90$ days.

a different means than permutation feature importance. Specifically, we train a tree with *only* the non-targeted features and trees with (all) the non-targeted features plus a single targeted feature. We do this to see which targeted feature most improves goodness-of-fit versus the non-targeted tree. Figure EC.7 shows the econ-to-contacts-ratio dominates the other trees, and is very close in RMSE to the tree using all features reported in Figure 10.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

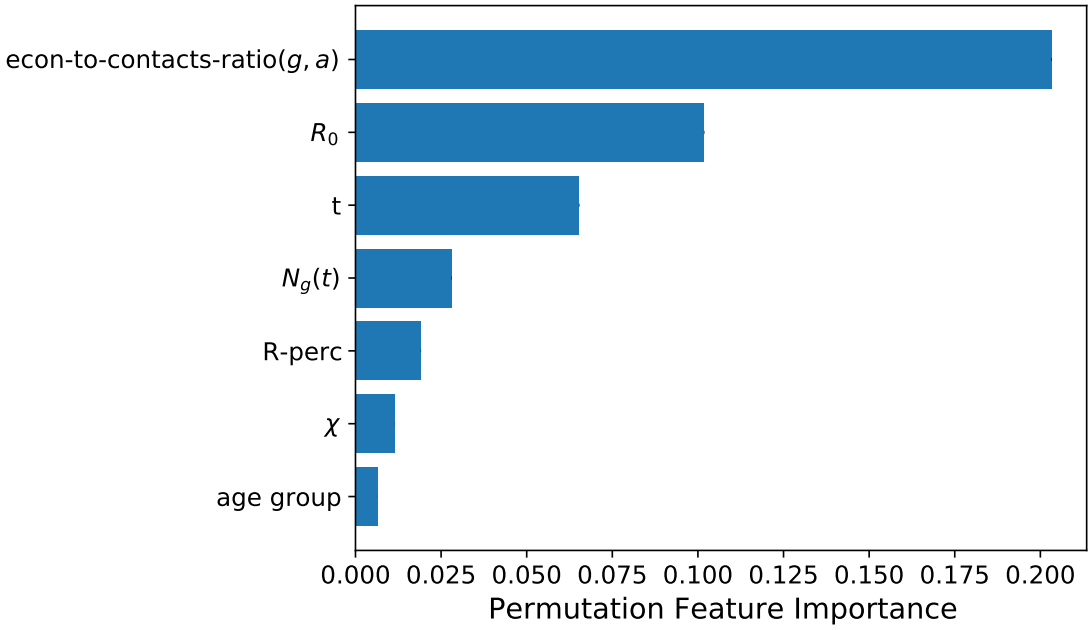

Figure EC.6 **Permutation feature importance for different variables, for the depth-10 decision tree for the ROLD AGE-ACT policy. We show all variables with importance score above 0.005.**

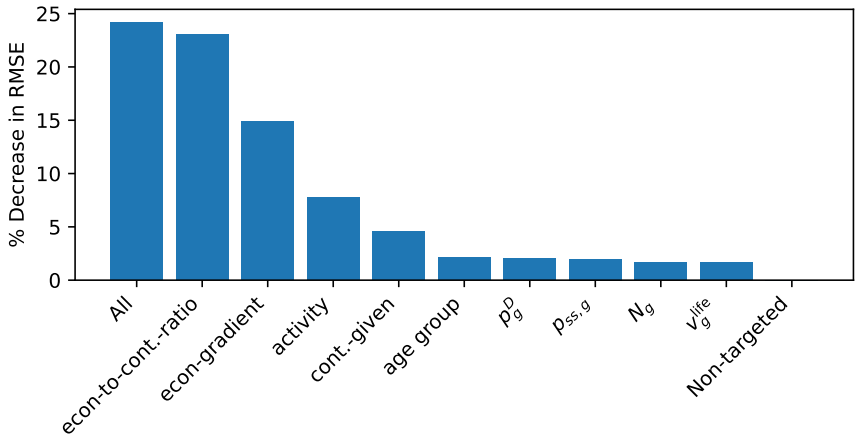

Figure EC.7 **Percent decrease in RMSE for the single targeted feature trees and the all features tree, versus the tree using only non-targeted features, for the ROLD AGE-ACT policy.**

EC.6.4.2. Additional Details for the Regression Decision Trees for the Gains of Dual Targeting The training set consists of 14,039 data samples (i.e., one sample for each problem instance), where we include several features, and targets corresponding to the relative gains of ROLD AGE-ACT over ROLD NO-TARGET and over ROLD ACT. The considered features are listed in Table EC.12. Note that the economic features are calculated as aggregated statistics (standard deviation) of the econ-to-contacts-ratio over group-activity pairs. When calculating these statistics, we only use relevant (age group, activity) pairs, where an activity is relevant for an age group if that age group engages in that activity, as detailed in Table EC.8.

Using this data as training set, we then employ the CART algorithm (implemented in the SCIKIT-LEARN Python library) to train regression decision trees to predict the relative gains of ROLD AGE-ACT over ROLD NO-TARGET, and ROLD AGE-ACT over ROLD ACT, as a function of

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.12 Features used for fitting trees for gains of dual targeting

Feature Name	Description
R_0	Basic reproduction number
contacts-elasticity	The percentage drop in total contacts when activity levels drop by 50%
χ	Cost of death
severity-scaling	Multiplicative scaling of probabilities $p_{ss,g}$ for all groups g
K^{ICU}	ICU capacity
econ-to-contacts-ratio-activity-stdev	The standard deviation of econ-to-contacts-ratio(g, a, t_0) across relevant age groups for an activity a. We include this feature for each activity $a \in \{\mathsf{work}, \mathsf{school}, \mathsf{transport}, \mathsf{leisure}, \mathsf{other}\}$.
econ-to-contacts-ratio-agegroup-stdev	The standard deviation of econ-to-contacts-ratio(g, a, t_0) across relevant activities for an age group g. We include this feature for each age group $g \in \mathcal{G}$.

the considered features, using the mean-squared-error (MSE) criterion as a goodness of fit metric. We use 100 as the minimum number of data points in order to split a node.

We report in Figures EC.8 and EC.9 the trees of depth five trained using the features described above, for the gains of AGE-ACT over NO-TARGET, and AGE-ACT over ACT, respectively.

EC.7. Stylized Model of Linearized Dynamics

This section generates further insight into the structure of the confinement policies produced by ROLD by examining a significantly simpler version of the problem we study, for which an analytical characterization is possible. Our goal is to understand how the optimal ROLD policy depends on various problem primitives (such as economic parameters, contact matrices, etc.), which in turn will confirm our selection of features for training the tree policies in Section 7 and the prominence of econ-to-contacts-ratio.

Simplified SEIR model. We consider a simplified compartmental model in which there is a single population group engaging in a single activity, and there are only susceptible (S), exposed (E), infectious (I), recovered (R) and deceased (D) compartments. Also for the sake of simplicity, we consider the *continuous* time dynamics of this model, namely:

$$\dot{S}(t) = -\beta C \ell(t) \frac{S(t)I(t)}{N} \tag{EC.56}$$

$$\dot{E}(t) = \beta C \ell(t) \frac{S(t)I(t)}{N} - \sigma E(t) \tag{EC.57}$$

$$\dot{I}(t) = \sigma E(t) - \mu I(t) \tag{EC.58}$$

$$\dot{R}(t) = (1-p)\mu I(t) \tag{EC.59}$$

$$\dot{D}(t) = p\mu I(t). \tag{EC.60}$$

We pause to explain the new dynamical equations. Since there is a single group and activity, the new infections term $\beta C \ell(t) S(t) I(t)$ is significantly simplified. Here, β is the transmission rate and C is the rate at which social contacts occur. Furthermore, σ and μ are the transition rates defined analogously to our original model. Since in this stylized model we remove hospitalized states, we have a direct transition from I to R or D; we denote by p the probability that an individual dies given that they are infectious. We denote by N the total population size.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

