

Optimized Targeted Confinements for Future Pandemic Response

Sergio Camelo

Stanford University, camelo@stanford.edu

Dragos Florin Ciocan

INSEAD, florin.ciocan@insead.edu

Dan A. Iancu

Stanford University, daniancu@stanford.edu

Xavier S. Warnes

Indiana University, xwarnes@iu.edu

Spyros I. Zoumpoulis

INSEAD, spyros.zoumpoulis@insead.edu

Effective preparation for future pandemics requires a clear understanding of how best to deploy non-pharmaceutical interventions, especially population confinements. Experience during COVID-19 shows that many jurisdictions tailor confinements by population group or by activity, yet such targeting is operationally demanding and politically sensitive, making rigorous cost–benefit quantification indispensable. We develop a modeling framework in which confinements can be targeted along two dimensions—age group and activity—to minimize a composite loss that combines mortality and foregone economic output. A stylized, analytically tractable version of the model yields closed-form optimal confinement rules and conditions that reveal when targeting generates welfare gains, and clarifies how these gains depend on key epidemiological and economic parameters. We find that targeting yields gains in certain L-shaped parameter regions, and gains behave non-monotonically in parameters. To translate these insights into practice and quantify their impact, we introduce a structured optimization procedure that couples model-predictive-control techniques with trust-region methods to derive high-quality solutions. A full-scale implementation for COVID-19 in Île-de-France demonstrates that targeting by age or by activity delivers Pareto improvements relative to non-targeted, uniform policies; and targeting along both dimensions delivers Pareto improvements relative to targeting along just one. We extend the model and the algorithm to deal with ambiguity in problem parameters through a distributionally-robust approach; we find that gains from targeting persist under parameter ambiguity and surprisingly, more ambiguity can increase these gains. Applying a structured optimization approach to derive optimized targeted confinements can therefore be highly beneficial even in the early stages of a pandemic, when estimates of epidemiological parameters are unreliable.

1. Introduction

Following the COVID-19 crisis, policymakers have refocused their attention on strengthening preparedness and response for future pandemics (Khor and Heymann 2021, Sirleaf and Clark 2021). The COVID-19 experience has underscored the pivotal contribution of non-pharmaceutical interventions—most notably, large-scale population confinement—in curbing transmission during the critical period before effective therapies and vaccines become available (Group et al. 2023). Designing such confinement policies requires recognizing that individuals engaged in different activities generate markedly heterogeneous health and economic externalities. Tailoring restrictions to this heterogeneity offers a powerful lever for reducing overall societal harm, yet it also raises concerns about equity and the potential for discriminatory implementation. To inform decision-making in the early stages of an emerging epidemic, we develop a quantitative framework that evaluates the benefits of targeted interventions and illustrate it with a detailed empirical case study.

Targeting has been implemented in several different ways during the COVID-19 pandemic. 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 indicate, targeted interventions offer significant advantages yet also entail noteworthy risks. When carefully calibrated, they can improve both health and economic outcomes, giving policymakers a more nuanced lever for navigating difficult trade-offs. Nevertheless, such fine-grained measures are harder to operationalize—especially in the early stages of a new pandemic, when key epidemiological parameters are uncertain and high-resolution data are scarce—and they may engender perceptions of inequity or even outright discrimination.

Because many real-world pandemic policies already incorporate some degree of targeting, it is crucial to model such measures transparently and to quantify both their benefits and their potential drawbacks as we prepare for future outbreaks. This prompts several natural research questions. First, to what extent can progressively finer targeting improve health and economic outcomes, and through what mechanisms are these improvements realized? Second, how do the gains from targeting vary with key epidemiological parameters such as the basic reproduction number R_0 and the infection fatality rate? Third, when decision makers face substantial uncertainty about such parameters, as is typical in the early stages of an epidemic, should they abandon targeting in favor of simpler uniform policies, or can targeting still confer meaningful gains despite the ambiguity?

1.1. Contributions

This study makes two core contributions. First, it offers policy-relevant insights into whether, and under what conditions, targeted confinement strategies can reduce the combined health and economic costs of a future pandemic relative to more uniform restrictions. Second, it introduces a quantitative framework, together with accompanying computational tools, that enables policymakers to operationalize these strategies and rigorously quantify their benefits. To elaborate, we make the following specific contributions.

- 1. Model for targeted interventions.** We anchor our analysis around an optimization model for targeted interventions, which we embed within a multi-period, multi-group SEIR epidemiological model that differentiates policies based on both population groups and activities, and balances the lost economic value with the cost of deaths.
- 2. Simplified theoretical model for stylized insights.** To assess whether targeting can generate welfare gains, we first analyze a stylized, analytically tractable model that collapses the pandemic evolution into two stages, while still accommodating all dimensions of heterogeneity that interest us. Within this framework, we derive closed-form optimal confinement rules: high activity levels are allowed for groups in activities with the largest ratio of (i) the net economic value they would generate if kept fully open—namely, their intrinsic economic output minus the externality death cost imposed on other groups at those groups’ prevailing activity levels—to (ii) their own intra-group mortality cost. The model further pinpoints the parameter regimes in which targeting improves outcomes and shows how these regimes shift with key epidemiological factors and with the policymaker’s relative weight on mortality versus economic losses. To the best of our knowledge, these are the first theoretical results that jointly characterize group- and activity-based targeted confinements.
- 3. Optimization framework for quantifying gains.** Although the stylized model yields valuable conceptual insights, it cannot capture the full complexity of multi-period pandemic dynamics or quantify the magnitude of targeting gains in realistic environments. To bridge this gap, we develop a structured optimization framework for the full model that operationalizes the theoretical prioritization rule and enables quantitative evaluation with empirical data. The framework integrates publicly available information with a robust-optimization formulation that explicitly accounts for parameter ambiguity. Because the resulting problem is highly non-convex, we introduce a novel solution method—Re-Optimization with Linearized Dynamics (ROLD)—which couples a model-predictive-control linearization with trust-region refinements to generate high-quality policies in real time. We apply ROLD in a detailed case

study of COVID-19 in Île-de-France (a region surrounding Paris, with 12 million inhabitants) and in additional scenarios involving other pandemics and geographies, thereby confirming the theoretical predictions and quantifying the potential gains from targeted confinement.

Insights for policymakers. From our findings, we synthesize the following messages for how policymakers should approach targeting:

1. The magnitude—and even the existence—of welfare gains from targeting hinges on the values of epidemiological parameters and on the relative importance that policymakers place on economic and mortality losses. Gains exist only inside an L-shaped region of the (R_0, χ) parameter space, where R_0 is the basic reproduction number and χ is the cost/weight that a policymaker places on a death. An analogous L-shaped frontier arises when R_0 is replaced by a group-specific disease progression parameter such as the infection fatality rate. Outside this region, a uniform, non-targeted policy is optimal; inside the region, interpretable conditions involving key epidemiological and economic parameters identify those age groups or activities prioritized for confinement and the gains from targeting can be significant, but the gains are also non-monotonic in problem parameters. To assess whether a specific situation calls for targeting and to quantify the gains from targeting therefore requires use of rigorous optimization-based approaches.
2. Using our ROLD algorithmic framework, we find that optimized targeted confinements can produce *large* gains. In our Île-de-France case study, even policies that target only based on age group or only based on activity already yield *Pareto improvements* over a uniform, non-targeted baseline: they lower aggregate economic losses without increasing pandemic fatalities. However, neither policy Pareto-dominates the other, and the choice between age group and activity-based targeting is subtle: our theoretical and empirical results suggest that activity-based targeting can be more effective at low-to-moderate basic reproduction numbers R_0 and when policy makers place a large weight χ on mortality losses, whereas age-based targeting is more effective at high reproduction numbers R_0 and low-to-moderate weight χ placed on

mortality losses. Crucially, targeting both age groups and activities goes further, Pareto-dominating every alternative that targets only one dimension.

3. The value of targeting persists even when there is uncertainty about important model parameters. Moreover, the common intuition that greater ambiguity should tilt policymakers toward uniform, non-targeted rules fails: as ambiguity levels rise, the welfare gains from targeting also *rise* in some important regions of the parameter space. Consequently, even in the early stages of an outbreak—when estimates of transmission, fatality, or other disease dynamics are most unreliable—optimized targeted confinement can still outperform uniform restrictions and deliver substantial benefits.

The paper is organized as follows. Section 2 reviews relevant literature. Section 3 describes our baseline pandemic model and formulates the underlying targeting problem. Section 4 presents the stylized version of the model and extracts theoretical insights. Section 5 then describes the ROLD algorithm to solve the more complex, general model. Section 6 documents our Île-de-France case study and the gains from targeting using ROLD. Section 7 extends our model and ROLD algorithm to a robust formulation that accounts for parameter ambiguity. Lastly, Section 8 concludes with a discussion. The Appendix provides modeling details (EC.1, EC.2), details and proofs for the theoretical results (EC.3), algorithmic details for ROLD (EC.4, EC.5), the parametrization for the empirical case study (EC.6), and robustness checks and sensitivity analyses (EC.7).

2. Literature Review

The literature on pandemic response, particularly following COVID-19, is vast, so we focus our 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. (2021), Matrajt et al. (2021), 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. 2021).

Optimization of interventions in epidemiological models. Our work relates to research that combines epidemiological modeling and optimization techniques to design improved interventions. In general, an epidemic is modeled by a compartmental model, where interventions change the parameters that describe the epidemic with the goal to minimize the health (and economic) burden. Although an analytical characterization of the optimal solution is possible in special cases or in those when the control is single-dimensional and thus dynamic programming approaches are amenable (Brandeau et al. 2003, Barnett et al. 2023, Calvia et al. 2024), the problem is generally intractable and, similar to our ROLD framework, research has focused on proposing heuristic algorithms and approximations for solving the general problem (Zaric and Brandeau 2001, 2002). The algorithm proposed in Bertsimas et al. (2020) and our heuristic crucially depend on solving linearized versions of the true SEIR dynamics that are tractable via commercial solvers. However, Bertsimas et al. (2020) focuses on vaccine allocation decisions, whereas we capture the dynamics of differential confinements and also allow activity-based targeting. Bose et al. (2021), Pataro et al. (2021), Morris et al. (2021) also borrow from the optimal control literature, but 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 does not account for time in the calculation of health or economic impact. Last, several studies from the operations research community have proposed optimization models to support the allocation of ventilators during epidemics (Huang et al. 2017, Adelman 2020, Mehrotra et al. 2020, Bertsimas et al. 2021).

Our paper is also related to a large stream of work that derives prescriptive insights for managing the COVID-19 pandemic. Kaplan (2020) summarizes modeling studies that supported local decisions on event crowd-size restrictions, hospital surge planning, and timing of activity restrictions during COVID-19 response. Several papers simulate a small number of candidate policies for social distancing, e.g., full lockdown versus school-only lockdown (Kucharski et al. 2020, El Housni et al. 2020, Bertsimas et al. 2021), compare a number of current and counterfactual lockdown policies that differ in their schedule of relaxations (Bolori and Saghafian 2023), or 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). These approaches do not use formal optimization—although they may allow for the solutions to change as the data changes—and, when considering a more complex policy space like in our targeting model, could lead to significantly sub-optimal results and misleading conclusions. Navabi-Shirazi et al. (2022) use multi-criteria optimization to select the mode (remote, in-person, hybrid) of university courses and assign classrooms, under severely reduced capacities due to COVID-19 social distancing measures. The study of Fotouhi et al. (2021) helps policymakers design curbside restrictions in meal delivery operations that reduce curbside crowding, thus increasing public safety during a pandemic, yet enable delivery companies to retain their profitability.

Parameter uncertainty. Our robust SEIR formulation is inspired by a longstanding stream in the operations research and optimal control literatures that seeks to optimize the worst case over a set that captures parameter uncertainty (Zhou and Doyle 1998, Dullerud and Paganini 2013, Ben-Tal and Nemirovski 2002, Bertsimas et al. 2011). On the epidemiology side, our approach relates to a stream of research that embeds uncertainty into deterministic compartmental models, with several approaches having emerged in the literature. One approach is related to stochastic SEIR models, adding stochastic noise terms to the deterministic dynamics of the SEIR (Greenwood and Gordillo 2009); such models have been used to estimate disease progression or control non-targeted interventions in the context of COVID-19 (Olivares and Staffetti 2022, Fu et al. 2021, He et al. 2020,

Lekone and Finkenstädt 2006). Another approach, which our paper adheres to, is to formulate the control problem in a way that accounts for parameter uncertainty. In this vein, Köhler et al. (2021), Bhardwaj et al. (2020), Wan et al. (2023), Lobato et al. (2021), Faranda and Alberti (2020) propose robust formulations of non-targeted interventions under SEIR-type models, where key problem parameters are allowed to belong to an uncertainty set. A closely related paper to ours is by Barnett et al. (2023), who also formulate a robust framework that allows for ambiguity in epidemiological and economic parameters. However, the ultimate focus of Barnett et al. (2023) is different: within the class of non-targeted interventions, they compare the lockdown profiles corresponding to an ambiguity-neutral decision maker versus one that is ambiguity-averse. In contrast, our focus is on a more complex and high-dimensional model in which we shed light on the benefits of targeting and how ambiguity impacts them.

3. Model and Optimization Problem

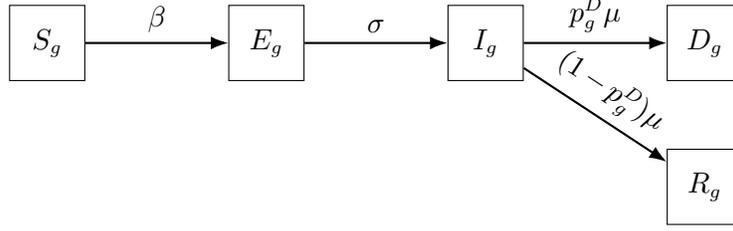
We develop a controlled, multi-group SEIR model that includes time-dependent confinements that can be targeted based on *age groups* and types of *activities* that individuals engage in. The framework is flexible and can be extended to capture other targeted interventions, such as testing or vaccinations, as well as additional restrictions that make targeting more fair or practical, such as fairness constraints, endogenous feedback, etc. We discuss several of these in Section 8.

3.1. Some Notation

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

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

population groups that interact with each other. In our case study we use three groups $g \in \mathcal{G}$ determined by age, with the youngest group capturing individuals with age of 0-19 years, the group in the middle capturing individuals aged 20-64 years, and the oldest group capturing individuals with age of 65 years or above. We use \mathcal{G} to denote the collection of considered age groups. Time is discrete, indexed by $t = 0, 1, \dots, T$ 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). The model also has states $R_g(t)$ (recovered) and $D_g(t)$ (deceased). Each state represents the number of individuals in that compartment at 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 rate¹ β . Exposed individuals transition to the infectious state at a rate σ and infectious individuals transition out of the infectious state at a rate μ . An infectious individual 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 , $N_g(t) := S_g(t) + E_g(t) + I_g(t) + R_g(t)$, and let $\mathbf{X}_t = [S_g(t), E_g(t), I_g(t), R_g(t), D_g(t)]_{g \in \mathcal{G}}$ denote the full state of the system (across groups) at time $0 \leq t \leq T$. We denote the number of compartments by $|\mathcal{X}|$, so the dimension of \mathbf{X}_t is $|\mathcal{G}||\mathcal{X}| \times 1$.

Controls. Individuals interact in activities $\mathcal{A} = \{\text{work, transport, leisure, school, home, other}\}$. These interactions generate *social contacts* that drive the rate of new infections.

¹ The reproduction number R_0 and the transmission rate β are related and comonotonic (Appendix EC.6.1).

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 $\ell_g^{\text{home}}(t) = 1$, meaning that the number of social contacts at home is unchanged irrespective of confinement policy.² 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

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

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).

Let $\mathbf{u}_t = [\boldsymbol{\ell}_g(t)]_{g \in \mathcal{G}}$ denote the vector of all decisions at time $t \in \{0, 1, \dots, T-1\}$, 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 \mathbf{u}_t is $|\mathcal{G}||\mathcal{U}| \times 1$ and the decisions are constrained: $\mathbf{u}_t \in [0, 1]^{|\mathcal{G}||\mathcal{U}|}$ for every $t \in \{0, 1, \dots, T-1\}$.

²The 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.

We can now write a complete set of discrete dynamical equations for the controlled SEIR model ((EC.1)-(EC.6) in Appendix EC.1) and summarize these using the function

$$\mathbf{F}_t(\mathbf{X}_t, \mathbf{u}_t) := \frac{\Delta \mathbf{X}_t}{\Delta t}, \quad (2)$$

where $\Delta \mathbf{X}_t := \mathbf{X}_{t+1} - \mathbf{X}_t$.

Targeting variants. We consider four versions of targeting: no targeting whatsoever (NO-TARGET), targeting age groups only (AGE), targeting activities only (ACT), or targeting both (AGE-ACT). We can think of each form of targeting as enforcing a constraint on activity levels—for example $\ell_g^a(t) = \ell_{g'}^{a'}(t)$ for all $g, g' \in \mathcal{G}$ and $a, a' \in \mathcal{A}$ in the case of NO-TARGET.