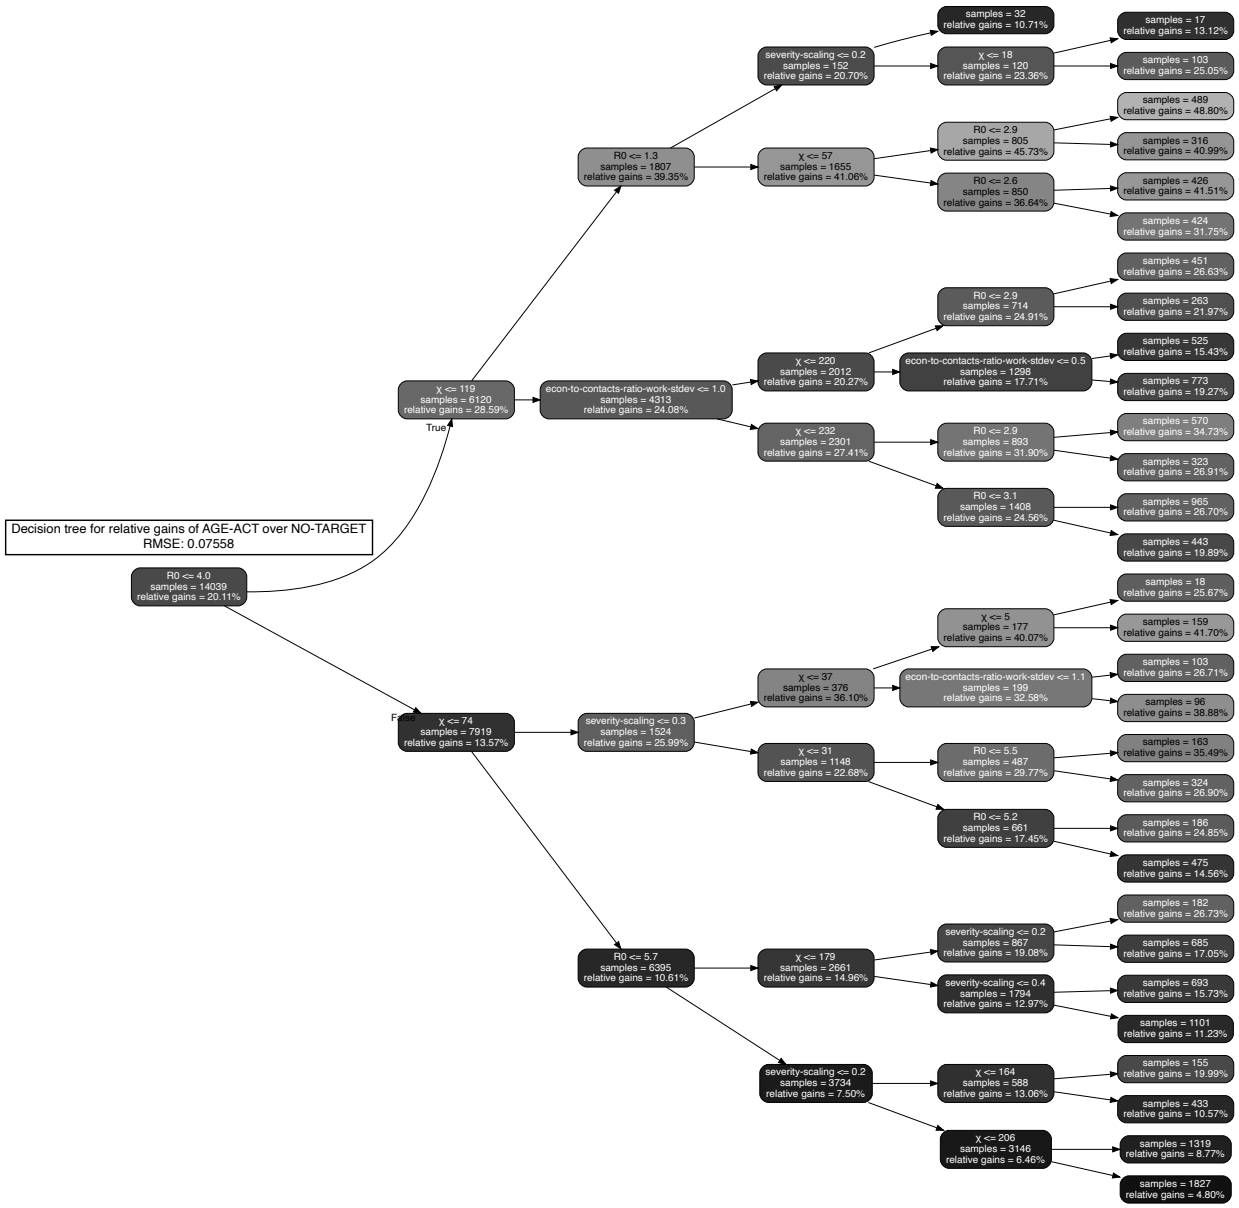

Figure EC.8 **Decision tree of depth five approximating the relative gains of the dual-targeted ROLD AGE-ACT over ROLD NO-TARGET for the total loss, trained on a total of 14,039 problem instances with an optimization horizon of $T' = 90$ days. The nodes are color-coded based on the relative gains, with darker colors corresponding to smaller gains.**

Activity levels. The control is the activity level of the population and is one-dimensional, so we denote it by ℓ, in analogy to the original model. We simplify the economic model by taking the economic value as a linear function of the activity level ℓ, i.e., $w\ell$, where w is the economic value generated per capita and per unit time under no confinement. Given the same cost of death

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

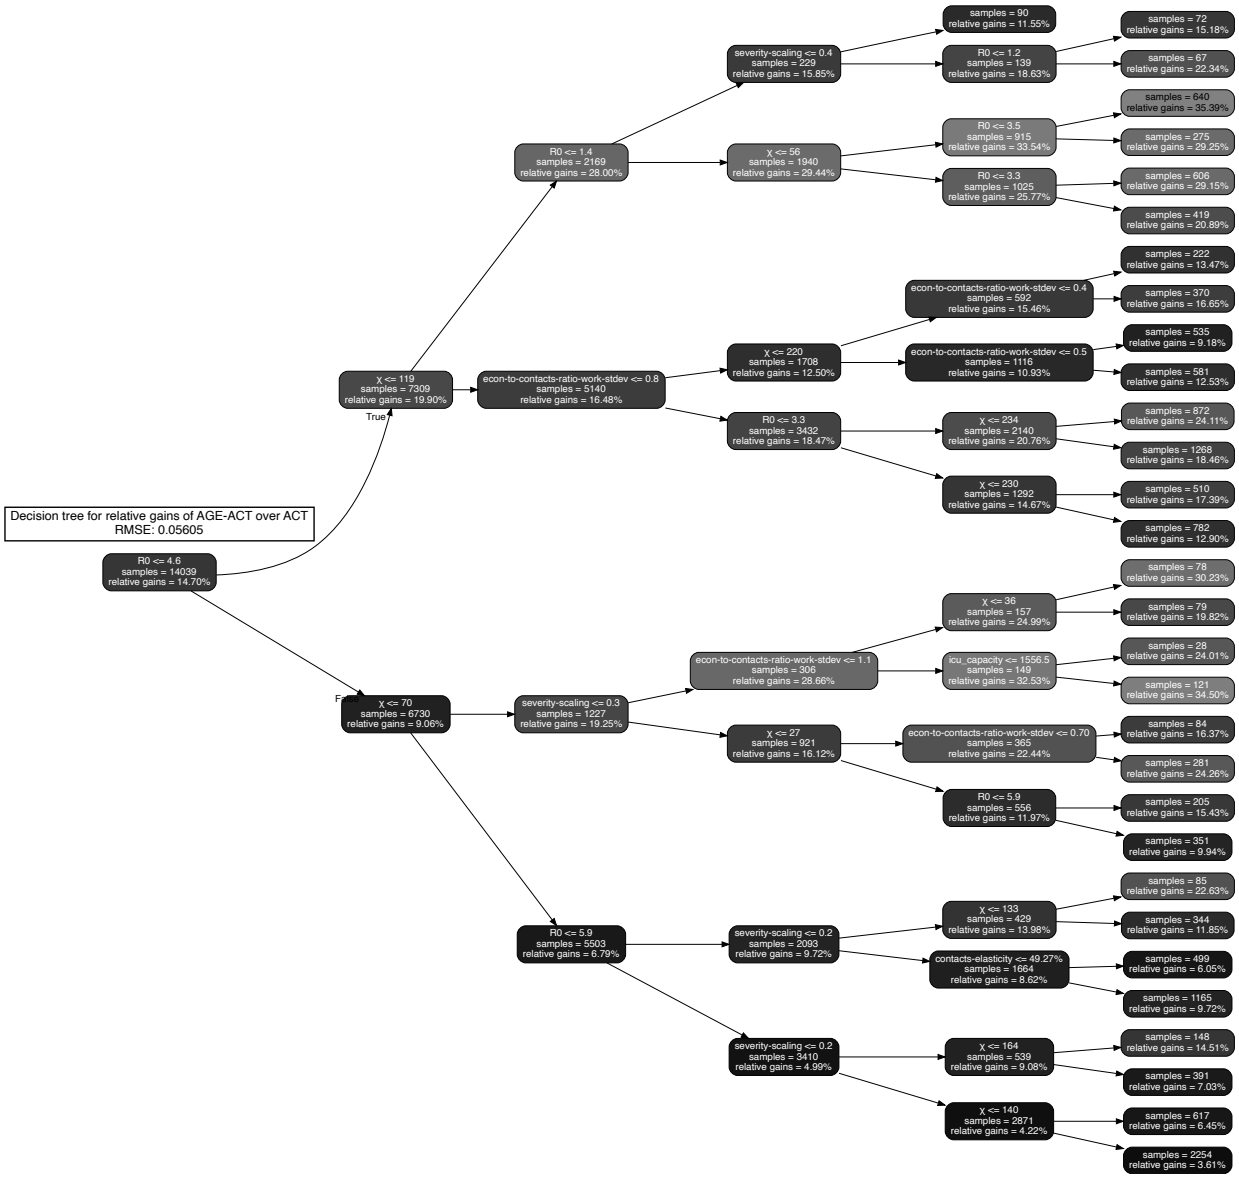

Figure EC.9 **Decision tree of depth five approximating the relative gains of the dual-targeted ROLD AGE-ACT over ROLD ACT for the total loss, trained on a total of 14,039 problem instances with an optimization horizon of $T' = 90$ days. The nodes are color-coded based on the relative gains, with darker colors corresponding to smaller gains.**

parameter χ, the objective is to maximize:[18]

$$\int_{t=0}^{T} w\ell(t)(S(t)+E(t)+I(t)+R(t))dt - \chi D(T). \tag{EC.61}$$

Solving the linearized system. To facilitate the analysis, we make the assumption that throughout the time horizon and for any activity level, $S(t) \approx N$. In that case, the new infections term can

[18] For simplicity, we fold the lost wages term v_g^{life} from our original model into χ.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

be approximated as

$$\beta C\ell(t)\frac{S(t)I(t)}{N} \approx \beta C\ell(t)I(t),$$

leading to the approximate dynamics:

$$\dot{S}(t) \approx -\beta C\ell(t)I(t) \tag{EC.62}$$
$$\dot{E}(t) \approx \beta C\ell(t)I(t) - \sigma E(t) \tag{EC.63}$$
$$\dot{I}(t) = \sigma E(t) - \mu I(t) \tag{EC.64}$$
$$\dot{R}(t) = (1-p)\mu I(t) \tag{EC.65}$$
$$\dot{D}(t) = p\mu I(t). \tag{EC.66}$$

In compact notation, these approximate SEIR dynamics can be written as $\dot{\boldsymbol{X}}(t) = \tilde{\boldsymbol{F}}(\boldsymbol{X}(t), \ell(t))$, where $\boldsymbol{X}(t) = (S(t), E(t), I(t), R(t), D(t))$ and the function $\tilde{\boldsymbol{F}}$ captures (EC.62)-(EC.66).

Additionally, under the assumption $S(t) \approx N$, the objective simplifies to:

$$\int_{t=0}^{T} w\ell(t)N dt - \chi D(T), \tag{EC.67}$$

since $S(t) + E(t) + I(t) + R(t) \approx N$.

With respect to this linearized dynamic and approximate objective, the control problem is:

$$\max_{\{\ell(t)\}} \int_{t=0}^{T} w\ell(t)N dt - \chi D(T) \tag{EC.68}$$
$$\text{s.t. } (\text{EC.62}) - (\text{EC.66}) \tag{EC.69}$$
$$\ell(t) \in [0,1], \quad \forall t. \tag{EC.70}$$

Note that the control problem under the simplified model does not include capacity constraints, as we have removed the hospitalized states from the model.

Although the dynamics in (EC.62) - (EC.66) are linear, the bilinear terms coming from multiplying ℓ with the SEIR states still make the problem difficult to solve. We thus proceed with a linearization of the problem which mimics the workings of the ROLD algorithm in Section 4, and characterize the optimal policy for this linearized problem. Analogously to Section 4, consider a nominal time-invariant control $\hat{\ell}$, and a nominal trajectory $\hat{\boldsymbol{X}}$. We build a linear approximation of the system dynamics as:

$$\dot{\boldsymbol{X}}(t) \approx \tilde{\boldsymbol{F}}(\hat{\boldsymbol{X}}, \hat{\ell}) + \nabla_{\boldsymbol{X}}\tilde{\boldsymbol{F}}(\hat{\boldsymbol{X}}, \hat{\ell})(\boldsymbol{X}(t) - \hat{\boldsymbol{X}}) + \nabla_{\ell}\tilde{\boldsymbol{F}}(\hat{\boldsymbol{X}}, \hat{\ell})(\ell(t) - \hat{\ell}) \tag{EC.71}$$
$$= \boldsymbol{A}\boldsymbol{X}(t) + \boldsymbol{B}\ell(t) + \boldsymbol{c}, \tag{EC.72}$$

where

$$\boldsymbol{A} = \begin{bmatrix} 0 & 0 & -\beta C\hat{\ell} & 0 & 0 \\ 0 & -\sigma & \beta C\hat{\ell} & 0 & 0 \\ 0 & \sigma & -\mu & 0 & 0 \\ 0 & 0 & (1-p)\mu & 0 & 0 \\ 0 & 0 & p\mu & 0 & 0 \end{bmatrix}, \quad \boldsymbol{B} = \begin{bmatrix} -\beta C\hat{I} \\ \beta C\hat{I} \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \boldsymbol{c} = \tilde{F}(\hat{\boldsymbol{X}}, \hat{\ell}) - \nabla_{\boldsymbol{X}}\tilde{F}(\hat{\boldsymbol{X}}, \hat{\ell})\hat{\boldsymbol{X}} - \nabla_{\ell}\tilde{F}(\hat{\boldsymbol{X}}, \hat{\ell})\hat{\ell}. \tag{EC.73}$$

This allows us to express the state at any t as a function of the confinement decision $\ell(\tau)$ for all $0 \le \tau \le t$. In particular, the solution of the dynamical system $\dot{\boldsymbol{X}}(t) = \boldsymbol{A}\boldsymbol{X}_t + \boldsymbol{B}\ell(t) + \boldsymbol{c}$ is

$$\boldsymbol{X}(t) = e^{\boldsymbol{A}t}\boldsymbol{X}(0) + \int_0^t e^{\boldsymbol{A}(t-\tau)}\left(\boldsymbol{B}\ell(\tau) + \boldsymbol{c}\right) d\tau. \tag{EC.74}$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Note that we can readily write the objective (EC.68) as maximizing

$$\int_0^T wN\ell(t)dt - \chi D(t) = \int_0^T \left(wN\ell(t) - \chi p\mu I(t)\right) dt, \tag{EC.75}$$

without resorting to any Taylor approximations for linearization.

By plugging in the solution for the dynamical system (EC.74), we have

$$\int_0^T I(t)dt = \int_0^T \left(\boldsymbol{e}_3^\top e^{\boldsymbol{A}t}\boldsymbol{X}(0) + \int_0^t \boldsymbol{e}_3^\top e^{\boldsymbol{A}(t-\tau)}\left(\boldsymbol{B}\ell(\tau)+\boldsymbol{c}\right)d\tau\right)dt \tag{EC.76}$$

$$= \int_0^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}t}\boldsymbol{X}(0)dt + \int_{t=0}^T\int_{\tau=0}^t \boldsymbol{e}_3^\top e^{\boldsymbol{A}(t-\tau)}\left(\boldsymbol{B}\ell(\tau)+\boldsymbol{c}\right)d\tau dt \tag{EC.77}$$

$$= \int_0^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}t}\boldsymbol{X}(0)dt + \int_{\tau=0}^T\left(\int_{t=\tau}^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}(t-\tau)}dt\right)\left(\boldsymbol{B}\ell(\tau)+\boldsymbol{c}\right)d\tau \tag{EC.78}$$

$$= \int_0^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}t}\boldsymbol{X}(0)dt + \int_{t=0}^T\left(\int_{\tau=t}^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}(\tau-t)}d\tau\right)\left(\boldsymbol{B}\ell(t)+\boldsymbol{c}\right)dt, \tag{EC.79}$$

where $\boldsymbol{e}_3 := [0,0,1,0,0]$.

We can now rewrite the objective in (EC.75) as

$$\int_0^T \left(wN - \chi p\mu\int_t^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}(\tau-t)}d\tau\cdot\boldsymbol{B}\right)\ell(t)dt + \int_0^T\left(\boldsymbol{e}_3^\top e^{\boldsymbol{A}t}\boldsymbol{X}(0) - \chi p\mu\int_t^T \boldsymbol{e}_3^\top e^{\boldsymbol{A}(\tau-t)}d\tau\cdot\boldsymbol{c}\right)dt. \tag{EC.80}$$

The second summand is a constant with respect to the control. The coefficient of $\ell(t)$ in the integral in the first summand calculates to

$$wN - \chi p\mu\cdot\frac{2\beta C\hat{I}\sigma\left(\frac{\exp\left(\frac{1}{2}(t-T)\left(\sqrt{4\beta C\hat{\ell}\sigma+(\mu-\sigma)^2}+\mu+\sigma\right)\right)-1}{\sqrt{4\beta C\hat{\ell}\sigma+(\mu-\sigma)^2}+\mu+\sigma} - \frac{\exp\left(-\frac{1}{2}(t-T)\left(\sqrt{4\beta C\hat{\ell}\sigma+(\mu-\sigma)^2}-\mu-\sigma\right)\right)-1}{-\sqrt{4\beta C\hat{\ell}\sigma+(\mu-\sigma)^2}+\mu+\sigma}\right)}{\sqrt{4\beta C\hat{\ell}\sigma+(\mu-\sigma)^2}}, \tag{EC.81}$$

and the optimal policy is to set ℓ to 1 if and only if the expression above is non-negative.