3.3. Objective

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

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 $\boldsymbol{\ell} := [\ell_g]_{g \in \mathcal{G}}$ 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 on activity levels only through the group's activity level in school (ℓ_g^{school}). For (b), we assume that an individual who is deceased generates no economic value. 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 V that would have been generated during the pandemic under a “no pandemic” scenario and the value generated during the pandemic, plus the future economic output lost due to deaths.

$$\text{Economic Loss}(\mathbf{u}_{0:T-1}) := V - \sum_{t=0}^{T-1} \sum_{g \in \mathcal{G}} \left(v_g(\ell(t)) \cdot N_g(t) \right) + \sum_{g \in \mathcal{G}} v_g^{\text{life}} \cdot D_g(T), \quad (3)$$

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

To allow policymakers 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 ($\chi \rightarrow \infty$). The policymaker aims to find control policies for confinement that minimize the sum of mortality and economic losses³:

$$J(\mathbf{u}_{0:T-1}) := \text{Economic Loss}(\mathbf{u}_{0:T-1}) + \chi \cdot \text{Total Deaths}(\mathbf{u}_{0:T-1}). \quad (4)$$

3.4. Optimization Problem

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

$$\begin{aligned} & \min_{\mathbf{u}_{0:T-1}} J(\mathbf{u}_{0:T-1}) \\ \text{s.t.} \quad & \mathbf{X}_{t+1} = \mathbf{F}_t(\mathbf{X}_t, \mathbf{u}_t), \quad \forall 0 \leq t \leq T-1 \\ & \mathbf{u}_{0:T-1} \in [0, 1]^{T|\mathcal{G}||\mathcal{U}|}. \end{aligned} \quad (5)$$

³ Similarly to Acemoglu et al. (2021), 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 for each value of χ ; this makes our results informative for policymakers who may have differing views on the right value for such a key parameter.

4. Analytical Results in a Simple Stylized Model

Before addressing (5), we first turn to investigating a simplified model in which we can analytically solve for optimal policies and tease out the role of targeting.

4.1. Two-Period, Discrete-Time SIR Model

We consider two groups, $\mathcal{G} = \{1, 2\}$, and two activities, $\mathcal{A} = \{1, 2\}$. Time is divided into two periods $t \in \{0, 1\}$, with length Δ_1 and Δ_2 respectively. The policymaker can only reduce activity levels in period $t = 0$, to a level no smaller than ε .⁴ In period $t = 1$, the economy reopens fully. Social mixing generates contacts according to (1), with $\alpha = 1$. A cure arrives at the end of period $t = 1$, so no infections or deaths are possible henceforth. The epidemiological model is identical to our base model, except that any susceptible individual that has a contact with an infected individual directly transitions into an infectious state, so we ignore the exposed chamber E_g .⁵ Economic value is generated according to the linear form $v_g(\boldsymbol{\ell}(t)) := \sum_a v_g^a \ell_g^a(t)$.

This simple model variation is rich enough to replicate the key trade-offs that motivate our study: the confinements during period $t = 0$ sacrifice economic value during that period for the benefit of reducing deaths (and associated economic losses) in period $t = 1$ and subsequently.

Let $\mathbf{u} = [u_1^1, u_2^1, u_1^2, u_2^2]^\top$ denote the activity levels in period $t = 0$, with components sorted by activity, so $u_g^a = \ell_g^a(0)$. The following result reformulates the policymaker's problem.

PROPOSITION 1 (Reformulation). *The problem of finding activity levels $\mathbf{u} \in [\varepsilon, 1]^{|\mathcal{G}||\mathcal{A}|}$ to minimize the total loss objective (4) in the two-period model is equivalent to the optimization problem:*

$$\max_{\mathbf{u} \in [\varepsilon, 1]^4} \mathbf{r}^\top \mathbf{u} - \frac{1}{2} \mathbf{u}^\top \mathbf{Q} \mathbf{u}, \quad (6)$$

⁴ Our base model only includes a lower bound for “home.” We can readily extend all our qualitative results here to bounds that are group- or activity-specific, but we retain the simpler assumption to streamline exposition.

⁵ Our qualitative insights hold even in a model with an E_g chamber, but that requires at least three periods so as to reflect the effects of confinements from $t = 0$ on deaths. For simplicity, we use two periods and eliminate E_g .

where $\mathbf{r} = [r_1^1, r_2^1, r_1^2, r_2^2]^\top$ with $r_g^a := \Delta_0 v_g^a N_g(0)$, and the matrix $Q \in \mathbb{R}^{|\mathcal{G}||\mathcal{A}| \times |\mathcal{G}||\mathcal{A}|}$ is given by

$$Q = \begin{pmatrix} Q^1 & 0 \\ 0 & Q^2 \end{pmatrix}, \text{ where } Q^a = \begin{pmatrix} B_1^a & B_{12}^a \\ B_{12}^a & B_2^a \end{pmatrix} \text{ for } a \in \mathcal{A} \text{ and} \quad (7a)$$

$$B_g^a = 2\beta \Delta_0 \Delta_1 A_g C_{gg}^a \frac{I_g(0)}{N_g(0)} \forall g \in \mathcal{G}, \quad B_{12}^a = \beta \Delta_0 \Delta_1 \left(A_1 C_{12}^a \frac{I_2(0)}{N_2(0)} + A_2 C_{21}^a \frac{I_1(0)}{N_1(0)} \right), \quad (7b)$$

$$A_g = (\chi + v_g^{\text{life}}) p_g^D \mu_g S_g(0) \text{ for } g \in \mathcal{G}. \quad (7c)$$

Proposition 1 exactly highlights the key trade-offs facing the policymaker, who is seeking to maximize the total economic value in period $t = 0$, $\mathbf{r}^\top \mathbf{u}$, net of the total cost of deaths created in period $t = 1$, $\frac{1}{2} \mathbf{u}^\top Q \mathbf{u}$. The parameter \mathbf{r} has components $r_g^a = \Delta_0 v_g^a N_g(0)$ for each group's *intrinsic economic value* in each activity, i.e., the total economic value that each group could generate when fully active in each activity a . To see why $\frac{1}{2} \mathbf{u}^\top Q \mathbf{u}$ captures the cost of deaths, note that A_g is a proxy for group g 's death cost per infected contact: the expected economic and mortality-related cost associated with a susceptible in group g getting infected (accounting for progression towards death, $\mu_g \cdot p_g^D$, and subsequent cost, $\chi + v_g^{\text{life}}$) multiplied by the number of susceptibles in group g . Thus, when activity a is fully open, $\frac{1}{2} B_g^a$ measures group g 's *intra-group death cost*, i.e., the total death cost that group g imposes upon itself through within-group contacts, and B_{12}^a measures the *inter-group death cost*, i.e., total death cost that groups impose on each other through across-group contacts. The quadratic term $\frac{1}{2} \mathbf{u}^\top Q \mathbf{u}$ adds such costs over all activities and all groups, accounting for the allowed activity levels \mathbf{u} in period $t = 0$.

Problem (6) is a constrained quadratic optimization problem that is separable across activities; within each activity $a \in \mathcal{A}$, the objective is component-wise concave in $\mathbf{u}^a = [u_1^a, u_2^a]^\top$, but it is jointly concave in \mathbf{u}^a if and only if $D_B^a := \det(Q^a) = B_1^a B_2^a - (B_{12}^a)^2 \geq 0$.

4.2. Which Groups and Activities to Target For Confinement?

We start by providing an intuitive characterization for the dual-targeted policy, showing how it prioritizes specific groups or activities for confinement. In this result and subsequently in this section, we denote NO-TARGET by NT, AGE-ACT by AA, and $h \in \mathcal{G} : h \neq \bar{g}$ by \bar{g} .

PROPOSITION 2 (**Optimal Dual Targeting; Comparative Statics**). *In the AA policy:*

- (i) For each activity $a \in \mathcal{A}$, the optimal solution $(\mathbf{u}^a)^*$ is given by Table 1 and the conditions therein if $D_B^a \geq 0$, and by one of the expressions in Cases 2-9 from Table 1 if $D_B^a < 0$.
- (ii) $(u_g^a)^*$ is increasing with $r_g^a, N_g(0), N_h(0)$, and is decreasing with any parameter in the set

$$\{\beta, \chi\} \cup \{v_h^{\text{life}}, p_h^D, \mu_h, S_h(0), I_h(0) : h \in \mathcal{G}\} \cup \{C_{gh}^a : g, h \in \mathcal{G}\} \cup \{r_{\bar{g}}^a\}.$$

Case	$(\mathbf{u}^a)^* = (u_1^a, u_2^a)^*$	Conditions on parameters
1. (I,I)	$\left(\frac{r_1^a B_2^a - r_2^a B_{12}^a}{D_B^a}, \frac{r_2^a B_1^a - r_1^a B_{12}^a}{D_B^a} \right)$	$\varepsilon < \frac{r_1^a B_2^a - r_2^a B_{12}^a}{D_B^a} < 1, \varepsilon < \frac{r_2^a B_1^a - r_1^a B_{12}^a}{D_B^a} < 1$
2. (L,I)	$\left(\varepsilon, \frac{r_2^a - B_{12}^a \varepsilon}{B_2^a} \right)$	$\varepsilon < \frac{r_2^a - B_{12}^a \varepsilon}{B_2^a} < 1, \quad r_1^a \leq B_1^a \varepsilon + B_{12}^a \left(\frac{r_2^a - B_{12}^a \varepsilon}{B_2^a} \right)$
3. (I,L)	$\left(\frac{r_1^a - B_{12}^a \varepsilon}{B_1^a}, \varepsilon \right)$	$\varepsilon < \frac{r_1^a - B_{12}^a \varepsilon}{B_1^a} < 1, \quad r_2^a \leq B_2^a \varepsilon + B_{12}^a \left(\frac{r_1^a - B_{12}^a \varepsilon}{B_1^a} \right)$
4. (U,I)	$\left(1, \frac{r_2^a - B_{12}^a}{B_2^a} \right)$	$\varepsilon < \frac{r_2^a - B_{12}^a}{B_2^a} < 1, \quad r_1^a \geq B_1^a + B_{12}^a \frac{r_2^a - B_{12}^a}{B_2^a}$
5. (I,U)	$\left(\frac{r_1^a - B_{12}^a}{B_1^a}, 1 \right)$	$\varepsilon < \frac{r_1^a - B_{12}^a}{B_1^a} < 1, \quad r_2^a \geq B_2^a + B_{12}^a \frac{r_1^a - B_{12}^a}{B_1^a}$
6. (L,L)	$(\varepsilon, \varepsilon)$	$r_1^a \leq \varepsilon(B_1^a + B_{12}^a), \quad r_2^a \leq \varepsilon(B_2^a + B_{12}^a)$
7. (L,U)	$(\varepsilon, 1)$	$r_1^a \leq B_1^a \varepsilon + B_{12}^a, \quad r_2^a \geq B_2^a + B_{12}^a \varepsilon$
8. (U,L)	$(1, \varepsilon)$	$r_2^a \leq B_2^a \varepsilon + B_{12}^a, \quad r_1^a \geq B_1^a + B_{12}^a \varepsilon$
9. (U,U)	$(1, 1)$	$r_1^a \geq B_1^a + B_{12}^a, \quad r_2^a \geq B_2^a + B_{12}^a$

Table 1 Optimal activity levels under dual targeting for a given activity a . The conditions identified are necessary for optimality and are sufficient if $D_B^a = B_1^a B_2^a - (B_{12}^a)^2 \geq 0$.

Proposition 2(i) shows that policymakers could rely on simple rules to decide which groups and activities to keep open, driven by the intuitive ratios in Table 1. Each ratio follows the same pattern. The numerator is the *net* economic value that group g could generate if fully open in the activity, given by the group's intrinsic economic value r_g^a net of the externality it imposes on the other group \bar{g} , i.e., the death cost $B_{g\bar{g}}^a u_{\bar{g}}^a$ that group g imposes on group \bar{g} through inter-group contacts, when taking into account group \bar{g} 's activity level.⁶ The denominator is the total death

⁶ Recall that $B_{g\bar{g}}^a = B_{\bar{g}g}^a$, so the externality term is also the death cost that group g itself incurs from inter-group contacts. For case (I,I), the proof shows that $(\mathbf{u}^a)^*$ exactly satisfies the equations $u_g^a = (r_g^a - B_{g\bar{g}}^a u_{\bar{g}}^a) / B_{g\bar{g}}^a, \forall g \in \mathcal{G}$.

cost B_g^a that group g imposes upon itself through intra-group contacts. The ratio thus captures an intuitive “bang-for-the-buck” metric for when a group should be active: it should generate a high ratio of net economic value per total death cost incurred. Notably, these ratios determine both *which* group-activity pairs to prioritize and *how active* to allow these to be.

Part (ii) of Proposition 2 shows that the optimal dual-targeted policy changes in natural ways with important problem parameters. Group g 's activity level increases with its intrinsic economic value r_g^a , but decreases with the cost of death χ , with the transmission rate β (and thus the basic reproduction number R_0), with any group's disease progression parameters μ_h, p_h^D or the economic value lost due to deceased individuals v_h^{lifc} , with the number of susceptible and infected individuals (in either group) at the start of the period, with all social contact values in that activity, and with the economic value generated by the other group r_g^a . The latter result is driven by the same substitution effect evident in the numerator of ratios driving the optimal policy: a group is more confined when another group with high economic value is active in a certain activity, in order to avoid creating extra costly infections (for both groups).

Similar results can be readily shown for each of the AGE and ACT policies (see Proposition EC.2 and Proposition EC.3, respectively). Our numerical study in Section 6 will confirm that these key qualitative insights also hold in a more complex model with several age groups and activities, and with many periods when confinement is possible.

4.3. Gains from Targeting

We next quantify the magnitude of gains from targeting and characterize how these depend on important problem parameters. For this analysis, we assume that $D_B^a > 0$ for each $a \in \mathcal{A}$, which is necessary and sufficient to ensure that problem (6) is a convex optimization problem, allowing us to fully characterize the optimal solution. The condition $D_B^a > 0$ is equivalent to:

$$B_1^a B_2^a > (B_{12}^a)^2 \Leftrightarrow 4A_1 A_2 C_{11}^a C_{22}^a \frac{I_1(0)}{N_1(0)} \frac{I_2(0)}{N_2(0)} > \left(A_1 C_{12}^a \frac{I_2(0)}{N_2(0)} + A_2 C_{21}^a \frac{I_1(0)}{N_1(0)} \right)^2, \quad (8)$$

which is an inequality with a natural interpretation: it requires the intra-group transmission costs to be “substantially larger” (in terms of their geometric mean) than the inter-group transmission

cost. The condition is likely to hold provided that the contact matrix C^a is strongly diagonally dominant, which is a standard assortativity property that contact matrices documented in the literature routinely satisfy (Mossong et al. 2008).

We next analyze how gains from targeting depend on the policy-relevant parameter χ and on the transmission parameter β . We consider β because it follows a strictly monotonic relationship with the basic reproduction number R_0 , arguably one of the most policy-relevant pandemic parameters.

PROPOSITION 3 (Conditions for Gains from Targeting). *For each optimal policy $\pi \in \{\text{NT}, \text{AGE}, \text{ACT}, \text{AA}\}$, there exist thresholds $\chi_{\min}^{\pi}(\beta) \leq \chi_{\max}^{\pi}(\beta)$ (with explicit closed-form expressions) such that:*

- (i) *There are zero gains from targeting if $0 \leq \chi \leq \chi_{\min}^{\pi}(\beta)$ or if $\chi \geq \chi_{\max}^{\pi}(\beta)$.*
- (ii) *$\chi_{\min}^{\pi}(\beta)$ and $\chi_{\max}^{\pi}(\beta)$ are piecewise linear, increasing functions of $1/\beta$.*
- (iii) *The four thresholds satisfy the inequalities:*

$$\chi_{\min}^{\text{AA}} \leq \chi_{\min}^{\text{AGE}}, \chi_{\min}^{\text{ACT}} \leq \chi_{\min}^{\text{NT}} \quad \text{and} \quad \chi_{\max}^{\text{AA}} \geq \chi_{\max}^{\text{AGE}}, \chi_{\max}^{\text{ACT}} \geq \chi_{\max}^{\text{NT}}.$$

- (iv) *For $\chi \in (\chi_{\min}^{\text{NT}}, \chi_{\max}^{\text{NT}})$, the gains of AGE, ACT, AA relative to NT are non-monotonic in problem parameters such as $\{r_g^a\}_{g \in \mathcal{G}, a \in \mathcal{A}}, \{\mu_g, p_g^D\}_{g \in \mathcal{G}}$.*

Proposition EC.4 in the Appendix restates this result more explicitly, by including expressions for all the thresholds and connecting them to other problem parameters.