It is perhaps most useful to focus on understanding the resulting policy for the case that $\hat{\ell}=0$, in which case the coefficient of $\ell(t)$ takes a simpler form equal to

$$wN - \chi p\mu\cdot\frac{2\beta C\hat{I}\sigma\left(\frac{\exp\left(\frac{1}{2}(t-T)(|\mu-\sigma|+\mu+\sigma)\right)-1}{|\mu-\sigma|+\mu+\sigma} - \frac{\exp\left(-\frac{1}{2}(t-T)(|\mu-\sigma|-\mu-\sigma)\right)-1}{-|\mu-\sigma|+\mu+\sigma}\right)}{|\mu-\sigma|}. \tag{EC.82}$$

The optimal decision then is

$$\ell^*(t) = \begin{cases} 1 & \text{if } \frac{w}{C} \geq f(\beta,\chi,p,\sigma,\mu,N,\hat{I},t,T) \\ 0 & \text{otherwise,} \end{cases} \tag{EC.83}$$

where

$$f(\beta,\chi,p,\sigma,\mu,N,\hat{I},t,T) := \frac{2\beta\chi p\sigma\mu\hat{I}}{N|\mu-\sigma|}\cdot\left(\frac{\exp\left(\frac{1}{2}(t-T)\left(|\mu-\sigma|+\mu+\sigma\right)\right)-1}{|\mu-\sigma|+\mu+\sigma} - \frac{\exp\left(-\frac{1}{2}(t-T)\left(|\mu-\sigma|-\mu-\sigma\right)\right)-1}{-|\mu-\sigma|+\mu+\sigma}\right). \tag{EC.84}$$

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

(We note that[19] $f(\beta,\chi,p,\sigma,\mu,N,\hat{I},t,T)\geq 0$ for $0\leq t\leq T$.)

Inequality (EC.83) uncovers a natural logic for the confinement decisions in this linearized formulation. Specifically, the left-hand side exactly corresponds to the econ-to-contacts-ratio that we identified in Section 7: it is given by the gradient of (per-capita) economic value generated with respect to the level of activity $\ell(t)$[20] divided by the rate of social contacts generated. Thus, the optimal policy is governed by a threshold on the value of the econ-to-contacts-ratio, allowing normal activity levels ($\ell(t)=1$) when the econ-to-contacts-ratio exceeds the threshold and completely confining the entire population ($\ell(t)=0$) otherwise. The threshold, given by the function $f(\beta,\chi,p,\sigma,\mu,N,\hat{I},t,T)$, is increasing in parameters such as the probability of infection given a contact β, the probability of death given infection p, and the cost of death χ, and is decreasing in the size of the overall population N, which matches intuition.

We remark that the simple threshold policy relying on econ-to-contacts-ratio emerges when we impose $\hat{\ell}=0$ in (EC.81); for general $\hat{\ell}$, the policy still follows a threshold rule based on the marginal economic value exceeding a function that depends on problem parameters and the rate of contacts, but the simple ratio no longer emerges in the rule.

There are two conclusions we highlight from this stylized model analysis. First, even with significant simplifications such as a single group and activity, the policies output by ROLD or similar procedures are quite complex and elusive to completely characterize in closed form. Second, there are however interesting regimes where these policies intuitively depend on a quantity resembling the econ-to-contacts-ratio. In this light, selecting the group-activity pairs to be confined in decreasing order of their econ-to-contacts-ratios, as our trees from Section 7 do, naturally parallels the policy in (EC.83). This provides some theoretical backing to the feature that the trees consider to be important in explaining ROLD's decisions.

EC.8. Details on the Practically Implementable Policies

EC.8.1. Retrieving Participation Rates from the Time-Use Survey Data

The time-use survey data by INSEE (2010) follows a data format that is harmonized with EU standards for time-use surveys. Each row in the main data file is a "diary" that corresponds to a specific day of a specific survey participant. Among other information, a diary has information on the self-reported location of the participant, as well as the self-reported main activity of the participant, per 10-minute interval (for a total of 144 intervals over the 24-hour day). The survey participant selects both the reported location and the reported main activity from a given menu of options.

Most of the reported locations in the diaries map naturally into one of our six activities in the set $\mathcal{A}=\{\mathsf{work},\mathsf{transport},\mathsf{leisure},\mathsf{school},\mathsf{home},\mathsf{other}\}$. However, for four of the reported locations in the diaries ("Unspecified location/transport mode", "Unspecified location (not travelling)", "Workplace or school", "Other specified location (not travelling)"), we use supplemental information on the main activity pursued, as well as on the activity (i.e., employment) status of the individual, to decide on the exact mapping. In particular:

[19] To see this, fix all the problem parameters and define

$$g(t):=\frac{\exp\left(\frac{1}{2}(t-T)\left(|\mu-\sigma|+\mu+\sigma\right)\right)-1}{|\mu-\sigma|+\mu+\sigma}-\frac{\exp\left(-\frac{1}{2}(t-T)\left(|\mu-\sigma|-\mu-\sigma\right)\right)-1}{-|\mu-\sigma|+\mu+\sigma}. \tag{EC.85}$$

Then $g(T)=0$, while

$$g'(t)=\frac{1}{2}\exp\left(\frac{1}{2}(t-T)\left(|\mu-\sigma|+\mu+\sigma\right)\right)-\frac{1}{2}\exp\left(\frac{1}{2}(t-T)\left(-|\mu-\sigma|+\mu+\sigma\right)\right)\leq 0, \tag{EC.86}$$

because $e^{x(a+b)}-e^{x(-a+b)}\leq 0$ for $a\geq 0, x\leq 0$. Therefore g is non-increasing in t in $[0,T]$, and $g(t)\geq 0$ for $0\leq t\leq T$.

[20] Since we use the simplification $S(t)+E(t)+I(t)+R(t)\approx N$, it stands to reason that a per capita value is sufficient here.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

- For reported locations "Unspecified location/transport mode", "Unspecified location (not travelling)", "Other specified location (not travelling)", we map the location to one or more of our activities in set $\mathcal{A}$ on the basis of the reported time-use activity, using our judgment on which of our activities are highly relevant for the respective pair of reported location-activity in the survey.
- For reported location "Workplace or school":

—If the activity (employment) status of the survey participant is "Pupil, student, further training, unpaid traineeship", then for most reported activities we assign to our activity school. For a few reported activities that relate to transport, we assign to both school and transport.

—If the activity (employment) status of the survey participant is "Employed full-time" or "Employed part-time", then for most reported activities we assign to our activity work. For a few reported activities, we assign either to school; work and school; work and transport; or school and transport.

—Otherwise (i.e., if the activity status of the survey participant is neither of the above), then for most reported activities we assign to both our activities work and school. For a few reported activities, we assign either to only work; only school; work and transport; or school and transport.

We detail how we calculate the participation rates $\mathcal{R}_g^a(\tau)$ in Algorithm 7. When the proposed mapping assigns to more than one of our activities, then we split the 10 minutes of the corresponding 10-minute interval equally across the proposed activities. We note that the time-use survey data does not include responses from participants aged below 10 y.o., and therefore for our 0-9 y.o. age group we propose an interpolation from survey participants aged 10, 11, and 12 y.o.

Algorithm 7 Calculating Participation Rates $\mathcal{R}_g^a(\tau)$

for age group $g \in \{$0-9, ..., 80+ y.o.$\}$ **do**
  for activity $a \in \mathcal{A}$ **do**
    for hourly slot $\tau = 1, 2, \ldots, 24$ **do**
      if $g \neq$ 0-9 y.o. **then**

$$\text{ratio_weekdays} := \frac{\sum_{\text{diaries from } g \text{ on weekdays}} \sum_{\text{10-min intervals of } \tau} \text{mins spent on act. } a \text{ during that interval in that diary}}{(\#\text{ diaries from } g \text{ on weekdays}) \cdot 60}$$

$$\text{ratio_weekend} := \frac{\sum_{\text{diaries from } g \text{ on weekend}} \sum_{\text{10-min intervals of } \tau} \text{mins spent on act. } a \text{ during that interval in that diary}}{(\#\text{ diaries from } g \text{ on weekend}) \cdot 60}$$

      else

$$\text{ratio_weekdays} := \frac{\sum_{\text{diaries from 10-12 y.o. on weekdays}} \sum_{\text{10-min intervals of } \tau} \text{mins spent on act. } a \text{ during that interval in that diary}}{(\#\text{ diaries from 10-12 y.o. on weekdays}) \cdot 60}$$

$$\text{ratio_weekend} := \frac{\sum_{\text{diaries from 10-12 y.o. on weekend}} \sum_{\text{10-min intervals of } \tau} \text{mins spent on act. } a \text{ during that interval in that diary}}{(\#\text{ diaries from 10-12 y.o. on weekend}) \cdot 60}$$

      end if

$$\mathcal{R}_g^a(\tau) := \frac{5}{7} \cdot \text{ratio_weekdays} + \frac{2}{7} \cdot \text{ratio_weekend}$$

    end for
  end for
end for

EC.8.2. Details on the Gradient Descent Procedure for Optimizing Curfews

We used the 'trust-constr' method in the SCIPY-OPTIMIZE Python library, and employed its default parameter settings. For a given problem instance, we initialized the optimization algorithm using two different strategies:

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

• Random initialization: For each activity, we sample the start and duration of the curfew uniformly at random, both from the $0-24$ hours interval. We repeat this exercise for a total of 120 iterations.