Proposition 3 implies that targeting only brings benefits inside a certain L-shaped range of the (β, χ) parameter space, when neither the cost of death χ nor the transmission rate β are excessively low or excessively high. When χ is very low (specifically, $\chi \leq \chi_{\min}^{\text{AA}}(\beta)$), the objective heavily weighs economic value, so every policy would leave all groups fully open in all activities, $u_g^a = 1$; conversely, when χ is very large ($\chi \geq \chi_{\max}^{\text{AA}}(\beta)$), all policies set minimum activity levels $u_g^a = \varepsilon$ to reduce the large penalty for deaths. In both cases, there are no gains from targeting. The same outcomes arise with a fixed χ if the transmission rate β becomes extreme: when β is very small—so essentially no transmission occurs—all policies set $u_g^a = 1$, whereas when the transmission rate is extremely

high, all policies set minimal activity levels, $u_g^a = \varepsilon$, to prevent an excessive number of costly infections. The region wherein gains from targeting arise is L-shaped because both its lower and upper envelopes decrease with β (by (ii), both decrease as $1/\beta$, in piecewise fashion).

Intuitively, part (iii) also emphasizes that the boundaries of this region are determined by the AGE-ACT policy: as χ rises just above $\chi_{\min}^{\text{AA}}(\beta)$, the AGE-ACT policy is the first to be able to generate gains, followed by AGE and ACT, and as χ is immediately below $\chi_{\max}^{\text{AA}}(\beta)$, AGE-ACT is the only policy that is still able to generate gains.

Lastly, part (iv) suggests that the gains from targeting—when positive—follow a complex, non-monotonic pattern with respect to problem parameters.

A similar analysis can be conducted with respect to other important disease parameters (such as p_g^D) together with χ , and in all such cases the parameter regime wherein gains from targeting exist can be shown to be L-shaped. These insights appear to be robust when we quantify the benefits of targeting under our more complex model in Section 6.3, for both β (equivalently, R_0) and p_g^D .

Proposition 3 does not provide any ordering or guidance regarding a choice between AGE and ACT. The following result gives a more direct comparison.

PROPOSITION 4 (AGE versus ACT). *Based on thresholds from Proposition 3, let $H^{\text{ACT-AGE}} := (\chi_{\max}^{\text{ACT}}, \chi_{\max}^{\text{AGE}})$ and $H^{\text{AGE-ACT}} := (\chi_{\max}^{\text{AGE}}, \chi_{\max}^{\text{ACT}})$. Also, with $\kappa_g := p_g^D \mu_g S_g(0)$, define:*

$$\begin{aligned} L_{g,a} &= \beta \Delta_0 \Delta_1 \left[\kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + \kappa_{\bar{g}} \cdot C_{g\bar{g}}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \\ K_{g,a} &= \beta \Delta_0 \Delta_1 \left[v_g^{\text{life}} \cdot \kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + v_{\bar{g}}^{\text{life}} \cdot \kappa_{\bar{g}} \cdot C_{g\bar{g}}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \end{aligned} \quad (9)$$

If

$$\exists a \in \mathcal{A} : \frac{\sum_g r_g^a}{\sum_g L_{g,a}} > \max_g \frac{\sum_a r_g^a}{\sum_a L_{g,a}} \quad \text{and} \quad \exists g \in \mathcal{G} : \frac{\sum_a K_g^a}{\sum_a L_{g,a}} < \min_{a \in \mathcal{A}} \frac{\sum_g K_{g,a}}{\sum_g L_{g,a}}, \quad (10)$$

then there exist positive thresholds $\beta^1 \leq \beta^2$ (possibly identical) such that:

- (i) if $\beta < \beta^1$, $H^{\text{ACT-AGE}} \neq \emptyset$ and ACT has larger optimal objective than AGE if $\chi \in H^{\text{ACT-AGE}}$;
- (ii) if $\beta > \beta^2$, $H^{\text{AGE-ACT}} \neq \emptyset$ and AGE has larger optimal objective than ACT if $\chi \in H^{\text{AGE-ACT}}$.

Moreover, if both inequalities in (10) are switched, (i) and (ii) hold with ACT switched with AGE.

The result states sufficient conditions under which ACT dominates AGE or vice-versa, per parts (i) and (ii), respectively. To gain intuition for the conditions in (10), consider first the factors $L_{g,a}$ and $K_{g,a}$. The relationship between these factors and the disease burden is quite straightforward: $B_g^a + B_{g,\bar{g}}^a = L_{g,a}\chi + K_{g,a}$. Therefore, $\frac{1}{2}L_{g,a}$ exactly accounts for the expected deaths created by an individual in group g through his social contacts in activity a ; this includes the deaths created within his own group (the term involving C_{gg}^a) and the average of the deaths created in groups g and \bar{g} through inter-group contacts (the terms with $C_{g\bar{g}}^a$ and $C_{\bar{g}g}^a$). The term $K_{g,a}$ quantifies the economic loss associated with these deaths due to lost lifetime economic value (from v_g^{life} and $v_{\bar{g}}^{\text{life}}$). The first inequality in (10) thus requires the existence of an activity a whose “average risk-adjusted intrinsic value”—the ratio of its total intrinsic economic value $\sum_g r_g^a$ prorated by the total expected deaths it creates when both groups are fully active in it, $\sum_g L_{g,a}$ —is strictly larger than the risk-adjusted intrinsic value of either group $g \in \mathcal{G}$. In contrast, the second inequality in (10) requires the existence of a group g whose “average economic loss from death”—the ratio of total economic losses from deceased group members, $\sum_a K_{g,a}$, prorated by the total expected deaths the group creates in all activities it engages in, $\sum_a L_{g,a}$ —is strictly smaller than that of either activity.

EXAMPLE 1. It is illuminating to consider an example inspired by our case study based on Ile-de-France data (Section 6.1). Group $g = “ya”$ corresponds to the young/adult population and group $g = “e”$ corresponds to an elderly population, in retirement; activity $a = “w”$ is work and activity $a = “o”$ summarizes other activities (leisure, transport, etc.). The data predictably suggest that the intrinsic economic value in work is much higher than that of other activities ($\sum_g r_g^w \gg \sum_g r_g^o$) and the lost economic output due to a death in the young/adult population is higher than that in the elderly population ($v_{ya}^{\text{life}} \gg v_e^{\text{life}}$). The elderly have higher mortality risk $p_e^D \gg p_g^D$. Lastly, the majority of social interactions occur in other activities, and at work among the young.⁷

With this example, one can verify that:

⁷ The social contact matrices satisfy $C_{ya,ya}^o > C_{e,ya}^o > C_{e,e}^o \approx C_{ya,ya}^w > C_{ya,e}^o \approx C_{e,ya}^w \approx C_{ya,e}^w \gg C_{e,e}^w$. To give a sense of relative magnitudes, a $>$ corresponds to doubling the contacts, and \gg corresponds to ten times the contacts.

1. The highest number of expected deaths (when fully open) would be incurred in other activities, for elderly and then the young, followed by work, in the order $L_{e,o} > L_{ya,o} \gg L_{e,w} \approx L_{ya,w}$.
2. The highest lifetime economic loss from deaths is incurred from the young/adult group engaging in other activities and then in work, following by the elderly engaging in other activities and work ($K_{ya,o} > K_{ya,w} > K_{e,o} > K_{e,w}$). This is driven largely by the large gap in $v_{ya}^{\text{life}} \gg v_e^{\text{life}}$ and by the large number of social contacts occurring in other activities.

Then, the work activity “ w ” and the elderly group “ e ” satisfy the inequalities in (10). Work satisfies the first inequality in (10) both because its intrinsic value is high (large numerator $\sum_g r_g^w$) and because the social contacts in work are small relative to other activities (small denominator $\sum_g L_{g,a}$). The elderly group $g = “e”$ satisfies the second inequality in (10) because of the disproportionately large number of deaths that it incurs when active (large denominator $L_{e,w} + L_{e,o}$) and due to the smaller loss in economic output from deaths (small numerator $K_{e,w} + K_{e,o}$).

The example highlights the parameter regimes that make the inequalities hold. Heterogeneity in the intrinsic economic value or the social contacts of activities would make the first inequality in (10) more likely to hold. If an activity with high intrinsic value also has fewer social contacts, as is expected from work, that is even more likely. Heterogeneity in the mortality risk of groups, in their lifetime economic value, or in the social contacts they generate make the second condition more likely to hold. If a particular group happens to have both a lower economic lifetime value and a high number of expected deaths, as is expected from the elderly, the condition very likely holds.

To appreciate why—under these conditions—ACT has better performance than AGE at a low β and high χ regime, whereas AGE has better performance at a high β and low χ regime, consider again our example. When the basic reproduction number R_0 is low (low β) and the policymaker cares enough about the health burden to consider confinements but not to completely confine everything (i.e., χ large but not excessive), then ACT is better than AGE: the policymaker could keep work open and confine other activities, which would maintain high economic output while also reducing deaths substantially (as $L_{e,o} > L_{ya,o} \gg L_{e,w} \approx L_{ya,w}$). In contrast, AGE would be forced

to confine both groups because the same inequality implies that confining only one group would still generate many deaths. When the reproduction is very large (β large), the number of infections and deaths is very large, and matter even for a policymaker who prioritizes economic value (i.e., low χ) due to the lifetime economic cost associated with deaths. In this regime, AGE would actually perform better: the inequality $K_{ya,o} > K_{ya,w} > K_{e,o} > K_{e,w}$ suggests that by confining the *young/adult group*, the policymaker would be able to suppress enough deaths and also preserve enough economic value in the process, whereas with ACT, they would have no alternative but to fully confine all activities.

The discussion and the example also highlight the ways in which an AGE-ACT policy improves upon both AGE and ACT. In fact, AGE-ACT can improve in *both regimes*: at low β and high χ , AGE-ACT could confine the elderly group in other activities and only if needed, confine the young group in other activities (and not necessarily to the same level). At high β and low χ , AGE-ACT could confine the young/adult group in other activities, followed by the young group in work, if need be; interestingly, note that the AGE-ACT policy would *not* target the elderly group with confinements in this case.

Our results based on the stylized model suggest a few important insights and implications. Targeted confinements follow intuitive rules—prioritizing groups with high ratios of intrinsic economic value to mortality cost—but the proper calculation of such ratios should take into account the externalities that groups create upon each other through cross-group infections. Gains from targeting only arise in specific parameter regimes and the choice of whether to target activities or demographic groups is subtle, governed by important disease parameters and by the policymaker’s prioritization between economic and mortality losses. As targeting may also entail implementation costs, it is critical to accurately quantify the magnitude of gains with more realistic models that capture long-term disease progression and the policymaker’s ability to adjust decisions over time.

5. Algorithm: Re-Optimization with Linearized Dynamics

With this motivation, we next focus on designing algorithms for finding optimal targeted policies for our general model in Section 3 and quantifying the gains of targeting. Solving problem (5) to

optimality, however, requires solving intractable problems. This arises because the key term in the dynamics of any SEIR-type model is the rate of new infections, which involves multiplying the susceptible state with the infected state. This introduces nonlinearity in the state dynamics; for instance, the change in the susceptible population in group g in our model from (EC.2) is:

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

Even with a social mixing parameter value $\alpha = 1$, like in our stylized model, expanding the expression of $S_g(t)$ for multiple periods produces a complex, non-convex dependency on past decisions $\ell(\tau)$ for $0 \leq \tau \leq t-1$, which makes the resulting problem intractable via convex optimization.

We therefore 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. Importantly, although we describe the ideas in the context of our simplified SEIR model, the ROLD framework can be readily generalized to deal with more complex, real-world settings that involve parameter uncertainty (see Section 7.1), complex SEIR models with more age groups or limited resources (such as hospital or ICU capacity), or limited compliance with the recommended lockdowns.

5.1. Linearization and Optimization

The key idea is to solve the problem in a shrinking-horizon fashion, where at each time step $k = 0, \dots, T$, we linearize the system dynamics and objective (over the remaining horizon), determine optimal decisions for all k, \dots, 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 (2). The typical approach in control theory 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 $\hat{\mathbf{u}}_{k:T-1}$ and let $\hat{\mathbf{X}}_{k:T}$ denote the resulting nominal system trajectory under the true dynamic (2) and under $\hat{\mathbf{u}}_{k:T-1}$. We approximate the original dynamics through a Taylor expansion around $(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t)$:

$$\frac{\Delta \mathbf{X}_t}{\Delta t} \approx \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) + \nabla_{\mathbf{X}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t)(\mathbf{X}_t - \hat{\mathbf{X}}_t) + \nabla_{\mathbf{u}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t)(\mathbf{u}_t - \hat{\mathbf{u}}_t), \quad (11)$$

where $\nabla_{\mathbf{X}} \mathbf{F}_t$ and $\nabla_{\mathbf{u}} \mathbf{F}_t$ denote the Jacobians with respect to \mathbf{X}_t and \mathbf{u}_t , respectively. Note that these Jacobians are evaluated at points on the nominal trajectory, so (11) is indeed a linear expression of \mathbf{X}_t and \mathbf{u}_t . By induction, every state \mathbf{X}_t under dynamic (11) will be a linear function of \mathbf{u}_τ for $\tau < t$, and all the constraints will also depend linearly on the decisions.

In a similar fashion, we linearize the objective (4). Since $v_g(\ell(t))$ is linear in \mathbf{u}_t for all $t = 0, \dots, T-1$, the objective contains *bilinear* terms and can be written compactly as:

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

for some matrix \mathbf{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.4). By linearizing this using a Taylor approximation, we consider the following objective instead:

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

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

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

1. Given the current state \mathbf{X}_k and a nominal control sequence $\hat{\mathbf{u}}_{k:T-1}^{(k)}$ for all remaining periods, calculate a nominal system trajectory $\hat{\mathbf{X}}_{k:T}$ under the true dynamic in (2). (The nominal control sequence is obtained by a random seeding procedure at $k = 0$, and from 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 $\hat{\mathbf{X}}_t$ and use (13) to approximate the objective-to-go function over the remaining periods $t \in \{k, \dots, T\}$.
3. Solve the linear program to obtain decision variables $\mathbf{u}_{k:T-1}^{*,(k)}$ 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 $\hat{\mathbf{u}}_{k+1:T-1}^{(k+1)} := \mathbf{u}_{k+1:T-1}^{*,(k)}$.
5. Update the states using the optimal control $\mathbf{u}_k^* := \mathbf{u}_k^{*,(k)}$ and the true dynamic in (2), i.e.

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \mathbf{F}_k(\mathbf{X}_k, \mathbf{u}_k^*).$$

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

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 $\mathbf{u}_{k:T-1}^{*(k)}$ determined in Step 3 diverge sufficiently from the nominal controls $\hat{\mathbf{u}}_{k:T-1}^{(k)}$ considered in the linearization in Step 2. This can lead to 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 $\hat{\mathbf{u}}_{k:T-1}^{(k)}$, 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.

6. The Gains from Targeting

We next apply our framework to quantify the magnitude of gains from targeting. We focus the discussion here on the context of COVID-19 in the Île-de-France region of France, but Appendix EC.7.1 presents results for two other countries with very different demographic profiles (Hong Kong and South Africa) and for other pandemics with different dynamics (Ebola and Influenza), showing that all the insights are robust.

6.1. Experimental Setup

We adopt values for disease progression parameters for Île-de-France directly from the study by Salje et al. (2020), and values for the number of social contacts for different age groups and

activities from the study of Béraud et al. (2015) focused on France, using the tool by Wille et al. (2020). 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. The values for all our model parameters are summarized in Tables EC.2–EC.5; Appendix EC.6 provides all the details for parameter specification.

We use an optimization horizon of $T' = 90$ days in the experiments reported in the main paper. The starting time of the horizon, $t = 0$, corresponds to March 17 2020, the day when lockdowns for the COVID-19 pandemic started in France. We obtain the initial configuration at $t = 0$ by considering the date of patient zero to be December 20 2019.⁸ We allow confinement decisions to change every two weeks. To optimize confinements, we set the total time horizon to be $T = T' + 41$ and further constrain the policies to be fully open on days $T', \dots, T' + 41$, while allowing for infections beyond day T' . We do this to mitigate possible end-of-horizon effects: any deaths and loss of economic value between day $T' + 1$ and $T' + 41$ will still count towards the objective, so no policy can allow for too many infections towards the end of the T' -day optimization horizon. Section 8 discusses further techniques for mitigating end-of-horizon effects.

To quantify the benefits of targeting, we calculate the four ROLD policies of interest—NO-TARGET, AGE, ACT, and AGE-ACT (or simply ROLD when no confusion can arise)—in problem instances covering a wide range of values for χ , from 0 to $1000\times$ the annual GDP per capita in France.⁹ For each χ value and each policy, we record separately the economic losses and the number of deaths generated. To obtain each policy, we run suitably constrained versions of the ROLD optimization problem initialized at several starting points (details in Appendix EC.4.4).

⁸ We assume that the SEIR process starts with an infected individual of the 20-64 y.o. age group (Mohammad 2020).

⁹ 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 EURO using the exchange rate on June 17, 2020.

6.2. Gains from Targeting

Figure 2 records each policy’s performance in several problem instances that differ in the 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 0-12.4 (0%-93.1%) in AGE, by EUR 0-13.1B (0%-86.1%) in ACT, and by EUR 0-16.6B (0.1%-96.8%) in AGE-ACT. This Pareto-dominance is unexpected because it is not explicitly required in our optimization procedure. This underlines that any form of targeting can lead to significant improvements in terms of *both* health and economic outcomes. Importantly, although neither AGE nor ACT Pareto-dominate each other, AGE-ACT Pareto-dominates *all* other policies. These results, which are very robust (see Appendix EC.7.2 for many more problem instances), suggest that dual targeting has the potential to significantly improve *both* health and economic outcomes.

We also compare the ROLD policies with two extreme benchmarks, corresponding to a “full confinement” (FC) policy that sets all activity levels (except home) to 0, and a “fully open” (FO) policy that sets all activity levels to 1. These benchmarks can be expected to perform well when completely prioritizing a reduction in the number of deaths ($\chi \rightarrow \infty$) or economic losses ($\chi = 0$), respectively. ROLD AGE-ACT meets or exceeds the performance of these two extreme policies: for a sufficiently large χ , ROLD exactly recovers the FC policy, resulting in 9,983 deaths and economic losses of EUR 33.8B; for a sufficiently low χ , ROLD actually *Pareto-dominates* the FO policy, reducing the number of deaths by 4,606 (9.11%) and reducing economic losses by 16.6%. The latter result, which may seem surprising, is driven by the natural premise captured in our model that deaths generate economic loss because of lost lifetime economic value. 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. That targeting can generate such improvements is unexpected because the ROLD framework’s objective does not stipulate this.

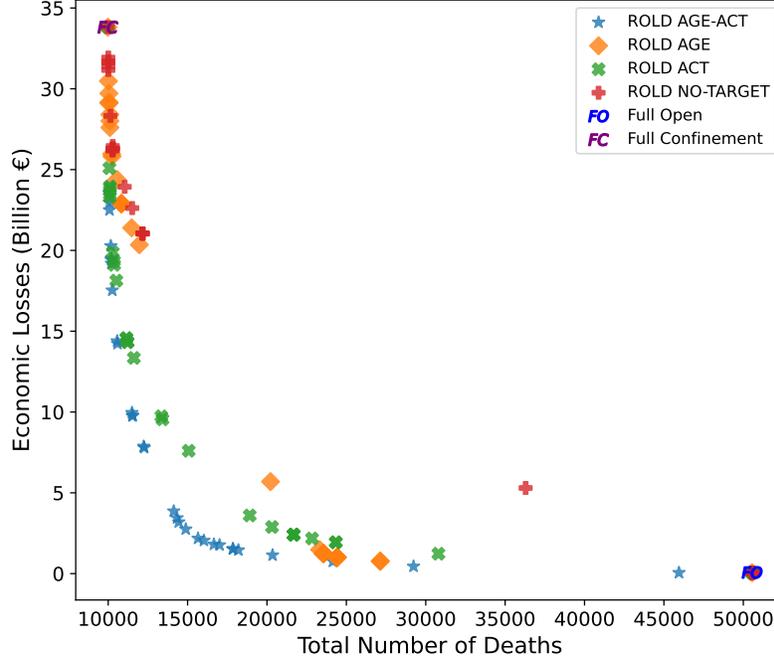


Figure 2 The total number of deaths and the economic losses generated by targeted ROLD policies. The plot also includes two important benchmarks: a “fully open” policy (FO) that sets all activity levels to 1, and a “full confinement” policy (FC) that sets all activity levels to 0. Each marker corresponds 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 we also include a very large value ($\chi = 10^{16}\times$, for which ROLD AGE-ACT recovers the full confinement policy FC).

To better understand how gains could arise from targeting, we next examine the structure of the optimal ROLD AGE-ACT decisions. Our main observation is that the policy maintains high activity levels for (age group, activity) pairs with high economic value and few social contacts. Specifically, the policy consistently raises activity levels in (age group, activity) pairs with a high econ-to-contacts-ratio, *defined* as the marginal economic value divided by the total social contacts generated by a group in an activity:

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

where $\text{Economic Value}(t) = \sum_{g \in \mathcal{G}} (v_g(\ell(t)) \cdot N_g(t))$. This ratio directly relates to our findings on the “bang-for-the-buck” ratio in Proposition 2 of Section 4, and confirms the key intuition that dual targeting raises activity levels based on ratios of economic value and potential death burden.

To understand how ROLD AGE-ACT generates *complementarities*, 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 20-64 y.o. group has the highest ratio in work, whereas the 0-19 y.o. and 65+ y.o. groups have the highest ratio in leisure. Accordingly, ROLD coordinates confinements to account for this (see Figure 3): the group 20-64 y.o. remains more open in work, but faces strict confinements in leisure, transport, and other; whereas groups 0-19 and 65+ y.o. remain more open in leisure, transport, and other, and group 0-19 y.o. is allowed some activity in work (but less than the work activity level of group 20-64 y.o.). 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.

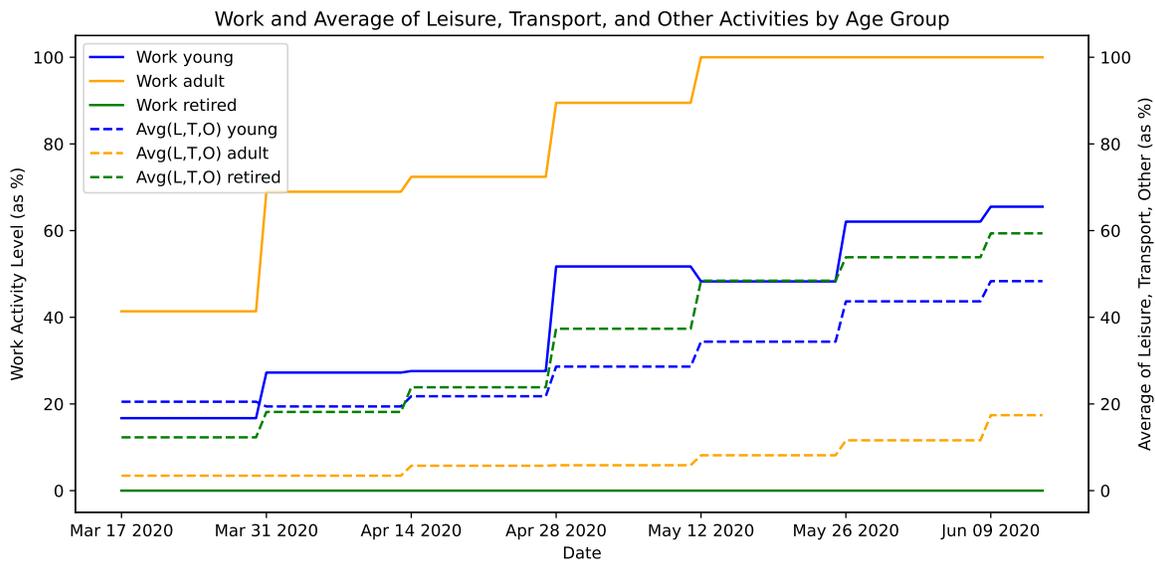


Figure 3 The activity levels of ROLD AGE-ACT policies for work, and for the average of leisure, transport, other, for each of the three age groups. The lines indicate mean activity levels, where for each time t the mean is taken across different values for the cost of death χ . (The figure is best viewed in color.)

6.3. The Dependence of Gains from Targeting on Key Parameters

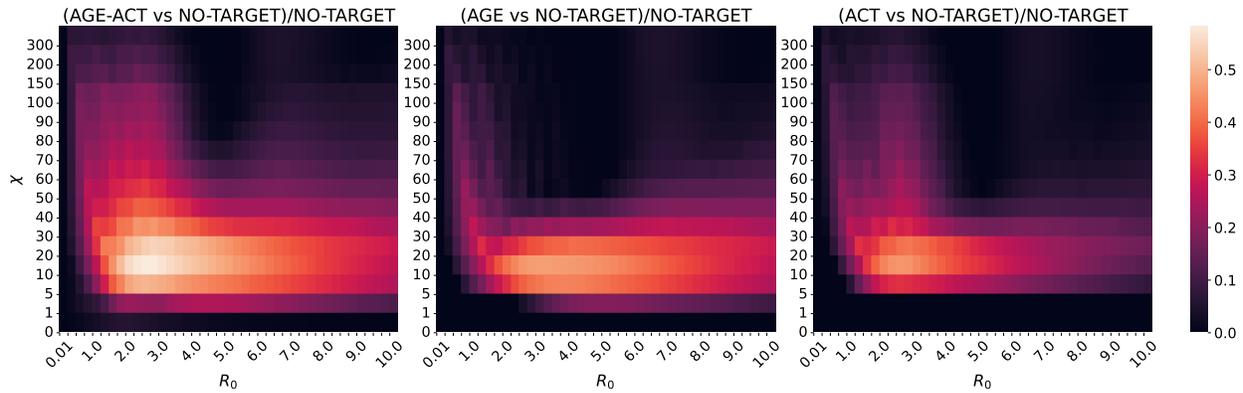
Next, we examine how the gains from targeting change with two key epidemiological parameters, the basic reproduction number R_0 and the probability of death conditioned on infection p_g^D , across different values of the cost of death χ .

Dependence on reproduction number. We vary R_0 starting from 0.01, 0.5 and then up to 10.0 in 0.5 increments. (For reference, the R_0 value in our COVID-19 study was 2.9, as defined in Appendix EC.6.) We also vary the key policy parameter χ from $0\times$ to $300\times$ and for each pair of (R_0, χ) values, we determine the optimal ROLD policies and calculate the gains of targeting relative to the optimal ROLD NO-TARGET policy.¹⁰ Figure 4a reports the gains of targeting, and Figure 4b reports the average activity levels for all ROLD variants.

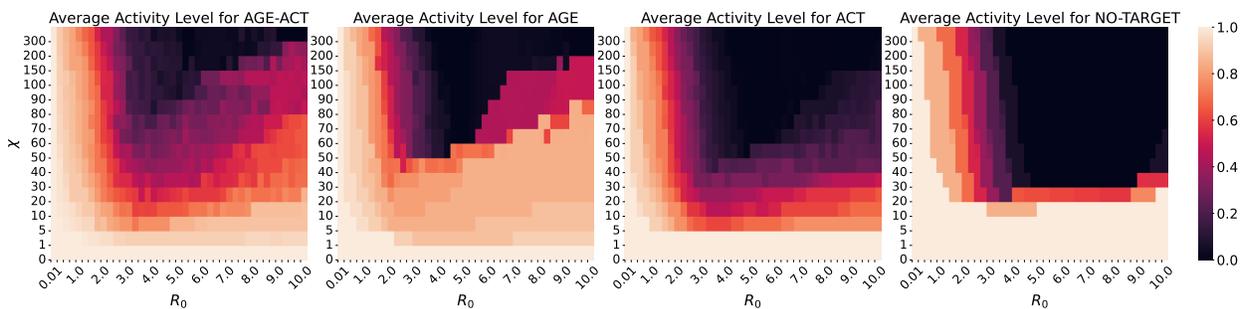
We focus on the gains of ROLD AGE-ACT with respect to ROLD NO-TARGET. Consistent with the predictions from Proposition 3 in Section 4, we see that ROLD AGE-ACT attains significant gains in an L-shaped region of the (R_0, χ) parameter space, wherein gains can exceed 50%. We distinguish three regimes depending on χ . For low χ , both AGE-ACT and NO-TARGET tend to remain open, and thus gains from targeting are zero for any R_0 . Similarly, for χ values exceeding a certain (R_0 -dependent) threshold, both AGE-ACT and NO-TARGET close down all activities, and thus gains from targeting are again zero. Both regimes replicate the predictions in Section 4, and can be seen in the average activity heatmaps in Figure 4b.

The dependence is more nuanced for medium χ (from $30\times$ to $100\times$), where two peaks emerge as R_0 increases from 0.01 to 10. For R_0 in the 0.5 to 4.0 range, both AGE-ACT and NO-TARGET progressively close down activity as R_0 increases. The advantage of AGE-ACT is that it can leverage finer targeting, which allows it to maintain a higher activity level than NO-TARGET—as evident from the more gradual reduction of activity shown in Figure 4b—which leads to higher economic value. However, as R_0 exceeds 4.5, NO-TARGET maintains full confinement, whereas AGE-ACT begins to gradually raise activity levels thanks to finer targeting, which leads to another peak in gains. At such high values of R_0 , one enters a herd immunity scenario where infections and

¹⁰ That is, one minus the ratio of total loss with a targeted policy divided by total loss with ROLD NO-TARGET.



(a) Gains of ROLD AGE-ACT, AGE, and ACT over NO-TARGET



(b) Average activity levels of ROLD AGE-ACT, AGE, ACT, and NO-TARGET

Figure 4 Heatmaps showing outcomes under optimal ROLD policies as a function of the basic reproduction number R_0 and the cost of death χ . Figure (a) provides heatmaps for the relative gains of targeting over NO-TARGET. Figure (b) provides heatmaps for the average activity levels of the ROLD variants. (The figure is best viewed in color.)

deaths are somewhat unavoidable¹¹, so all ROLD policies seems to exhibit a switch in strategy from full confinement to policies that enforce confinements over fewer periods of time and allow more activity subsequently (as evidenced in the pattern of average activity levels that increase with R_0 , seen in Figure 4b). In this regime, targeting is again beneficial: AGE-ACT can confine groups and activities differentially, which raises economic value without substantially increasing deaths. In contrast, NO-TARGET is forced to maintain significantly stricter confinements to control the excessive number of deaths. (This is best seen in Figure EC.3 in the Appendix, which plots separate heatmaps for economic losses and deaths.)

¹¹ Because we do not restrict the home activity, the entire population becomes infected for very large R_0 .

The complex dependency of gains (specifically, the two peaks emerging due to herd immunity) highlights the difficulty of making accurate predictions in realistic settings based solely on a stylized model, and underscores the need for a structured optimization approach to drive policy.

We also examine the relative gains of targeting based solely on age groups or on activities. Figure 4a shows the gains of AGE and ACT over NO-TARGET. The plots exhibit qualitatively similar behavior to that of AGE-ACT, so our prior comments directly apply, albeit in a more narrow region of the parameter space. Importantly, when comparing AGE and ACT, it is noticeable that ACT tends to deliver larger gains for pairs with low R_0 and high χ values, whereas AGE delivers larger gains at pairs with high R_0 and lower χ values. This behavior, which was predicted in Proposition 4 based on the stylized model, becomes even more evident in Figure 5, which shows the relative gains of AGE over ACT. Note that there is a region of R_0 between 1 and 4.0, with moderate-to-high values of χ , where ACT dominates over AGE. Conversely, AGE delivers larger gains than ACT in the regime of high R_0 and low-to-moderate χ . Dual targeting leverages the advantages of both AGE and ACT, which enlarges the L-shaped region of gains.

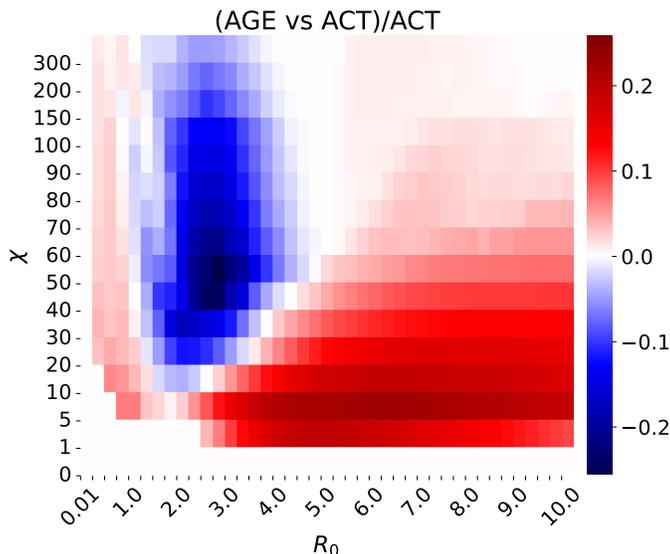
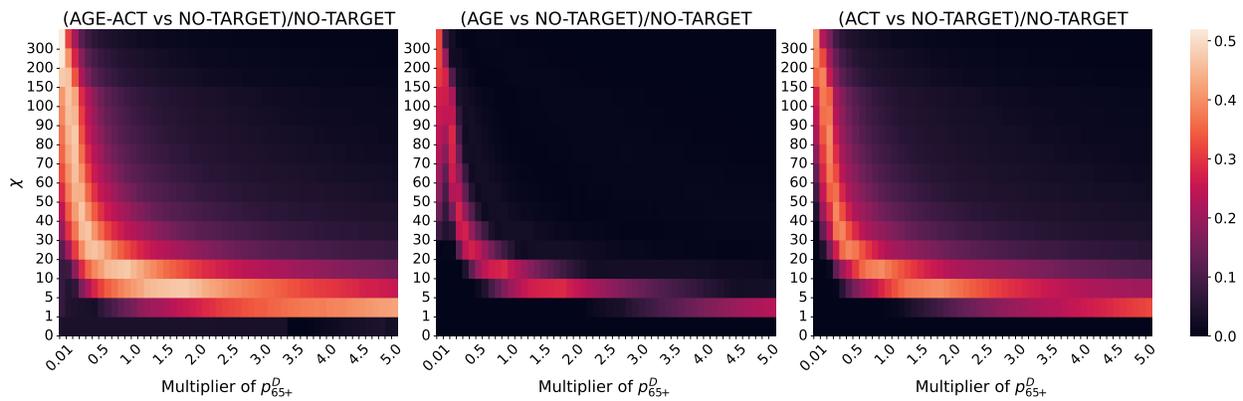
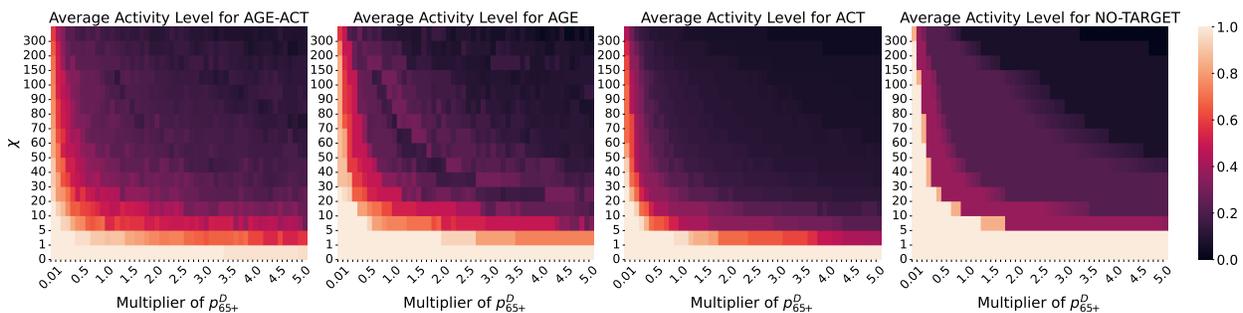


Figure 5 Heatmap with the relative gains of ROLD AGE vs. ROLD ACT, as a function of the reproduction number R_0 and the cost of death χ . Large positive (negative) values indicate that AGE (respectively, ACT) dominates. (The figure is best viewed in color.)