• ROLD ACT solution initialization, randomly perturbed: For each activity, we draw an i.i.d. Bernoulli variable with probability of success p, where p is a parameter explained below. In case of success, the curfew is initialized at the ROLD ACT solution. In case of failure, we sample the start and duration of the curfew uniformly at random, both from the $0-24$ hours interval. We vary the parameter p in the range $[0, 1/6, 2/6, 3/6, 4/6, 5/6]$ and repeat this exercise for 20 iterations for each of the values for parameter p. This gives a total of 120 iterations.

Each of the iterations is run for 24 hours using two strategies:

• Full completion: The algorithm runs until local optimality, which, on average, is achieved after eight hours.

• Partial completion: The algorithm runs for two hours and is restarted with a new random seed for a total of 12 times.

In total, we used 600 CPUs, each running for 24 hours, where the processors used in the cluster were Intel E5-2640v4, Intel 5118, and AMD 7502. We note that we did not optimize the code for more efficient computation.

EC.8.3. Time-Average Aggregate Activity Levels for Age Groups

We calculate the time-average activity level of each age group, averaged over the activities relevant to that age group (Table EC.8), and weighting different activities by the respective participation rates. Specifically, for AGE-ACT, ACT, and NO-TARGET, the proposed calculation is the one in (EC.55) (and similarly for the other age groups). For CURFEW, the proposed calculation just replaces $\ell_g^a(t)$ in these equations with $\ell_g^{a,\text{eff}}(t)$ from (15).

EC.8.4. Details on Limited Compliance Policies with Age-Group-Specific Compliance Fractions

We use the estimates for age-group-specific non-compliance probabilities by Ganslmeier et al. (2022) to estimate a relative scaling of compliance probabilities across age groups. Ganslmeier et al. (2022) estimate non-compliance probabilities based on data from the first COVID-19 lockdown in the UK. Because the age groups used by Ganslmeier et al. (2022) are different from ours, to find the compliance probability p_g for an age group $g \in \mathcal{G}$ we take the conservative approach of using the minimum over the compliance probabilities of the age groups of Ganslmeier et al. (2022) that intersect with age group g. For our 0-9 y.o. age group, for which Ganslmeier et al. (2022) do not estimate compliance, we use the compliance probability from Ganslmeier et al. (2022) for their youngest group, 18-24 y.o. Table EC.13 summarizes the age-group specific compliance probabilities we interpolate for our age groups. Based on these, we estimate a relative scaling between compliance probabilities of different age groups, assuming that the relative scaling in France and in particular Île-de-France would be similar to the relative scaling in the UK.

Restricting to that fixed relative scaling, we refer to the resulting effective activity levels as the ROLD COMP-NH(p) policy, where NH stands for non-homogeneous across age groups, and where p is the compliance fraction for the least compliant age group according to the scaling. We look at the performance of the ROLD COMP-NH(p) policy as the compliance fraction p varies. The results are similar to the results for limited compliance policies with a population-wide compliance fraction (ROLD COMP(p)), and are summarized in Figure EC.10.

E-companion References

Banque de France (2020a) 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`, accessed June 19, 2020.

Banque de France (2020b) 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`, accessed June 19, 2020.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Table EC.13 Fraction of the population that complies with recommended restrictions by age group as interpolated from study by Ganslmeier et al. (2022)

Age Group g (y.o.)	Compliance Fraction p_g (%)
0-9	96.10
10-19	96.10
20-29	94.66
30-39	94.66
40-49	95.50
50-59	96.27
60-69	96.40
70-79	96.45
80+	96.45

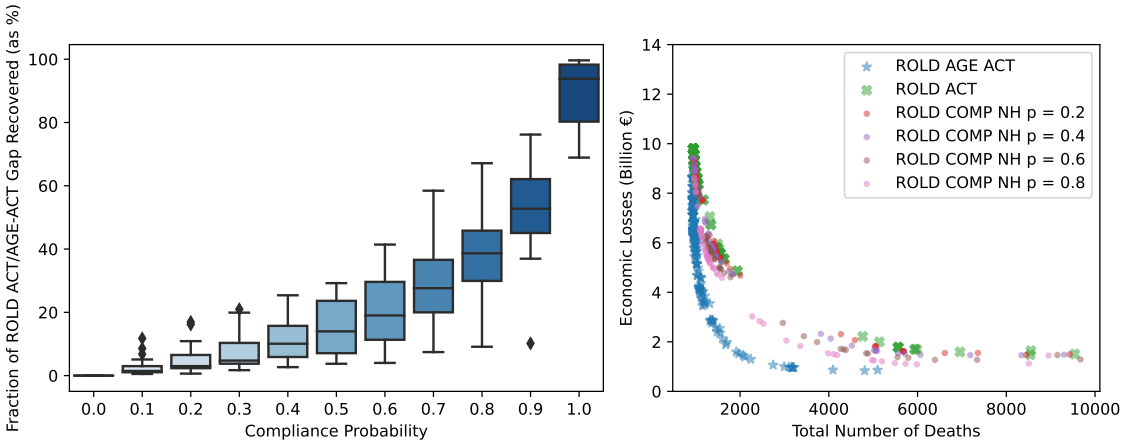

Figure EC.10 The performance of limited compliance policies with age-group-specific compliance, for different compliance levels. We estimate a relative scaling of compliance fractions across age groups using the estimates in Ganslmeier et al. (2022). Keeping the relative scaling across age groups fixed, and denoting the compliance fraction for the least compliant age group by p, we look at the performance of ROLD COMP-NH(p) as p varies. In the left panel, the boxplot at each p summarizes the fraction of the ROLD ACT/AGE-ACT performance gap that ROLD COMP-NH(p) recovers, across different values of the cost of death χ. We include 201 distinct values of χ from 0 to 1000$\times$. The right panel illustrates the total number of deaths and the economic losses generated by ROLD COMP-NH(p) for different compliance levels p, compared to ROLD AGE-ACT and ROLD ACT. Each marker corresponds to a different problem instance parameterized by χ.

Bi Q, et al. (2020) Epidemiology and transmission of COVID-19 in 391 cases and 1286 of their close contacts in Shenzhen, China: a retrospective cohort study. *The Lancet Infectious Diseases* 20(8).

Diekmann O, Heesterbeek JAP, Roberts MG (2010) The construction of next-generation matrices for compartmental epidemic models. *Journal of The Royal Society Interface* 7(47):873–885.

Du Z, Xu X, Wu Y, Wang L, Cowling BJ, Meyers L (2020) Serial interval of COVID-19 among publicly reported confirmed cases. *Emerging Infectious Diseases* 26(6).

Duque D, Morton DP, Singh B, Du Z, Pasco R, Meyers LA (2020) Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proceedings of the National Academy of Sciences* 117(33):19873–19878.

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

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

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

INSEE (2016a) T402 : Salaire brut en équivalent temps plein, par secteur d'activité, région et département. `https://www.insee.fr/fr/statistiques/fichier/4204500/T402.xls`, accessed June 19, 2020.

INSEE (2016b) 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`, accessed June 19, 2020.

INSEE (2019) Au quatrieme trimestre 2019, le taux de chômage passe de 8,5 % a 8,1 %. `https://www.insee.fr/fr/statistiques/4309346`, accessed June 19, 2020.

INSEE (2020) 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`, accessed June 19, 2020.

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

Mohammad H (2020) Coronavirus : un habitant de Bobigny considéré comme le nouveau "patient zéro". France Bleu, URL `https://www.francebleu.fr/infos/sante-sciences/coronavirus-un-habitant-de-bobigny-considere-comme-le-nouveau-patient-zero-1588796619`, accessed January 9, 2021.

Perasso A (2018) An introduction to the basic reproduction number in mathematical epidemiology. *ESAIM: Proceedings and Surveys* 62:123–138.

Salje H, et al. (2020) Estimating the burden of SARS-CoV-2 in France. *Science* 38.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC 4.0 International license.