Dependence on probability of death multiplier. We next examine how targeting gains depend on *group-specific* parameters. We focus on the impact of the mortality risk for the 65+ y.o. age group, which we change by multiplying the probability of death given infection p_g^D for that group with a factor taking values from 0.01, 0.05, 0.1, and then up to 5.0 in 0.1 increments. We keep the probabilities of death for the other two groups fixed. The results are depicted in Figure 6a, which plots the gains relative to NO-TARGET, and in Figure 6b, which plots the average activity levels for all ROLD variants.



(a) Gains of ROLD AGE-ACT, AGE, and ACT over NO-TARGET



(b) Average activity levels of ROLD AGE-ACT, AGE, ACT, and NO-TARGET

Figure 6 Heatmaps showing outcomes under optimal ROLD policies, as a function of the multiplier of the probability of death p_g^D for the 65+ y.o. age group and the cost of death χ . Figure (a) provides heatmaps for the relative gains of targeting over NO-TARGET. Figure (b) provides heatmaps for the average activity levels of the ROLD variants. (The figure is best viewed in color.)

The L-shaped parameter region where gains exist is again readily apparent, and gains can be as high as 50% with AGE-ACT. For low χ , all policies maintain full activity levels regardless of

the mortality risk (p_g^D), so gains from targeting are negligible. At intermediate χ values, as the mortality risk rises, all policies lower the activity levels, but to different degrees: NO-TARGET must lower them significantly, to control the large number of deaths in the elderly group, whereas AGE-ACT can differentially confine groups, and allow a higher average activity level, which creates gains. Those gains vanish as mortality risk (p_g^D) increases further and all policies tend to enforce complete confinement. The main difference compared to our earlier study focused on R_0 is that no second peak in gains emerges. Regardless of χ , a higher mortality risk would make every policy reduce its activity levels, so no switch towards herd immunity occurs, and the relative gains simply vanish as mortality risk rises sufficiently. This also makes χ and p_g^D rough substitutes in terms of how they affect outcomes and gains: both χ and p_g^D ultimately scale the overall cost of deaths, so gains decrease roughly with the product of χ and p_g^D , consistent with the predictions based on the stylized model in Section 4.

7. Dealing with Parameter Uncertainty

At the time when policymakers are forced to make important confinement decisions, they may still face uncertainty regarding the value of important problem parameters such as the basic reproduction number R_0 or the mortality rates for different age groups—and more so at the onset of a new pandemic. However, some of this uncertainty could be resolved with time, as more information becomes available. In this section, we provide an extension of our framework to accommodate this situation and we investigate how gains from targeting depend on uncertainty.

7.1. Robust Formulation

To capture this realistic situation, let $\theta \in \mathbb{R}^m$ denote the vector of all parameters with unknown values at time $t = 0$; we model θ as a random vector following an unknown discrete probability distribution. Specifically, we consider a discrete set of scenarios $\Omega = \{1, 2, \dots, |\Omega|\}$ indexed by ω and assume that θ is supported on the set $\Theta = \{\theta^\omega : \omega \in \Omega\} \subset \mathbb{R}^m$, taking value θ^ω with probability p_ω , for $\omega \in \Omega$. We let $\mathbf{p} := (p_1, p_2, \dots, p_{|\Omega|})^\top$ denote the probability distribution of θ . To model

ambiguity, we assume that \mathbf{p} is only known to reside in an ambiguity set of potential distributions \mathcal{P} , which is given as the convex hull of n known extreme beliefs $\mathbb{P}^1, \mathbb{P}^2, \dots, \mathbb{P}^n$:

$$\mathbf{p} \in \mathcal{P} := \text{conv}(\{\mathbb{P}^1, \mathbb{P}^2, \dots, \mathbb{P}^n\}), \quad \text{where } \mathbb{P}^s \in \left\{ \mathbf{p} \in \mathbb{R}^{|\Omega|} : \mathbf{p} \geq 0, \sum_{\omega \in \Omega} p_\omega = 1 \right\}, \quad \forall s \in \{1, \dots, n\}.$$

Our robust model allows for limited learning while remaining tractable. Although the value of $\boldsymbol{\theta}$ is unknown at $t = 0$, we assume that it becomes known at the start of a future period $T^{\text{new}} \in \{1, \dots, T - 1\}$, and the policymaker can adjust their subsequent decisions after learning the value. That is, when $\boldsymbol{\theta} = \boldsymbol{\theta}^\omega$, the control during periods $t \in \{T^{\text{new}}, T^{\text{new}} + 1, \dots, T - 1\}$, notated by \mathbf{u}_t^ω , can depend on the realized scenario ω . To formalize this, we allow the decisions to depend on the scenario ω for *every* period $t \in \{0, \dots, T - 1\}$, but we impose non-anticipativity constraints to ensure that the policymaker only uses information that is available at each time:

$$\mathbf{u}_t^\omega = \mathbf{u}_t^{\omega'}, \quad \forall t \in \{0, 1, \dots, T^{\text{new}} - 1\}, \quad \forall \omega, \omega' \in \Omega. \quad (14)$$

The constraints ensure that the same control policy is followed up until the time T^{new} when new information is learned, while subsequent decisions are allowed to depend on the learned information.

Correspondingly, we revisit our earlier notation to show the explicit dependency on parameters $\boldsymbol{\theta}$. In scenario $\omega \in \Omega$, we use $\mathbf{u}_{0:T-1}^\omega$ to denote all decisions, we rewrite the state dynamics as:

$$\mathbf{X}_{t+1}^\omega = \mathbf{X}_t^\omega + \mathbf{F}_t(\mathbf{X}_t^\omega, \mathbf{u}_t^\omega, \boldsymbol{\theta}^\omega), \quad \forall 0 \leq t \leq T - 1, \quad \forall \omega \in \Omega, \quad (15)$$

and use $\text{Economic Loss}(\mathbf{u}_{0:T-1}^\omega, \boldsymbol{\theta}^\omega)$, $\text{Total Deaths}(\mathbf{u}_{0:T-1}^\omega, \boldsymbol{\theta}^\omega)$, and $J(\mathbf{u}_{0:T-1}^\omega, \boldsymbol{\theta}^\omega)$ to denote the resulting economic losses, total deaths, and total loss objective, respectively.

When deciding confinements in period $t = 0$, the policymaker seeks to minimize the worst-case expected total loss. Formally, the policymaker's optimization problem is:

$$\begin{aligned} & \underset{\{\mathbf{u}_{0:T-1}^\omega\}_{\omega \in \Omega}}{\text{minimize}} && \max_{\mathbf{p} \in \mathcal{P}} \sum_{\omega \in \Omega} p_\omega \cdot J(\mathbf{u}_{0:T-1}^\omega, \boldsymbol{\theta}^\omega) \\ & \text{s.t.} && \mathbf{X}_{t+1}^\omega = \mathbf{X}_t^\omega + \mathbf{F}_t(\mathbf{X}_t^\omega, \mathbf{u}_t^\omega, \boldsymbol{\theta}^\omega) \quad \forall t \in \{0, \dots, T - 1\}, \forall \omega \in \Omega \\ & && \mathbf{u}_{0:T-1}^\omega \in [0, 1]^{T|\mathcal{G}||\mathcal{U}|} \quad \forall \omega \in \Omega \\ & && \mathbf{u}_t^\omega = \mathbf{u}_t^{\omega'} \quad \forall t \in \{0, 1, \dots, T^{\text{new}} - 1\}, \forall \omega, \omega' \in \Omega. \end{aligned} \quad (16)$$

Discussion of modeling assumptions. Modeling the support Θ and ambiguity set \mathcal{P} allows us to capture varying degrees of information and attitudes towards uncertainty that policymakers may have. In practice, Θ could contain a few scenarios, for instance, a central, representative scenario and some important deviations derived from the confidence intervals obtained when estimating the parameters. The ambiguity set \mathcal{P} models the accuracy of the information, but also the policymaker’s attitude towards ambiguity. With $n = 1$, the policymaker is risk-neutral. With $n > 1$, the policymaker is ambiguity averse, and a higher n increases the aversion towards ambiguity and makes the model more robust. As an extreme example, taking $n = |\Omega|$ and \mathbb{P}^s as unit mass on scenario θ^s (i.e., 1 in position s and 0 otherwise) would be akin to considering a robust optimization model that minimizes the worst-case losses calculated over all values $\theta \in \Theta$.

Robust ROLD. We can readily extend the ROLD algorithm described in Section 5 to the robust problem in (16). First, we linearize the states under each scenario $\omega \in \Omega$, i.e., we use a nominal control sequence $\hat{\mathbf{u}}_{0:T-1}^\omega$ and its associated nominal system trajectory $\hat{\mathbf{X}}_{0:T}^\omega$ to linearize the states $\mathbf{X}_{0:T}^\omega$ corresponding to scenario ω —with evolution described by (15)—around the nominal trajectories. With this, the problem of finding a non-anticipative control policy that minimizes the worst-case expected losses can be rewritten as a large-scale linear program. For further algorithmic details, we refer the reader to Appendix EC.5.

7.2. Gains of Robust ROLD Under Ambiguity

A natural conjecture to make is that a policymaker faced with ambiguity might prefer stricter confinements and that the benefits of targeting would *decrease* with *increased* ambiguity. To examine such questions, we apply our robust ROLD algorithm to the Île-de-France case study and consider two sets of experiments: one where ambiguity affects the reproduction factor R_0 and one where it affects the multiplier of the probability of death p_g^D for the elderly group. Focusing on these two parameters makes sense because they are key drivers of pandemic progression and more likely to be pandemic-specific and thus unknown at the start of a new pandemic. In contrast, other parameters

such as social contacts or economic values are less likely to change with a new pandemic, so ambiguity seems less pressing in their case. (Our robust framework readily allows modeling ambiguity over any problem parameters, should that be meaningful.)

Experimental setup. We describe the robust formulation with ambiguity in R_0 ; the one for the p_g^D multiplier is similarly set up. We consider values for R_0 as in Section 6.3, namely $\{0.01, 0.5, 1, \dots, 9.5, 10\}$. For each experiment, we consider $|\Omega| = 3$ scenarios: a focal R_0 scenario plus the two adjacent R_0 values from the range above. For simplicity we refer to these as the L (lower R_0), M (medium/focal R_0), and H (higher R_0) scenarios. We consider $n = 3$ extreme beliefs over these scenarios, where each belief is a three-point distribution over the three scenarios:

- Representative belief: the probability distribution on L, M, H is $(0, 1, 0)$;
- Optimistic belief: the distribution is $(\epsilon, 1 - \epsilon, 0)$;
- Pessimistic belief: the distribution is $(0, 1 - \epsilon, \epsilon)$.

The parameter ϵ , which we refer to as the ambiguity level, allows controlling the amount of ambiguity in the formulation: note that a larger ϵ increases the size of the ambiguity set. We consider ϵ values from $\{0, 0.33, 0.67, 1\}$, with $\epsilon = 1$ yielding the full probability simplex.

The naming of scenarios is consistent with the effect on the total loss objective. Because total losses strictly increase with R_0 , the three scenarios are naturally sorted: the optimistic belief yields the smallest expected loss, followed by the representative belief and the pessimistic belief. This also allows us to readily identify the belief that achieves the worst-case expected loss in the problem: that is always the *pessimistic* one.

We set up the robust formulation with ambiguity in the p_g^D multiplier of the 65+ y.o. age group in a similar fashion, considering the values $\{0.01, 0.05, 0.1, \dots, 1.45, 1.5\}$.

Dependence on ambiguity. In Figure 7 we report the gains of AGE-ACT over NO-TARGET as a function of the ambiguity parameter ϵ for $\chi \in \{30\times, 60\times, 90\times\}$, for ambiguity in R_0 and for ambiguity in the p_g^D multiplier.

Note that the magnitude of gains from targeting remains quite high, exceeding 40% in some cases with $\chi = 30\times$ and exceeding 20% in some cases with $\chi = 90\times$, and across all ambiguity levels ϵ . This suggests that targeting remains quite relevant even under ambiguity.

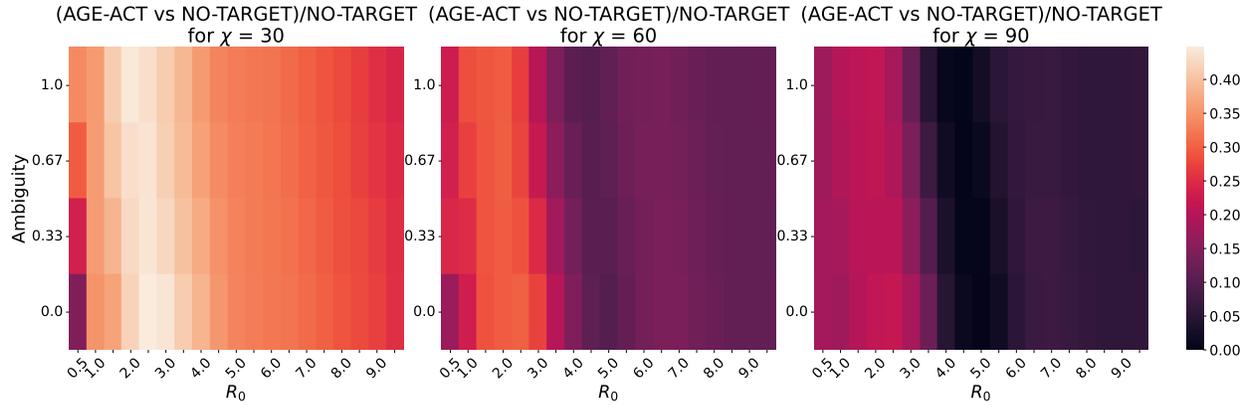
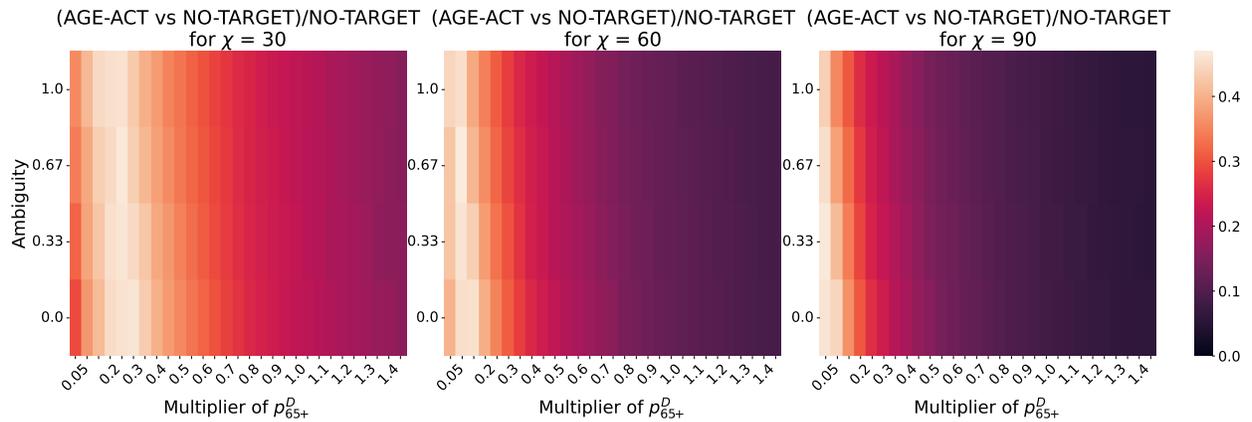
(a) Gains of ROLD AGE-ACT over NO-TARGET with ambiguity in R_0 (b) Gains of ROLD AGE-ACT over NO-TARGET with ambiguity in the multiplier of p_g^D for the 65+ y.o.