Table EC.14: Examples of enforced targeted restrictions

Country	Region	Type of Targeting	Involved Activities	Involved Age Groups	Summary of Restriction	Dates of Restriction	Source
Azerbaijan	All of Azerbaijan	AGE	ALL	65+ y.o.	All residents not allowed to leave their homes unless given permission through SMS. Persons aged 65 and above completely banned from leaving their homes.	Lockdown in effect March 24, 2020. SMS permissions in effect April 5, 2020.	Jabbarova (2020)
Bosnia	All of Bosnia	AGE	ALL	Below 18 y.o.; 65+ y.o.	All children below 18 and adults above 65 not allowed to leave home. Restrictions eased on April 24. Adults above 65 allowed to leave home between 9 a.m. to 1 p.m. on Monday, Wednesday and Friday. Children younger than 18 allowed out on Tuesday, Thursday and Saturday from 2 p.m. to 8 p.m.	Ease of restrictions in effect April 24, 2020.	Staff; Sadikovic and Synovitz (2020)
Colombia	All of Colombia	AGE	ALL	70+ y.o.	All persons older than 70 required to stay at home. Exceptions include purchasing groceries and medicines, using health services, and accessing financial services.	In effect March 20 to May 31, 2020.	Reuters Staff (2020a)
Colombia	All of Colombia	AGE	ALL	6–17 y.o.; 18–69 y.o.; 70+ y.o.	Mandatory nationwide confinement imposed. Measures gradually eased for different age groups and at different times. Measures eased for children aged 6–17. Children aged 6–13 allowed to leave their homes for half an hour three times a week, between 8 a.m. to 11 a.m. Teenagers aged 14–17 allowed out for half an hour between 2 p.m. to 5 p.m. Measures eased for adults. Persons aged 18–69 allowed to leave the house for 2 hours everyday. Persons aged 70 and above allowed to leave the house for 1 hour three times a week. Stricter measures for persons aged 70 and above seen as discriminatory and challenged in Colombian court. Courts ruled in favor of persons aged 70 and above. Elderly permitted to leave their homes for the same amounts of time as persons aged 18-69	Mandatory confinement with effect March 25, 2020. Restrictions for children aged 6-17 eased with effect May 11, 2020. Restrictions for persons aged 18-69 eased with effect June 1, 2020. Restrictions for persons aged 70 and above eased with effect July 3, 2020.	Pertuz-Cruz et al. (2021); Portafolio (2020); President of the Republic of Colombia (2020)
Denmark	All of Denmark	ACT; AGE-ACT	Schools	Children above 12 y.o.	After schools had been closed for one month, restrictions were eased by reopening schools for children aged 2-12.	In effect April 15, 2020	Reuters Staff (2020c)
Finland	All of Finland	AGE	ALL	70+ y.o.	General guideline: Elderly over the age of 70 to refrain from contact with others (quarantine-like conditions). Members of parliament, state leadership, and elected officials in local government exempted.	In effect March 16, 2020.	Finnish Government (2020)
Georgia	All of Georgia	AGE	ALL	70+ y.o.	Nationwide curfew imposed from 9 p.m. to 6 a.m. Persons aged 70 and above banned from leaving their houses unless necessary.	In effect March 31, 2020 to April 21, 2020	Civil.ge (2020)

AGE (ACT) refers to restrictions that target age groups (activities). AGE-ACT refers to restrictions that target both age groups and activities.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC 4.0 International license.

Examples of enforced targeted restrictions

Country	Region	Type of Targeting	Involved Activities	Involved Age Groups	Summary of Restriction	Dates of Restriction	Source
Hungary	All of Hungary	AGE-ACT	Shopping	65+ y.o.	Retail shopping reserved from 9 a.m. to 11 a.m. on weekdays and 8 a.m. to 10 a.m. on weekends for elderly who are 65 y.o. and above.	In effect November 23, 2020 to December 14, 2020.	A3M Global Monitoring (2020); Reuters Staff (2020b)
Hungary	All of Hungary	AGE-ACT	Schools	Older children (grade 5 and above)	Kindergarten and grades 1 to 4 in primary school to reopen. Classes for grades 5 to 6 as well as high schools to remain online.	In effect April 19, 2021	Xinhua (2020)
India	All of India	ACT	Schools	ALL	Closure of schools.	In effect March 2020.	Mogul and Sud (2022)
Israel	All of Israel	AGE-ACT	Workplaces	67+ y.o.	Quarantine restrictions rolled back on April 19. 30 percent of workers allowed back into the office, but persons aged 67 and above requested to stay at home.	Differentiated restrictions in effect April 19, 2020. Ban on the elderly returning to work ended May 2, 2020.	CTECH CTech (2020); Times of Israel Magid (2020)
Jordan	All of Jordan	AGE	ALL	Below 16 y.o.; 60+ y.o.	Children under 16 and elderly above 60 not allowed to leave their homes except for essential needs.	In effect from March 17, 2020. Eased restrictions began on March 21, 2020. Government affirmed on April 27, 2020 the restrictions on those below 16 and those above 60.	Picheta and Qiblawi (2020); Human Rights Watch (2020); Saraya News (2020)
Kenya	All of Kenya	ACT	Schools	ALL	Schools closed for six months from March 2020. Schools reopened in October 2020 only for the grades preparing for standardised tests which had been postponed from November 2020 to March 2021 (grades 4, 8, 12).	Schools in Kenya closed on March 15, 2020. Partial reopening began on October 12, 2020.	Parsitau and Jepkemei (2020); Reuters (2020)
Moldova	All of Moldova	AGE	ALL	63+ y.o.	Restriction of elderly persons aged 63 and above from leaving their homes unless for specified reasons (eg. buying food, medical reasons, etc.).	In effect March 25, 2020	CSOmeter (2020); Timpul (2020)
New Zealand	Auckland and Waikato	AGE-ACT	Schools	Children (school years 1-10)	Staggered school re-openings in Auckland and Waikato after 10 week closure. Children in years 11-13 to return to school first, followed by children in years 1-10.	Years 11-13 returned to school on October 26, 2021. Years 1-10 returned to school on November 17, 2021.	CGTN (2021); 1News (2021)
Romania	All of Romania	AGE	ALL	65+ y.o.	Citizens allowed to leave the house only for work, medical attention, or to purchase food. Elderly aged 65 and above only allowed to leave the house between 11 a.m. to 1 p.m.	In effect March 25, 2020	Crisis 24 (2020)
Russia	Moscow	AGE	ALL	65+ y.o.; 60+ y.o.	Residents in Moscow over 65 y.o. and those with underlying health concerns requested to remain at home. Update: 60+ y.o. and above.	In effect September 28, 2020. Update on 60+ y.o.: in effect October 25, 2021 to February 25, 2022	Foy (2020)
Serbia	All of Serbia	AGE	ALL	65+ y.o.; 70+ y.o.	Nationwide curfew from 8 p.m. to 5 a.m. Nationwide ban on the elderly from leaving home: elderly above 65 in urban areas and above 70 in rural areas forbidden to go out (with exceptions).	In effect March 18, 2020.	Andrić (2020); Vasovic (2020)
Serbia	All of Serbia	AGE-ACT	Shopping	65+ y.o.	Persons above the age of 65 allowed to leave home to buy groceries from 3 a.m. to 8 a.m. only on specified day of the week. Assigned day shifted from Sunday to Saturday, then to Friday.	In effect March 18, 2020. Shift to Saturday in effect April 3, 2020	Stevanović (2020) ; The Government of the Republic of Serbia (2020) ; Reuter's Staff (2020)

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC 4.0 International license .

Examples of enforced targeted restrictions

Country	Region	Type of Targeting	Involved Activities	Involved Age Groups	Summary of Restriction	Dates of Restriction	Source
Singapore	All of Singapore	AGE-ACT	Senior Activity Centres	50+ y.o.	Suspension of all senior-centric activities organised by Government agencies (e.g., senior activities centres and residents' committees).	In effect March 11, 2020.	The Straits Times Zhang and Han (2020)
Singapore	All of Singapore	AGE-ACT	Social/family visiting	Senior citizens and vulnerable populations	Lockdown had been imposed since April 7, 2020. From June 2, 2020, visits to parents and grandparents allowed. Senior citizens should not leave their houses to visit anyone, but their children are allowed to visit them. Exceptions permitted for siblings, nieces and nephews to visit elderly who do not have children.	In effect June 2, 2020	Goh (2020)
The Bahamas	All of the Bahamas	AGE-ACT	Shopping	60+ y.o.	Nationwide curfew implemented from 9 p.m. to 5 a.m.. Residents not allowed to leave homes unless to purchase essential items. Shopping schedule implemented. All adults assigned specific times and days according to last name. Those with disabilities and persons aged 60 and above provided additional time slot from 6 a.m. to 12 p.m. on Tuesdays regardless of last name.	In effect April 6, 2020	Crisis24 (2020); Office of the Prime Minister (2020)
Turkey	All of Turkey	AGE	ALL	65+ y.o.	Citizens aged 65 and above restricted from leaving their homes.	In effect March 21, 2020	Zorlu (2020)
Turkey	Bursa	AGE-ACT	Public transportation	65+ y.o.	Senior citizens not allowed on public transport during morning and afternoon rush hours. Senior citizens also prohibited from attending social gatherings and from entering crowded places such as markets and open air bazaars after 2 p.m.	In effect August 17, 2020	Xinhua (2020)
Turkey	Bursa	AGE	ALL	65+ y.o.	Further restrictions imposed on the elderly by Bursa authorities. Elderly aged 65 and above only allowed out between 10 a.m. to 4 p.m. instead of 10 a.m. to 8 p.m.	Effective November 3, 2020	Xinhua (2020a)
Uganda	All of Uganda	ACT	Schools	ALL	Closure of schools.	School closures from March 2020 to January 10, 2022, with some openings and closures in between October 2020 to June 2021.	Atuhaire (2022) ; Ntabadde (2020)
USA	Many US states	AGE-ACT	Shopping	60+ y.o.	Dedicated hours when only the senior population is allowed to shop at super markets.	In effect March 2020.	Time Aguilera (2020)
USA	Illinois	AGE-ACT	Shopping	55+ y.o.	Special shopping hours reserved for elderly aged 55 and above and vulnerable populations implemented by some shops affiliated to the Illinois Retail Merchants Association.	No specified date	Illinois Retail Merchants Association (2022)
Wales	All of Wales	ACT	All educational establishments	ALL	Closure of all educational establishments.	In effect from March 20, 2020. Openings and closures through April 2021.	Welsh Government (2020)