Figure 7 Heatmap with the relative gains of ROLD AGE-ACT vs. ROLD NO-TARGET, as a function of the ambiguity probability parameter and **(a)** the reproduction number R_0 ; **(b)** the multiplier of the probability of death given infection p_g^D of age group 65+ y.o. (The figure is best viewed in color.)

A critical insight is that the magnitude of gains does *not* depend monotonically on the ambiguity level ϵ , contradicting the naive intuition that simpler, less targeted policies would be more effective in highly ambiguous environments. In fact, drawing attention to how gains depend on R_0 and ambiguity in Figure 7a, we observe that there are certain values of R_0 where gains increase in ambiguity, and other values where the opposite happens. Moreover, the monotonicity switches back and forth as R_0 increases. For example, for $\chi = 60\times$, the gains are first increasing in ambiguity (for $R_0 = 0.5$ or 1), then decreasing (for $R_0 = 3$), then increasing again (for $R_0 = 5.5$), and finally decreasing (for $R_0 = 8$).

This pattern is driven entirely by the non-monotonic dependency of gains on (R_0, χ) that we documented in Figure 4a. To appreciate this, recall from our earlier discussion that the “pessimistic” belief, which places probability ϵ on the H scenario, is the belief that yields the worst-case expected cost. Therefore, when the M scenario (with the focal value for R_0) and the H scenario (with the higher value for R_0) and the considered χ correspond to a region of the heatmap in Figure 4a where gains are *increasing* in R_0 , then gains will also *increase* with ambiguity in the robust model. And similarly, when the M and H scenarios (and χ) are such that gains in Figure 4a *decrease* with R_0 , the gains will also *decrease* with ambiguity in the robust model. As a concrete example, for $\chi = 60$, the values $R_0 = 0.5$ and $R_0 = 5.5$ both correspond to points in the left panel of Figure 4a where gains are increasing in R_0 (the first peak and the “beginning” of the second peak, respectively); this explains why the corresponding vertical bands in the middle panel of Figure 7a show gains that are increasing in ambiguity at $R_0 = 0.5$ and $R_0 = 5.5$. The insight that gains in the robust model are increasing (decreasing) in ambiguity when gains increase (decrease) with the parameter remains robust when considering ambiguity in p_g^D instead of R_0 (Figure 7b).

Importantly, we note that our qualitative findings would hold even if we considered other robust formulations. In effect, any robust model where the support of the distributions in the ambiguity set involves both low and high values of R_0 would yield the same qualitative behavior, whereby gains from targeting can either increase or decrease with ambiguity. What drives this is simply the fact that total loss increases in R_0 and the fact that gains depend in *non-monotonic* ways on R_0 , as our stylized results in Section 4 proved and our empirical results in Section 6.3 confirmed.

These results show that targeting may achieve significant gains even in the context of early pandemic decision-making, when parameters are uncertain and policymakers are ambiguity-averse. That ambiguity has subtle effects on the gains from targeting—which may be challenging to intuit or quantify—further underscores the need for a structured optimization approach such as ROLD, which can help develop a precise understanding and make important policy decisions.

8. Discussion

We conclude by discussing a number of directions for future research that our current framework and results do not tackle.

Implementability. A potential hurdle to using highly targeted policies is that they are more complex to operationalize and more contentious than uniform policies. This is particularly the case with age-based targeting. Although outside our scope here, one could design policies that function as proxies to ROLD AGE-ACT without explicit age targeting. For example, one could 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. Another direction would be to consider recommending—as opposed to enforcing—separate activity levels for each age group, while accounting for imperfect population compliance. Our preliminary experiments indicate that both of these variants could still retain a significant fraction of the benefits of targeting discussed in this paper.

Endogenous feedback. There is growing evidence of the presence of behavioral dynamics changing pandemic parameters such as transmission rates (Rahmandad et al. 2021, 2022). For example, Rahmandad et al. (2022) build a forecast model where transmission rates parametrically depend on perceived risk of death (using lagged per capita deaths as proxy) and show that endogenizing feedback into a compartmental model can lead to significant improvements in accuracy. In principle, such endogeneous feedback loops could be built into our decision model, at the expense of increased computational and estimation complexity.

Other pandemic mitigation tools. Our paper is a case study in the effectiveness of targeting, focused on a single intervention mechanism. Clearly, confinements are only one of the intervention tools available to policymakers and should be used in conjunction with vaccinations, mask wearing, testing and contact tracing in any real-world policy. In fact, Group et al. (2023), Jentsch et al. (2021) observe the importance of an integrated, system-aware approach for future pandemic preparedness. A benefit of our framework is flexibility—though outside the scope of our paper, we believe these other interventions, as well as more nuanced interactions between them, could be built into our model and the ROLD framework.

References

- Acemoglu D, Chernozhukov V, Werning I, Whinston MD (2021) Optimal targeted lockdowns in a multigroup SIR model. *American Economic Review: Insights* 3(4):487–502, URL <http://dx.doi.org/10.1257/aeri.20200590>.
- Adelman D (2020) Thousands of lives could be saved in the US during the COVID-19 pandemic if states exchanged ventilators. *Health Affairs* 39(7):1247–1252, URL <http://dx.doi.org/10.1377/hlthaff.2020.00505>, pMID: 32352846.
- 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*.
- Anderson RM, May RM (1992) *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press).
- Barnett M, Buchak G, Yannelis C (2023) Epidemic responses under uncertainty. *Proceedings of the National Academy of Sciences* 120(2):e2208111120.
- Bastani H, Drakopoulos K, Gupta V, Vlachogiannis I, Hadjicristodoulou C, Lagiou P, Magiorakis 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>.
- Ben-Tal A, Nemirovski A (2002) Robust optimization—methodology and applications. *Mathematical programming* 92:453–480.
- 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, Boussioux L, Cory-Wright R, Delarue A, Digalakis V, Jacquillat A, Kitane DL, Lukin G, Li M, Mingardi L, Nohadani O, Orfanoudaki A, Papalexopoulos T, Paskov I, Pauphilet J, Lami OS, Stellato B, Bouardi HT, Carballo KV, Wiberg H, Zeng C (2021) From predictions to prescriptions: A data-driven response to COVID-19. *Health Care Management Science* 24(2):253–272, URL <http://dx.doi.org/10.1007/s10729-020-09542-0>.
- Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM review* 53(3):464–501.
- 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*.
- Bhardwaj A, Ou HC, Chen H, Jabbari S, Tambe M, Panicker R, Raval A (2020) Robust lock-down optimization for covid-19 policy guidance. *AI4SG@ AAAI Fall Symposium*.
- 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.
- Bloori A, Saghafian S (2023) Health and economic impacts of lockdown policies in the early stage of COVID-19 in the United States. *Service Science* 0(0):null, URL <http://dx.doi.org/10.1287/serv.2023.0321>.
- Bose S, Souyris S, Mukherjee UK, Ivanov A, Seshadri S, England AC, Xu Y (2021) Control of epidemic spreads via testing and lock-down. *2021 60th IEEE Conference on Decision and Control (CDC)*, 4272–4279 (IEEE).
- Brandeau ML, Zaric GS, Richter A (2003) Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *Journal of Health Economics* 22(4):575–598, ISSN 0167-6296, URL [http://dx.doi.org/https://doi.org/10.1016/S0167-6296\(03\)00043-2](http://dx.doi.org/https://doi.org/10.1016/S0167-6296(03)00043-2).
- Calvia A, Gozzi F, Lippi F, Zanco G (2024) A simple planning problem for covid-19 lockdown: a dynamic programming approach. *Economic Theory* 77(1):169–196.
- 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.
- 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.
- Dullerud GE, Paganini F (2013) *A course in robust control theory: a convex approach*, volume 36 (Springer Science & Business Media).
- 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.
- 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*.
- Faranda D, Alberti T (2020) Modeling the second wave of covid-19 infections in france and italy via a stochastic seir model. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30(11).
- Favero CA, Ichino A, Rustichini A (2020) Restarting the economy while saving lives under COVID-19. C.E.P.R Discussion Paper 14664, CEPR.

- Fotouhi H, Mori N, Miller-Hooks E, Sokolov V, Sahasrabudhe S (2021) Assessing the effects of limited curbside pickup capacity in meal delivery operations for increased safety during a pandemic. *Transportation Research Record* 2675(5):436–452, URL <http://dx.doi.org/10.1177/0361198121991840>.
- 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.
- 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.
- Fu C, Sim M, Zhou M (2021) Robust epidemiological prediction and optimization. *Available at: SSRN* 3869521.
- 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.
- 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.
- Greenwood PE, Gordillo LF (2009) Stochastic epidemic modeling. *Mathematical and statistical estimation approaches in epidemiology*, 31–52 (Springer).
- Group RSEW, et al. (2023) Covid-19: examining the effectiveness of non-pharmaceutical interventions .
- 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.
- He S, Tang S, Rong L, et al. (2020) A discrete stochastic model of the covid-19 outbreak: Forecast and control. *Math. Biosci. Eng* 17(4):2792–2804.
- Huang HC, Araz OM, Morton DP, Johnson GP, Damien P, Clements B, Meyers LA (2017) Stockpiling ventilators for influenza pandemics. *Emerging Infectious Diseases* 23(6):914–921, URL <http://dx.doi.org/10.3201/eid2306.161417>.
- Jentsch PC, Anand M, Bauch CT (2021) Prioritising covid-19 vaccination in changing social and epidemiological landscapes: a mathematical modelling study. *The Lancet Infectious Diseases* 21(8):1097–1106.
- Kanbur N, Ankgül S (2020) Quaranteenagers: A single country pandemic curfew targeting adolescents in Turkey. *Journal of Adolescent Health* 67(2).
- Kaplan EH (2020) Om forum—covid-19 scratch models to support local decisions. *Manufacturing & Service Operations Management* 22(4):645–655, URL <http://dx.doi.org/10.1287/msom.2020.0891>.
- Khor SK, Heymann DL (2021) Pandemic preparedness in the 21st century: which way forward? *The Lancet Public Health* 6(6):e357–e358.
- Köhler J, Schwenkel L, Koch A, Berberich J, Pauli P, Allgöwer F (2021) Robust and optimal predictive control of the covid-19 outbreak. *Annual Reviews in Control* 51:525–539.
- 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.
- Lekone PE, Finkenstädt BF (2006) Statistical inference in a stochastic epidemic seir model with control intervention: Ebola as a case study. *Biometrics* 62(4):1170–1177.
- Lobato FS, Libotte GB, Platt GM (2021) Mathematical modelling of the second wave of covid-19 infections using deterministic and stochastic sidr models. *Nonlinear Dynamics* 106(2):1359–1373.
- 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 (2021) Vaccine optimization for COVID-19: Who to vaccinate first? *Science Advances* 7(6):eabf1374, URL <http://dx.doi.org/10.1126/sciadv.abf1374>.
- Mehrotra S, Rahimian H, Barah M, Luo F, Schantz K (2020) A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. *Naval Research Logistics (NRL)* 67(5):303–320, URL <http://dx.doi.org/https://doi.org/10.1002/nav.21905>.
- Morris DH, Rossine FW, Plotkin JB, Levin SA (2021) Optimal, near-optimal, and robust epidemic control. *Communications Physics* 4(1):1–8.
- Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, Massari M, Salmaso S, Tomba GS, Wallinga J, Heijne J, Sadkowska-Todys M, Rosinska M, Edmunds WJ (2008) Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLOS Medicine* 5(3):1–1, URL <http://dx.doi.org/10.1371/journal.pmed.0050074>.
- Navabi-Shirazi M, El Tonbari M, Boland N, Nazzal D, Steimle LN (2022) Multicriteria course mode selection and classroom assignment under sudden space scarcity. *Manufacturing & Service Operations Management* 24(6):3252–3268, URL <http://dx.doi.org/10.1287/msom.2022.1131>.

- Olivares A, Staffetti E (2022) Robust optimal control of compartmental models in epidemiology: Application to the covid-19 pandemic. *Communications in Nonlinear Science and Numerical Simulation* 111:106509.
- 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.
- 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.
- Rahmandad H, Lim TY, Sterman J (2021) Behavioral dynamics of covid-19: estimating underreporting, multiple waves, and adherence fatigue across 92 nations. *System dynamics review* 37(1):5–31.
- Rahmandad H, Xu R, Ghaffarzagdegan N (2022) Enhancing long-term forecasting: Learning from covid-19 models. *PLoS computational biology* 18(5):e1010100.
- 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.
- Sirleaf EJ, Clark H (2021) Report of the independent panel for pandemic preparedness and response: making covid-19 the last pandemic. *The Lancet* 398(10295):101–103.
- 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.
- Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press).
- Wan J, Mirghorbani S, Wu NE, Cheng C (2023) Model predictive control in optimal intervention of covid-19 with mixed epistemic-aleatoric uncertainty. *2023 Winter Simulation Conference (WSC)*, 722–733 (IEEE).
- 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).
- Zaric GS, Brandeau ML (2001) Resource allocation for epidemic control over short time horizons. *Mathematical Biosciences* 171(1):33–58, ISSN 0025-5564, URL [http://dx.doi.org/https://doi.org/10.1016/S0025-5564\(01\)00050-5](http://dx.doi.org/https://doi.org/10.1016/S0025-5564(01)00050-5).
- Zaric GS, Brandeau ML (2002) Dynamic resource allocation for epidemic control in multiple populations. *Mathematical Medicine and Biology: A Journal of the IMA* 19(4):235–255, URL <http://dx.doi.org/10.1093/imamb/19.4.235>.
- Zhou K, Doyle JC (1998) *Essentials of robust control*, volume 104 (Prentice hall Upper Saddle River, NJ).

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.

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:

$$\Delta N_g(t) = -\mu \cdot p_g^D \cdot I_g(t) \quad (\text{EC.1})$$

$$\Delta S_g(t) = -\beta \cdot S_g(t) \cdot \left(\sum_{h \in \mathcal{G}} c_{g,h}(\ell_g(t), \ell_h(t)) \cdot \frac{I_h(t)}{N_h(t)} \right) \quad (\text{EC.2})$$

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

$$\Delta I_g(t) = \sigma \cdot E_g(t) - \underbrace{\mu \cdot p_g^D \cdot I_g(t)}_{\text{deceased}} - \underbrace{\mu \cdot (1 - p_g^D) \cdot I_g(t)}_{\text{recovered}} \quad (\text{EC.4})$$

$$\Delta R_g(t) = \mu \cdot (1 - p_g^D) \cdot I_g(t) \quad (\text{EC.5})$$

$$\Delta D_g(t) = \mu \cdot p_g^D \cdot I_g(t). \quad (\text{EC.6})$$

We now provide justification for how we account for social contacts and, in particular, for the expressions in (EC.2) and (EC.3). Fix a person i in age group $g \in \mathcal{G}$, in state S_g . Then:

$$Pr \left(\bigcup_{h \in \mathcal{G}} \{i \text{ got infected through socializing with age group } h\} \right) \quad (\text{EC.7})$$

$$= Pr \left(\bigcup_{h \in \mathcal{G}} \{i \text{ got infected through socializing with individuals in } I_h\} \right) \quad (\text{EC.8})$$

$$= 1 - Pr(\{i \text{ did not get infected through socializing with individuals in any } I_h \text{ for any } h \in \mathcal{G}\}) \quad (\text{EC.9})$$

$$= 1 - \prod_{h \in \mathcal{G}} Pr(\{i \text{ did not get infected through socializing with individuals in } I_h\}) \quad (\text{EC.10})$$

$$= 1 - \prod_{h \in \mathcal{G}} \left(1 - \beta \frac{I_h(t)}{N_h(t)} \right)^{c_{g,h}} \quad (\text{EC.11})$$

$$\approx 1 - \prod_{h \in \mathcal{G}} \left(1 - \beta c_{g,h} \frac{I_h(t)}{N_h(t)} \right) \quad (\text{EC.12})$$

$$\approx \beta \sum_{h \in \mathcal{G}} c_{g,h} \frac{I_h(t)}{N_h(t)}. \quad (\text{EC.13})$$

In (EC.11) we use the following reasoning. Having fixed person i in age group g , (a) any contact with a randomly chosen individual in group h will result in person i getting infected with probability $\beta \frac{I_h(t)}{N_h(t)}$, 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.12) and (EC.13) 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 3.3, economic losses come from three separate sources:

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

The $v_g(\ell)$ 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

$$v_g(\ell) := \begin{cases} v_g^{\text{schooling}}(\ell) + v_g^{\text{employment}}(\ell) & \text{if } g = 0\text{-}19 \text{ y.o.} \\ v_g^{\text{employment}}(\ell) & \text{if } g = 20\text{-}64 \text{ y.o.} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{EC.14})$$

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

- **Value from employment** $v_g^{\text{employment}}(\ell)$. 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, $v_g^{\text{employment}}(\ell)$ is linear in ℓ_g^{work} and the weighted average of $\ell^{\text{transport}}$, ℓ^{leisure} and ℓ^{other} across these three activities and all groups $g \in \mathcal{G}$:

$$v_g^{\text{employment}}(\ell) := w_g \cdot \left(\nu^{\text{work}} \cdot \ell_g^{\text{work}} + \nu^{\text{other activities}} \cdot \left(\frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} w^{\text{transport}} \ell_h^{\text{transport}} + w^{\text{leisure}} \ell_h^{\text{leisure}} + w^{\text{other}} \ell_h^{\text{other}} \right) + \nu^{\text{fixed}} \right). \quad (\text{EC.15})$$

Additionally, ν^{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.15). 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^{leisure} , w^{other} capture the relative importance of each of these three activities for employment value. For our baseline setting, we take $w^{\text{transport}} = w^{\text{leisure}} = w^{\text{other}} = 1/3$.