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.
It is made available under a CC-BY-NC 4.0 International license.

Table EC.15: Examples of recommended targeted restrictions

Country	**Region**	**Type of Targeting**	**Involved Activities**	**Involved Age Groups**	**Summary of Recommended Restriction**	**Dates of Restriction**	**Source**
China	All of China	AGE-ACT	Travel	Elderly	Elderly and high-risk groups advised against travelling during China's week long May holiday	In effect April 23, 2020	Xinhua (2020b); National Working Committee on Children and Women under State Council (2020)
Croatia	All of Croatia	AGE	ALL	60+ y.o.	Elderly above the age of 60 advised to avoid crowded places. Routine visits to the doctor should be skipped unless urgent.	In effect March 3, 2020	Government of the Republic of Croatia (2020)
Czech Republic	All of Czech Republic	AGE	ALL	60+ y.o.	Seniors advised to leave the house as little as possible.	In effect March 20, 2020	Ministry of Health (2020)
Ireland	All of Ireland	AGE	ALL	70+ y.o.	Elderly above the age of 70 and medically vulnerable persons advised to cocoon (i.e., stay home as much as possible).	In effect March 27, 2020	Department of Health (2020); McGreevy and Gleeson (2020)
New Zealand	All of New Zealand	AGE	ALL	Older people	Older and vulnerable people encouraged to "hunker down".	Peak of Omicron outburst, early 2022.	?
Thailand	All of Thailand	AGE	ALL	Elderly	Elderly and their primary caregivers advised to stay at home as much as possible	In effect April 6, 2020	Department of Medical Services
The Bahamas	All of the Bahamas	AGE-ACT	Work	65-74 y.o.	Elderly between the age of 65-74 recommended to work from home.	In effect March 29, 2020	Office of the Prime Minister Communications Unit (2020)
The Bahamas	All of the Bahamas	AGE	ALL	75+ y.o.	Elderly above the age of 75 asked not to leave home.	In effect March 29, 2020	Office of the Prime Minister Communications Unit (2020)

AGE (ACT) refers to restrictions that target age groups (activities). AGE-ACT refers to restrictions that target both age groups and activities.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

References for Examples of Targeted Restrictions

1News (2021) Auckland high school seniors prepare to return to class. *1News* URL https://www.1news.co.nz/2021/10/25/auckland-high-school-seniors-prepare-to-return-to-class/, accessed April 29, 2022.

A3M Global Monitoring (2020) Covid-19 pandemic - hungary. *A3M Global Monitoring* URL https://global-monitoring.com/gm/page/events/epidemic-0001988.5I3CRSSqviVB.html?lang=en, accessed April 30, 2022.

Aguilera J (2020) Some supermarkets are launching senior-only hours during the coronavirus pandemic. not all retailers think that's a good idea. *Time* URL https://time.com/5804574/grocery-store-chains-seniors-coronavirus/, accessed April 28, 2022.

Andrić G (2020) Corona virus: "you work all day, so you run to get home before curfew - how serbia stopped for 50 days". *BBC* URL https://www-bbc-com.translate.goog/serbian/cyr/srbija-60733573?_x_tr_sl=sr&_x_tr_tl=en&_x_tr_hl=en&_x_tr_pto=sc, accessed April 29, 2022.

Atuhaire P (2022) Uganda schools reopen after almost two years of covid closure. *BBC* URL https://www.bbc.com/news/world-africa-59935605, accessed April 29, 2022.

CGTN (2021) Schools in new zealand's auckland to reopen from november 17. *CGTN* URL https://newsaf.cgtn.com/news/2021-11-11/Schools-in-New-Zealand-s-Auckland-to-reopen-from-November-17-155kca6Frdm/index.html, accessed on April 29, 2022.

Civilge (2020) Covid-19: Georgia announces nationwide lockdown, partial curfew. *UNA.ge* URL https://civil.ge/archives/344761, accessed April 30, 2022.

Crisis 24 (2020) Romania: Government announces lockdown measures on March 25 /update 2. *Crisis 24* URL https://crisis24.garda.com/alerts/2020/03/romania-government-announces-lockdown-measures-on-march-25-update-2, accessed April 30, 2022.

Crisis24 (2020) Bahamas: Additional restrictions for shopping introduced as of April 6 /update 4. *Crisis24* URL https://crisis24.garda.com/alerts/2020/04/bahamas-additional-restrictions-for-shopping-introduced-as-of-april-6-update-4, accessed April 30, 2022.

CSOmeter (2020) State of emergency in the republic of moldova. *CSOmeter* URL https://csometer.info/updates/state-emergency-republic-moldova, accessed April 30, 2022.

CTech (2020) Israel rolls back covid-19 restrictions as number of sick stabilizes. *CTech* URL https://www.calcalistech.com/ctech/articles/0,7340,L-3809065,00.html, accessed April 28, 2022.

Department of Health (2020) Guidance on cocooning to protect people over 70 years and those extremely medically vulnerable from covid-19 - updated guidance from 29 june. *Government of Ireland* URL https://www.gov.ie/en/publication/923825-guidance-on-cocooning-to-protect-people-over-70-years-and-those-extr/?fbclid=IwAR2TkpZADAmZpaFd6MCkmk2OElQzZ1gEuIq9BOPF3p9K2AF2yRU29PuqIKY, accessed April 30, 2022.

Department of Medical Services (????) Guidelines for caring for the elderly during the covid-19 outbreak .

Finnish Government (2020) Government, in cooperation with the president of the republic, declares a state of emergency in finland over coronavirus outbreak. *Finnish Government* URL https://valtioneuvosto.fi/en/-/10616/hallitus-totesi-suomen-olevan-poikkeusoloissa-koronavirustilanteen-vuoksi, accessed April 30, 2022.

Foy H (2020) Russian businesses prepare for fresh lockdowns as covid-19 cases soar. *The Financial Times* URL https://www.ft.com/content/f0696f19-fe0b-425c-8e73-3703bd44dec3, accessed April 30, 2022.

Goh YH (2020) Limited visits to parents or grandparents to be allowed from June 2 after circuit breaker. *Straits Times* URL https://www.straitstimes.com/singapore/limited-visits-to-parents-or-grandparents-to-be-allowed-from-june-2-seniors-urged-to-stay, accessed April 30, 2022.

Government of the Republic of Croatia (2020) Coronavirus protection measures. *Government of the Republic of Croatia* URL https://vlada.gov.hr/coronavirus-protection-measures/28950, accessed April 30, 2022.

Human Rights Watch (2020) Jordan: State of emergency declared. *Human Rights Watch* URL https://www.hrw.org/news/2020/03/20/jordan-state-emergency-declared, accessed April 30, 2022.

Illinois Retail Merchants Association (2022) Special shopping hours for vulnerable populations. *IRMA* URL https://irma.org/covid-19-senior-shopping/, accessed April 29, 2022.

Jabbarova A (2020) Azerbaijanis required to send sms notification before leaving their homes. *Global Voices* URL https://globalvoices.org/2020/04/10/azerbaijanis-required-to-send-sms-notification-before-leaving-homes/, accessed April 30, 2022.

Magid J (2020) Cabinet approves removal of age restriction for those returning to work. *Times of Israel* URL https://www.timesofisrael.com/liveblog-may-2-2020/, accessed April 28, 2022.

McGreevy R, Gleeson C (2020) Cocooning is advisory, not mandatory, government confirms. *The Irish Times* URL https://www.irishtimes.com/news/ireland/irish-news/cocooning-is-advisory-not-mandatory-government-confirms-1.4229569, accessed April 30, 2022.

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Ministry of Health (2020) Senioři by měli co nejvíce omezit vycházení ven. *Ministry of Health of the Czech Republic* URL https://koronavirus.mzcr.cz/semiori-by-meli-co-nejvice-omezit-vychazeni-ven/, accessed April 30, 2022.

Mogul R, Sud V (2022) After more than 600 days shut out, delhi's students just want to go back to school. *CNN* URL https://edition.cnn.com/2022/01/27/india/india-delhi-schools-reopen-600-days-intl-hnk/index.html, accessed April 29, 2022.

National Working Committee on Children and Women under State Council (2020) Experts from the chinese center for disease control and prevention remind that high-risk groups such as pregnant women, the elderly, and patients with chronic diseases are not recommended to travel. *National Working Committee on Children and Women under State Council* URL http://www.nwccw.gov.cn/2020-04/24/content_283708.htm, accessed April 30, 2022.

Ntabadde C (2020) "i wish i could have some magic, i would just say let schools open right away". *UNICEF* URL https://www.unicef.org/uganda/stories/i-wish-i-could-have-some-magic-i-would-just-say-let-schools-open-right-away, accessed April 29, 2022.

Office of the Prime Minister (2020) Food shopping schedule. *Office of the Prime Minister* URL https://www.bahamas.gov.bs/wps/wcm/connect/df5a4569-9c5b-40e4-85f3-d764584aef5c/Statement+-+Food+shopping+schedule.pdf?MOD=AJPERES, accessed April 30, 2022.

Office of the Prime Minister Communications Unit (2020) Prime minister announces three new cases of covid-19 and plans to extend state of emergency for eight days. *The Government of the Bahamas* URL https://tinyurl.com/3vppjmm6, accessed April 30, 2022.

Parsitau DS, Jepkemei E (2020) How school closures during covid-19 further marginalize vulnerable children in kenya. *Brookings; Education Plus Development* URL https://www.brookings.edu/blog/education-plus-development/2020/05/06/how-school-closures-during-covid-19-further-marginalize-vulnerable-children-in-kenya/#:~:text=On%20March%2015%2C%202020%2C%20the,devastating%20consequences%20for%20marginalized%20learners., accessed April 30, 2022.

Pertuz-Cruz SL, Molina-Montes E, Rodríguez-Pérez C, Guerra-Hernández EJ, Cobos de Rangel OP, Artacho R, Verardo V, Ruiz-Lopez MD, García-Villanova B (2021) Exploring dietary behavior changes due to the covid-19 confinement in colombia: A national and regional survey study. *Frontiers in Nutrition* URL https://www.frontiersin.org/articles/10.3389/fnut.2021.644800/full#B5, accessed April 29, 2022.

Picheta R, Qiblawi T (2020) Jordan eases lockdown after total curfew leads to chaos. *CNN* URL https://edition.cnn.com/2020/03/25/middleeast/jordan-lockdown-coronavirus-intl/index.html, accessed April 30, 2022.

Portafolio (2020) Tumban cuarentena obligada para adultos mayores de 70 años. *Portafolio* URL https://www.portafolio.co/economia/gobierno/tumban-cuarentena-obligatoria-para-adultos-mayores-de-70-anos-coronavirus-en-colombia-hoy-12-agosto-20 accessed April 29, 2022.

President of the Republic of Colombia (2020) Estas son las medidas que se deben tener en cuenta para la salida de los niños durante la cuarentena. *President of the Republic of Colombia* URL https://id.presidencia.gov.co/Paginas/prensa/2020/Estas-son-las-medidas-que-se-deben-tener-en-cuenta-para-la-salida-de-los-ninos-durante-la-cuarentena-200507.aspx3, accessed April 29, 2022.

Reuters (2020) Kenya partially reopens schools, 6 months after covid shuttered them. *Vox* URL https://www.voanews.com/a/covid-19-pandemic_kenya-partially-reopens-schools-6-months-after-covid-shuttered-them/6197024.html, accessed April 30, 2022.

Reuters Staff (2020a) Colombia declares coronavirus state of emergency, orders elderly to stay home. *Reuters* URL https://www.reuters.com/article/us-health-coronavirus-colombia-idUSKBN2150AI, accessed April 30, 2022.

Reuters Staff (2020b) Hungary imposes restricted shopping hours to protect elderly in pandemic. *Reuters* URL https://www.reuters.com/article/uk-health-coronavirus-hungary-idUKKBN2831OW, accessed April 30, 2022.

Reuters Staff (2020c) Reopening schools in Denmark did not worsen outbreak, data shows. *Reuters* URL https://www.reuters.com/article/us-health-coronavirus-denmark-reopening/reopening-schools-in-denmark-did-not-worsen-outbreak-data-shows-idUSKBN2341N7, accessed April 30, 2022.

Reuter's Staff (2020) Serbia's elderly venture out for dawn food run. *Reuters* URL https://www.reuters.com/article/health-coronavirus-serbia-elderly-idUSL8N2BM0A5, accessed April 29, 2022.

Sadikovic M, Synovitz R (2020) Coronavirus in court: Bosnia's age-based lockdowns are ruled discriminatory. *RadioFreeEurope/RadioLiberty* URL https://www.rferl.org/a/coronavirus-in-court-bosnia-s-age-based-lockdowns-are-ruled-discriminatory/30574453.html, accessed April 30, 2022.

Saraya News (2020) Adaileh: The government continues to ban children and the elderly from leaving. *Saraya News* URL https://www.sarayanews.com/article/611788?fbclid=IwAR0B__fhu348B9y0BAnVM6UN5tIOSj-kRa9r3i3wYenXbrIeQSmMN-GgE6c, accessed April 30, 2022.

Staff R (????) Bosnian region eases lockdown on seniors, children after court ruling .

medRxiv preprint doi: https://doi.org/10.1101/2021.03.23.21254155; this version posted November 29, 2022. The copyright holder for this preprint **(which was not certified by peer review)** is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity.

It is made available under a CC-BY-NC 4.0 International license .

Stevanović K (2020) Corona virus and over 65 in serbia: "it never occurs to me to vampire". *BBC* URL https://www.bbc.com/serbian/cyr/srbija-52266134, accessed April 29, 2022.

The Government of the Republic of Serbia (2020) Senior citizens to go grocery shopping on saturday instead of sunday. *The Government of the Republic of Serbia* URL https://www.srbija.gov.rs/vest/en/153257/senior-citizens-to-go-grocery-shopping-on-saturday-instead-of-sunday.php, accessed April 29, 2022.

Timpul (2020) Drastic decisions /// the government prohibits people over the age of 63 from being outside. it is forbidden for all people to be in public spaces (parks, recreational areas) / quarantine in the villages of bălceana and sofia. *Timpul* URL https://timpul.md/articol/decizii-drastice-guvernul-interzice-persoanelor-in-varsta-de-peste-63-ani-aflarea-in-afara-domiciliului-se-interzice-tut html, accessed April 30, 2022.

Vasovic A (2020) Serbia imposes night curfew, orders elderly indoors. *Reuters* URL https://www.reuters.com/article/us-health-coronavirus-serbia/serbia-imposes-night-curfew-orders-elderly-indoors-idUSKBN2143XR, accessed April 29, 2022.

Welsh Government (2020) Timeline of school closures during the coronavirus (covid-19) pandemic, march 2020 to april 2021. *Welsh Government* URL https://gov.wales/timeline-school-closures-during-coronavirus-covid-19-pandemic, accessed April 30, 2022.

Xinhua (2020) Hungary set to cautiously reopen schools, kindergartens. *Xinhua* URL http://www.xinhuanet.com/english/europe/2021-04/16/c_139885633.htm, accessed April 30, 2022.

Xinhua (2020) Turkey's bursa province adopts new covid-19 restrictions for elderly. *Xinhua* URL http://www.xinhuanet.com/english/2020-08/17/c_139297485.htm, accessed April 30, 2022.

Xinhua (2020a) Turkey's bursa tightens covid-19 lockdown on elderly. *Xinhua* URL http://www.news.cn/english/2020-11/04/c_139488890.htm, accessed April 29, 2022.

Xinhua (2020b) Xinhua headlines: China's tourism rebounds during may day holiday as coronavirus eases. *Xinhua* URL http://www.xinhuanet.com/english/2020-05/05/c_139032433.htm, accessed April 30, 2022.

Zhang LM, Han GY (2020) Coronavirus: Govt agencies to suspend activities for seniors for 14 days to cut risk of transmission. *The Straits Times* URL https://www.straitstimes.com/singapore/health/govt-agencies-to-suspend-activities-for-seniors-for-14-days-starting-march-11-to, accessed April 27, 2022.

Zorlu F (2020) Thousands of citizens age 65 and above, restricted to their homes due to covid-19, take breath of fresh air. *Anadolu Agency* URL https://www.aa.com.tr/en/latest-on-coronavirus-outbreak/turkey-elderly-given-one-time-permission-to-go-out/1835850, accessed April 30, 2022.