We estimate the coefficients of this model from data, as we describe in detail in Appendix EC.6.

- **Value from schooling** $v_g^{\text{schooling}}(\ell)$. A day of schooling for the individuals in the 0-19 year-old age group 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 20-64 year-old group multiplied by a factor,

and we discount for a number of years corresponding to the difference between the midpoint of the 0-19 year-old age group and the beginning of the 20-64 year-old group. Concretely, the discounting factor we apply is

$$\delta_{0-19 \text{ y.o.}} = (1 + r)^{-10},$$

where r is the discount rate. We further multiply the wage by $f_{0-19 \text{ y.o.}}$, which is an estimate of the fraction of the 0-19 year-old age group that is in school.¹² 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 take $\theta = 0.5$ in our experiments.

Thus, the definition for value of school days is

$$v_{0-19 \text{ y.o.}}^{\text{schooling}}(\ell) := \theta \cdot f_{0-19 \text{ y.o.}} \cdot \delta_{0-19 \text{ y.o.}} \cdot v_{20-64 \text{ y.o.}}^{\text{employment}}(\mathbf{1}) \cdot \ell_{0-19 \text{ y.o.}}^{\text{school}}. \quad (\text{EC.16})$$

Effect of deaths during the pandemic. We capture the economic effect of deaths during the pandemic by assuming that if at some time period, an individual in group g is in SEIR chamber D_g , then they generate no economic value.

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} , and we have

$$v_{0-19 \text{ y.o.}}^{\text{life}} := \sum_{\tau=10}^{64} \frac{1}{(1+r)^{\tau-10}} \cdot (\mathbb{1}\{10 \leq \tau \leq 19\} \cdot v_{0-19 \text{ y.o.}}^{\text{employment}}(\mathbf{1}) + \mathbb{1}\{20 \leq \tau \leq 64\} \cdot v_{20-64 \text{ y.o.}}^{\text{employment}}(\mathbf{1})), \quad (\text{EC.17})$$

$$v_{20-64 \text{ y.o.}}^{\text{life}} := \sum_{\tau=42}^{64} \frac{1}{(1+r)^{\tau-42}} \cdot v_{20-64 \text{ y.o.}}^{\text{employment}}(\mathbf{1}). \quad (\text{EC.18})$$

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,

$$V := \sum_{t=0}^{T-1} \sum_{g \in \mathcal{G}} v_g(\mathbf{1}) \cdot N_g(0). \quad (\text{EC.19})$$

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

¹² 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. In the experiments for Île-de-France, we estimate $f_{0-19 \text{ y.o.}} = 1 - (0.1876 \cdot 764916 / 3140965) \approx 0.9543$, where 18.76% is the regression-estimated employment rate within the 15-19 year olds, 764916 is the number of 15-19 year olds, and 3140965 is the number of 0-19 year olds.

EC.3. Additional Details and Proofs for Section 4

We first define all state updates and then express the objective.

State updates. Recall that confinements are only applicable in the first period, and in the second period all activities are fully open. Let $\mathbf{u} = [u_1^1, u_2^1, u_1^2, u_2^2]^\top$, where $u_g^a = \ell_g^a(0)$ denote the activity levels in period $t=0$. (Note that we sort components by activity, so that the first two components are for activity 1 and the next two for activity 2.) Then, defining the force of infection during period t for any group g as:

$$\lambda_g(0) = \beta \sum_{a \in \mathcal{A}} \sum_{h \in \mathcal{G}} C_{gh}^a u_g^a u_h^a S_g(0) \frac{I_h(0)}{N_h(0)}, \quad \lambda_g(1) = \beta \sum_{a \in \mathcal{A}} \sum_{h \in \mathcal{G}} C_{gh}^a S_g(1) \frac{I_h(1)}{N_h(1)}, \quad (\text{EC.20})$$

we can express the state evolution for each period $t \in \{0, 1\}$ as:

$$S_g(t+1) = S_g(t) - \Delta_t \lambda_g(t) \quad I_g(t+1) = I_g(t) + \Delta_t (\lambda_g(t) - \mu_g I_g(t)), \quad (\text{EC.21a})$$

$$R_g(t+1) = R_g(t) + \Delta_t (1 - p_g^D) \mu_g I_g(t) \quad D_g(t+1) = D_g(t) + \Delta_t p_g^D \mu_g I_g(t). \quad (\text{EC.21b})$$

Importantly for the subsequent developments, the total number of deaths in group g during the planning horizon is:

$$D_g(2) = \beta p_g^D \mu_g [\Delta_0 I_g(0) + \Delta_1 I_g(1)] = \text{constant} + \beta \Delta_0 \Delta_1 p_g^D \mu_g \lambda_g(0),$$

which only depends on the decisions \mathbf{u} through the force of infection $\lambda_g(0)$ from (EC.20).

Economic Value, Total Deaths. Let v_g^a denote the per-person, per-unit time economic value generated by one individual in group g engaging in activity a and define $r_{(g,a)} := \Delta_0 v_g^a N_g(0)$ for every $g \in \mathcal{G}$ and $a \in \mathcal{A}$ and $\mathbf{r} = [r_1^1, r_2^1, r_1^2, r_2^2]^\top$. The economic value generated in period t , $\text{EV}(t)$, is:

$$\begin{aligned} \text{EV}(0) &= \sum_{g \in \mathcal{G}} \sum_{a \in \mathcal{A}} \Delta_0 v_g^a N_g(0) u_g^a = \mathbf{r}^\top \mathbf{u} \\ \text{EV}(1) &= \sum_{g,a} \Delta_1 v_g^a N_g(1) = \sum_{g,a} \Delta_1 v_g^a [N_g(0) - \Delta_0 p_g^D \gamma_g I_g(0)] \end{aligned}$$

Note that only $\text{EV}(0)$ depends on the decisions \mathbf{u} , and $\text{EV}(1)$ is a constant. We can then express the total economic loss as:

$$\text{Economic Loss}(\mathbf{u}) = V - \text{EV}(0) - \text{EV}(1) + \sum_{g \in \mathcal{G}} v_g^{\text{life}} \cdot D_g(2).$$

Our objective (4) corresponding to total losses $J(\mathbf{u})$ then becomes equivalent to:

$$J(\mathbf{u}) = V - \text{EV}(0) - \text{EV}(1) + \sum_{g \in \mathcal{G}} (\chi + v_g^{\text{life}}) \cdot D_g(2).$$

EC.3.1. Proofs for Section 4.2

Proof of Proposition 1 The objective of minimizing total losses $J(\mathbf{u})$ is equivalent here to maximizing the total *net welfare* $W(\mathbf{u})$, given by the total economic value generated during the two periods minus the total penalty associated with the deaths incurred:

$$W(\mathbf{u}) = \sum_{t=0}^{T-1} \text{EV}(t) - \sum_{g \in \mathcal{G}} (\chi + v_g^{\text{life}}) D_g(2). \quad (\text{EC.22})$$

Replacing the expressions for $EV(0), EV(1), D_g(2)$ and using $\lambda_g(0)$ from (EC.20) yields:

$$W(\mathbf{u}) = W_{\text{const}} + \mathbf{r}^\top \mathbf{u} - \frac{1}{2} \mathbf{u}^\top Q \mathbf{u}, \quad (\text{EC.23})$$

where W_{const} is independent of \mathbf{u} and the matrix Q is exactly as described in (7a). \square

Proof of Proposition 2. Recall that problem (6) is separable across activities. For each $a \in \mathcal{A}$, letting $\mathbf{u}^a = [u_1^a, u_2^a]^\top$ and $\mathbf{r}^a = [r_1^a, r_2^a]^\top$, the resulting problem can be written compactly as:

$$\max_{\mathbf{u}^a \in [\varepsilon, 1]^2} (\mathbf{r}^a)^\top \mathbf{u}^a - \frac{1}{2} (\mathbf{u}^a)^\top Q^a \mathbf{u}^a \quad (\text{EC.24})$$

The proof of part (i) follows directly from Proposition EC.1 below.

(ii) To simplify notation, we omit the superscript a and prove the results for $g = 1$, without loss of generality. With the change of variables $\tilde{u}_1 = -u_1$, the objective in (EC.24) becomes:

$$W(\tilde{u}_1, u_2) = -r_1 \tilde{u}_1 + r_2 u_2 - \frac{1}{2} [B_1 u_1^2 - 2 B_{12} \tilde{u}_1 u_2 + B_2 u_2^2].$$

Because all the coefficients appearing are positive, it can be readily checked that:

$$\frac{\partial^2 W}{\partial \tilde{u}_1 \partial u_2} = B_{12} > 0 \quad \frac{\partial^2 W}{\partial \tilde{u}_1 \partial r_1} = -1 \quad \frac{\partial^2 W}{\partial \tilde{u}_1 \partial B_{12}} = u_2 \geq 0 \quad \frac{\partial^2 W}{\partial \tilde{u}_1 \partial B_2} = 0.$$

W is therefore supermodular in (\tilde{u}_1, u_2) and in (\tilde{u}_1, θ) for $\theta \in \Theta := \{-r_1, B_1, B_{12}\}$. Because we are maximizing the function $W(\mathbf{u})$ over the lattice $[-1, -\varepsilon] \times [\varepsilon, 1]$, classic results in complementarity (Theorem 2.8.2 in Topkis 1998) can be invoked to prove that \tilde{u}_1^* is increasing in θ , which in turn means u_1^* is decreasing in θ , for every $\theta \in \Theta$. Recalling the expressions of B_1, B_{12} and the expressions of A_1, A_2 then yields all the claimed comparative statics for u_1^* , except those with respect to $N_h(0)$ ($h = 1, 2$), C_{22} , and c_2^a . To see that u_1^* increases with $N_{1,2}(0)$, it suffices to note that u_1^* increases with r_1 (which increases with $N_1(0)$) and decreases with B_1 and B_{12} (which decrease with $N_1(0), N_2(0)$). The results with respect to C_{22} and r_2^a follow because although their direct effect on \tilde{u}_1 is zero (as the second order mixed derivatives are zero), an increase in these variables would increase u_2^* , and the supermodularity of W in (\tilde{u}_1, u_2) would then lead to the increase in \tilde{u}_1 and the decrease in u_1 .

PROPOSITION EC.1 (Optimal Dual-Targeted Policy). *For each activity $a \in \mathcal{A}$, the optimal dual-targeted solution $(u_g^a)^*$ is given by Table 1 and the conditions therein if $D_B^a \geq 0$, and is given by one of the expressions in Cases 2-9 from Table 1 if $D_B^a < 0$.*

Proof. Subsequently, we rewrite this as a minimization and drop the superscript a .

$$\begin{aligned} \min_{u_1, u_2} & \frac{1}{2} [B_1 u_1^2 + 2 B_{12} u_1 u_2 + B_2 u_2^2] - r_1 u_1 - r_2 u_2 \\ \text{s.t.} & \quad \varepsilon \leq u_1 \leq 1, \quad \varepsilon \leq u_2 \leq 1. \end{aligned} \quad (\text{EC.25})$$

The objective function is component-wise concave in each u_g because $B_g > 0$; it is jointly concave in $[u_1, u_2]$ if $D_B = B_1 B_2 - B_{12}^2 \geq 0$.

Because our problem trivially admits a Slater point (e.g., $u_g = (1 + \varepsilon)/2$ for every $g \in \mathcal{G}$), the constraint qualification conditions are satisfied and the KKT conditions are necessary for optimality. Moreover, if $D_B \geq 0$, the problem is convex and these conditions are sufficient for optimality.

We analyze the KKT optimality conditions. With ν and η as the dual variables corresponding to the constraints $-x \leq 0$ and $x - 1 \leq 0$, respectively, we can write the KKT system as:

$$\begin{aligned} \text{Stationarity: } & B_1 u_1 + B_{12} u_2 - r_1 - \nu_1 + \eta_1 = 0, \\ & B_{12} u_1 + B_2 u_2 - r_2 - \nu_2 + \eta_2 = 0; \\ \text{Feasibility: } & \varepsilon \leq u_g \leq 1, \quad \nu_g, \eta_g \geq 0 \quad (g = 1, 2); \\ \text{Complementary slackness: } & \nu_g(\varepsilon - u_g) = 0, \quad \eta_g(1 - u_g) = 0 \quad (g = 1, 2). \end{aligned} \tag{KKT}$$

We consider each case separately, depending on which bounds on the variables are active. ‘‘I’’ refers to an interior solution, ‘‘L’’ refers to the case when lower bound $u_g \geq \varepsilon$ is active (which implies $\eta_g = 0$) and ‘‘U’’ refers to the case when the upper bound $u_g \leq 1$ is active (which implies $\nu_g = 0$).

Case (I,I). All bounds inactive: $\nu_g = \eta_g = 0$. We can solve the system to obtain:

$$u_1 = \frac{r_1 B_2 - r_2 B_{12}}{D_B}, \quad u_2 = \frac{r_2 B_1 - r_1 B_{12}}{D_B}.$$

Feasibility. An interior solution $D_B > 0$, which is equivalent to requiring that the objective function is concave in \mathbf{u} (without this, a strictly interior solution is not possible). Moreover, this also requires $\varepsilon < u_g < 1$, i.e.

$$\varepsilon < \frac{r_1 B_2 - r_2 B_{12}}{D_B} < 1, \quad \varepsilon < \frac{r_2 B_1 - r_1 B_{12}}{D_B} < 1.$$

If these hold, (u_1, u_2) above is optimal.

Case (L,I). Fix $u_1 = \varepsilon$ ($\eta_1 = 0, \nu_1 \geq 0$), keep $\varepsilon < u_2 < 1$ ($\nu_2 = \eta_2 = 0$). Stationarity gives

$$u_2 = \frac{r_2 - B_{12}\varepsilon}{B_2}, \quad \nu_1 = B_1\varepsilon + \frac{B_{12}(r_2 - B_{12}\varepsilon)}{B_2} - r_1.$$

Conditions: $\varepsilon < \frac{r_2}{B_2} < 1$, $r_1 \leq B_1\varepsilon + \frac{B_{12}(r_2 - B_{12}\varepsilon)}{B_2}$, ensuring $u_2 \in (\varepsilon, 1)$ and $\nu_1 \geq 0$.

Case (I,L). Symmetric to (L,I). **Case (U,I).** Fix $u_1 = 1$ ($\nu_1 = 0, \eta_1 \geq 0$), $0 < u_2 < 1$ ($\nu_2 = \eta_2 = 0$).

$$u_2 = \frac{r_2 - B_{12}}{B_2}, \quad \eta_1 = r_1 - B_1 - \frac{B_{12}}{B_2}(r_2 - B_{12}).$$

Conditions: $\varepsilon < \frac{r_2 - B_{12}}{B_2} < 1$ and $r_1 B_2 - r_2 B_{12} \geq D_B$, which ensures that $\eta_1 \geq 0$.

Case (I,U). Symmetric to (U,I).

Case (L,L). Fix $u_1 = u_2 = \varepsilon$. This implies $\eta_1 = \eta_2 = 0$. We can then express

$$\nu_1 = (B_1 + B_{12})\varepsilon - r_1, \quad \nu_2 = (B_2 + B_{12})\varepsilon - r_2.$$

Conditions: $r_1 \leq \varepsilon(B_1 + B_{12})$ and $r_2 \leq \varepsilon(B_2 + B_{12})$.

Case (L,U). Fix $u_1 = \varepsilon$ (so $\eta_1 = 0$) and $u_2 = 1$ (so $\nu_2 = 0$). We get:

$$\nu_1 = B_1\varepsilon + B_{12} - r_1, \quad \eta_2 = r_2 - B_{12}\varepsilon - B_2.$$

Conditions: $r_1 \leq B_1\varepsilon + B_{12}$ and $r_2 \geq B_2 + B_{12}\varepsilon$.

Case (U,L). Symmetric to (L,U).

Case (U,U). Fix $u_1 = u_2 = 1$ ($\nu_g = 0, \eta_g \geq 0$).

$$\eta_1 = B_1 + B_{12} - r_1, \quad \eta_2 = B_2 + B_{12} - r_2.$$

Conditions. $r_1 \geq B_1 + B_{12}, r_2 \geq B_2 + B_{12}$.

When $D_B > 0$, these conditions are necessary and sufficient for optimality. When $D_B \leq 0$, note that no interior solution is possible (that is, case (I,I) does not admit a feasible solution). Therefore, in an optimal solution u_g^* , there exists at least one $g \in \mathcal{G}$ such that $u_g^* \in \{\varepsilon, 1\}$. Our objective restricted to the other variable $u_h, h \neq g$ will be convex, so the same optimality conditions discussed in Cases 2-9 would arise and the optimal solution would be given by one of the expressions corresponding to those cases.

PROPOSITION EC.2 (Age Group-Based Targeting). *The optimal group-based targeted policy $[u_1^*, u_2^*]^\top$ can be obtained from Table 1 by replacing $[r_1^a, r_2^a]^\top$ with $[\sum_{a \in \mathcal{A}} r_1^a, \sum_{a \in \mathcal{A}} r_2^a]^\top$, replacing B_g^a with $B_g^\Sigma := \sum_{a \in \mathcal{A}} B_g^a$, replacing B_{12}^a with $B_{12}^\Sigma = \sum_{a \in \mathcal{A}} B_{12}^a$, and replacing D_B with $D_B^\Sigma = B_1^\Sigma B_2^\Sigma - (B_{12}^\Sigma)^2$. With these replacements, the optimal solution is given by Table 1 and the conditions therein if $D_B^\Sigma \geq 0$, and is given by one of the expressions in Cases 2-9 from that table if $D_B^\Sigma < 0$.*

Proof of Proposition EC.2. The optimization problem for finding group-specific decisions has the same structure as the activity-specific problem we considered in (EC.24), with the only difference that the first-stage reward becomes $\mathbf{r}_g := \sum_{a \in \mathcal{A}} \mathbf{r}^a$ and the matrix Q^a is replaced with $\sum_{a \in \mathcal{A}} Q^a$. As such, the optimal solution is characterized by a counterpart of Table 1 obtained with the replacements mentioned in the statement of the proposition. The comparative statics also readily follow from an analogous argument to that used in the proof of Proposition 2, because the optimization problems are structurally identical. We omit details for brevity.

PROPOSITION EC.3 (Activity-Based Targeting). *The optimal activity-based targeted policy is separable by activity; for each activity a , the optimal activity level u_a^* and resulting welfare are given in Table EC.1, where $R^a := r_1^a + r_2^a$ and $B^a := B_1^a + 2B_{12}^a + B_2^a$. The comparative statics of the optimal policy are the same as that of the optimal AGE-ACT policy, in Proposition 2.*

Condition	Optimal intensity u^{a*}	Welfare W_{ACT}^a
$R^a \geq B^a$	$u^{a*} = 1$	$R^a - \frac{1}{2}B^a$
$B^a \varepsilon < R^a < B^a$	$u^{a*} = \frac{R^a}{B^a}$	$\frac{(R^a)^2}{2B^a}$
$R^a \leq B^a \varepsilon$	$u^{a*} = \varepsilon$	$R^a \varepsilon - \frac{1}{2}B^a \varepsilon^2$

Table EC.1 Optimal activity levels and resulting welfare for each action a .

Proof. With activity-based targeting, the optimization problem (EC.25) reduces to maximizing the concave quadratic function $R^a u - \frac{1}{2} B^a u^2$ with $u \in [\varepsilon, 1]$, so the results immediately follow. The comparative statics also readily follow because the objective is trivially supermodular in (u, R^a) and also $(u, -B^a)$ and the feasible set is a lattice, so u^* is increasing in R^a and decreasing in B^a .

EC.3.2. Proofs for Section 4.3

PROPOSITION EC.4 (**Gains From Targeting**). For each policy $\pi \in \{\text{NT}, \text{AGE}, \text{ACT}, \text{AA}\}$, define the following thresholds:

$$\chi_{\min}^{\pi}(\beta) := \begin{cases} \sum_{g,a} \omega_{g,a} \chi_{g,a}^{\min} & \pi = \text{NT}, \\ \min_g \sum_a \lambda_{g,a} \chi_{g,a}^{\min} & \pi = \text{AGE}, \\ \min_a \sum_g \rho_{g,a} \chi_{g,a}^{\min} & \pi = \text{ACT}, \\ \min_{g,a} \chi_{g,a}^{\min} & \pi = \text{AA}. \end{cases} \quad \chi_{\max}^{\pi}(\beta) := \begin{cases} \sum_{g,a} \omega_{g,a} \chi_{g,a}^{\max} & \pi = \text{NT}, \\ \max_g \sum_a \lambda_{g,a} \chi_{g,a}^{\max} & \pi = \text{AGE}, \\ \max_a \sum_g \rho_{g,a} \chi_{g,a}^{\max} & \pi = \text{ACT}, \\ \max_{g,a} \chi_{g,a}^{\max} & \pi = \text{AA}. \end{cases} \quad (\text{EC.26})$$

where for every group-activity pair $(g, a) \in \mathcal{G} \times \mathcal{A}$, we define:

$$\chi_{g,a}^{\min}(\beta) := \frac{r_g^a - \beta \Delta_0 \Delta_1 K_{g,a}}{\beta \Delta_0 \Delta_1 L_{g,a}}, \quad \chi_{g,a}^{\max}(\beta) := \frac{r_g^a - \varepsilon \beta \Delta_0 \Delta_1 K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 L_{g,a}} \quad (\text{EC.27})$$

and:

$$\lambda_{g,a} = \frac{L_{g,a}}{\sum_{a'} L_{g,a'}}, \quad \rho_{g,a} = \frac{L_{g,a}}{\sum_{g'} L_{g',a}}, \quad \omega_{g,a} = \frac{L_{g,a}}{\sum_{g',a'} L_{g',a'}},$$

and $L_{g,a}, K_{g,a}$ are given by

$$\begin{aligned} L_{g,a} &:= \left[\kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + \kappa_{\bar{g}} \cdot C_{\bar{g}g}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \\ K_{g,a} &:= \left[v_g^{\text{life}} \cdot \kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + v_{\bar{g}}^{\text{life}} \cdot \kappa_{\bar{g}} \cdot C_{\bar{g}g}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \\ \kappa_g &:= p_g^D \mu_g S_g(0) = A_g / (\chi + v_g^{\text{life}}). \end{aligned} \quad (\text{EC.28})$$

Then, the following holds:

- (i) there are zero gains from targeting if $0 \leq \chi \leq \chi_{\min}^{\pi}(\beta)$ or if $\chi \geq \chi_{\max}^{\pi}(\beta)$.
- (ii) $\chi_{\min}^{\pi}(\beta)$ and $\chi_{\max}^{\pi}(\beta)$ are piecewise linear increasing functions of $1/\beta$.
- (iii) The four thresholds satisfy the inequalities:

$$\chi_{\min}^{\text{AA}} \leq \chi_{\min}^{\text{AGE}}, \quad \chi_{\min}^{\text{ACT}} \leq \chi_{\min}^{\text{NT}} \quad \text{and} \quad \chi_{\max}^{\text{AA}} \geq \chi_{\max}^{\text{AGE}}, \quad \chi_{\max}^{\text{ACT}} \geq \chi_{\max}^{\text{NT}}.$$

Proof of Proposition EC.4. (i) Fix a transmission rate $\beta > 0$ and for a group $g \in \mathcal{G}$, let \bar{g} denote the other group (i.e., $h \in \mathcal{G}$ with $h \neq g$).

AGE-ACT. Recall the optimal AA policy from Proposition EC.1. Specifically:

$$\mathbf{u}^* = \varepsilon \cdot \mathbf{1} \quad \Leftrightarrow \text{case (L,L) for every } a \in \mathcal{A} \quad \Leftrightarrow r_g^a \leq \varepsilon (B_g^a + B_{12}^a), \quad \forall (g, a) \in \mathcal{G} \times \mathcal{A} \quad (\text{EC.29a})$$

$$\mathbf{u}^* = \mathbf{1} \quad \Leftrightarrow \text{case (U,U) for every } a \in \mathcal{A} \quad \Leftrightarrow r_g^a \geq B_g^a + B_{12}^a, \quad \forall (g, a) \in \mathcal{G} \times \mathcal{A}. \quad (\text{EC.29b})$$

Recall the definitions of B_g^a and B_{12}^a , repeated below for convenience:

$$\begin{aligned} B_g^a &= 2\beta \Delta_0 \Delta_1 A_g C_{gg}^a \frac{I_g(0)}{N_g(0)} \quad \forall g \in \mathcal{G}, \quad B_{12}^a = \beta \Delta_0 \Delta_1 \left(A_1 C_{12}^a \frac{I_2(0)}{N_2(0)} + A_2 C_{21}^a \frac{I_1(0)}{N_1(0)} \right), \\ A_g &= (\chi + v_g^{\text{life}}) p_g^D \mu_g S_g(0) \quad \text{for } g \in \mathcal{G}. \end{aligned}$$

Then, the sum $B_g^a + B_{12}^a$ can be written to highlight the dependency on β and χ as:

$$\begin{aligned} B_g^a + B_{12}^a &= \beta \Delta_0 \Delta_1 \left[L_{g,a} \cdot \chi + K_{g,a} \right], \quad \text{where} \\ L_{g,a} &= \beta \Delta_0 \Delta_1 \left[\kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + \kappa_{\bar{g}} \cdot C_{\bar{g}g}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \\ K_{g,a} &= \beta \Delta_0 \Delta_1 \left[v_g^{\text{life}} \cdot \kappa_g \left(2C_{gg}^a \cdot \frac{I_g(0)}{N_g(0)} + C_{g\bar{g}}^a \cdot \frac{I_{\bar{g}}(0)}{N_{\bar{g}}(0)} \right) + v_{\bar{g}}^{\text{life}} \cdot \kappa_{\bar{g}} \cdot C_{\bar{g}g}^a \cdot \frac{I_g(0)}{N_g(0)} \right] \\ \kappa_g &:= p_g^D \mu_g S_g(0) = A_g / (\chi + v_g^{\text{life}}). \end{aligned} \quad (\text{EC.30})$$

In that case, we can readily define for each $(g, a) \in \mathcal{G} \times \mathcal{A}$ the thresholds

$$\chi_{g,a}^{\min}(\beta) := \frac{r_g^a - \beta \Delta_0 \Delta_1 K_{g,a}}{\beta \Delta_0 \Delta_1 L_{g,a}}, \quad \chi_{g,a}^{\max}(\beta) := \frac{r_g^a - \varepsilon \beta \Delta_0 \Delta_1 K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 L_{g,a}} \quad (\text{EC.31})$$

and note that (EC.29a) holds if and only if $\chi \geq \chi_{g,a}^{\max}(\beta)$, whereas (EC.29b) holds if and only if $\chi \leq \chi_{g,a}^{\min}(\beta)$. Therefore, we conclude that $\chi_{\max}^{\pi}(\beta)$ and $\chi_{\min}^{\pi}(\beta)$ exist for $\pi = \text{AA}$, with:

$$\chi_{\min}^{\pi}(\beta) := \min_{(g,a) \in \mathcal{G} \times \mathcal{A}} \chi_{g,a}^{\min}, \quad \chi_{\max}^{\pi}(\beta) := \max_{(g,a) \in \mathcal{G} \times \mathcal{A}} \chi_{g,a}^{\max}.$$

AGE. We can repeat the same reasoning for the AGE policy. From Proposition EC.1, the optimization problem for AGE is identical to that for AGE-ACT, with the sole difference that the AGE objective involves quantities $r_g = \sum_a r_g^a$, $B_g = \sum_a B_g^a$, $B_{12} = \sum_a B_{12}^a$ aggregated over activities. Thus, the conclusions for AGE follow from those for AGE-ACT by replacing the $L_{g,a}$ with $\sum_{a \in \mathcal{A}} L_{g,a}$ and $K_{g,a}$ with $\sum_{a \in \mathcal{A}} K_{g,a}$. Specifically, with weights $\lambda_{g,a} = \frac{L_{g,a}}{\sum_{a'} L_{g,a'}}$, we define:

$$\begin{aligned} \chi_g^{\min}(\beta) &:= \frac{\sum_a r_g^a - \beta \Delta_0 \Delta_1 \sum_a K_{g,a}}{\beta \Delta_0 \Delta_1 \sum_a L_{g,a}} = \sum_{a \in \mathcal{A}} \lambda_{g,a} \cdot \chi_{g,a}^{\min}(\beta), \quad \forall g \in \mathcal{G} \\ \chi_g^{\max}(\beta) &:= \frac{\sum_a r_g^a - \varepsilon \beta \Delta_0 \Delta_1 \sum_a K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 \sum_a L_{g,a}} = \sum_{a \in \mathcal{A}} \lambda_{g,a} \cdot \chi_{g,a}^{\max}(\beta), \quad \forall g \in \mathcal{G}. \end{aligned}$$

Then, we obtain the desired result for $\pi = \text{AGE}$ by setting:

$$\chi_{\min}^{\pi}(\beta) := \min_{g \in \mathcal{G}} \chi_g^{\min}, \quad \chi_{\max}^{\pi}(\beta) := \max_{g \in \mathcal{G}} \chi_g^{\max}.$$

ACT. The same reasoning applies to ACT. For activity $a \in \{1, 2\}$ the policymaker chooses a single control u^a . Let $R^a = r_1^a + r_2^a$ and $B^a = B_1^a + 2B_{12}^a + B_2^a$. Then, from Proposition EC.3, we have that

$$(u^a)^* = \varepsilon \Leftrightarrow R^a \leq \varepsilon B^a \quad \text{and} \quad (u^a)^* = 1 \Leftrightarrow R^a \geq B^a.$$

We can readily see that $B^a = \sum_{g \in \mathcal{G}} L_{g,a} \cdot \chi + \sum_{g \in \mathcal{G}} K_{g,a}$, where $L_{g,a}$ and $K_{g,a}$ are defined as in AGE-ACT. Then, following the same logic as for AGE but with weights $\rho_{g,a}$ given by $\rho_{g,a} = \frac{L_{g,a}}{\sum_{g'} L_{g',a}}$, we define the following thresholds:

$$\begin{aligned} \chi_a^{\min}(\beta) &:= \frac{\sum_g r_g^a - \beta \Delta_0 \Delta_1 \sum_g K_{g,a}}{\beta \Delta_0 \Delta_1 \sum_g L_{g,a}} = \sum_{a \in \mathcal{A}} \rho_{g,a} \cdot \chi_{g,a}^{\min}(\beta), \quad \forall a \in \mathcal{A} \\ \chi_a^{\max}(\beta) &:= \frac{\sum_g r_g^a - \varepsilon \beta \Delta_0 \Delta_1 \sum_g K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 \sum_g L_{g,a}} = \sum_{a \in \mathcal{A}} \rho_{g,a} \cdot \chi_{g,a}^{\max}(\beta), \quad \forall a \in \mathcal{A}. \end{aligned}$$

Then, we obtain the desired result for $\pi = \text{ACT}$ by setting:

$$\chi_{\min}^{\pi}(\beta) := \min_{a \in \mathcal{A}} \chi_a^{\min}, \quad \chi_{\max}^{\pi}(\beta) := \max_{a \in \mathcal{A}} \chi_a^{\max}.$$

NO-TARGET. The same reasoning applies to NO-TARGET. With $R = \sum_{g,a} r_g^a$ and $B = \sum_{a \in \mathcal{A}} (B_1^a + 2B_{12}^a + B_2^a)$, we have that

$$u^* = \varepsilon \Leftrightarrow R \leq \varepsilon B \quad \text{and} \quad u^* = 1 \Leftrightarrow R \geq B.$$

With the same logic as for AGE and ACT, with weights $\omega_{g,a}$ given by $\omega_{g,a} = \frac{L_{g,a}}{\sum_{g',a'} L_{g',a'}}$, we obtain our desired thresholds for policy $\pi = \text{NT}$:

$$\begin{aligned}\chi_{\min}^{\pi}(\beta) &:= \frac{\sum_{g,a} r_g^a - \beta \Delta_0 \Delta_1 \sum_{g,a} K_{g,a}}{\beta \Delta_0 \Delta_1 \sum_{g,a} L_{g,a}} = \sum_{a \in \mathcal{A}} \omega_{g,a} \cdot \chi_{g,a}^{\min}(\beta) \\ \chi_{\max}^{\pi}(\beta) &:= \frac{\sum_{g,a} r_g^a - \varepsilon \beta \Delta_0 \Delta_1 \sum_{g,a} K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 \sum_{g,a} L_{g,a}} = \sum_{a \in \mathcal{A}} \omega_{g,a} \cdot \chi_{g,a}^{\max}(\beta).\end{aligned}$$

In conclusion, the desired thresholds for each policy $\pi \in \{\text{NT}, \text{AGE}, \text{ACT}, \text{AA}\}$ are

$$\chi_{\min}^{\pi}(\beta) := \begin{cases} \sum_{g,a} \omega_{g,a} \chi_{g,a}^{\min} & \pi = \text{NT}, \\ \min_g \sum_a \lambda_{g,a} \chi_{g,a}^{\min} & \pi = \text{AGE}, \\ \min_a \sum_g \rho_{g,a} \chi_{g,a}^{\min} & \pi = \text{ACT}, \\ \min_{g,a} \chi_{g,a}^{\min} & \pi = \text{AA}. \end{cases} \quad \chi_{\max}^{\pi}(\beta) := \begin{cases} \sum_{g,a} \omega_{g,a} \chi_{g,a}^{\max} & \pi = \text{NT}, \\ \max_g \sum_a \lambda_{g,a} \chi_{g,a}^{\max} & \pi = \text{AGE}, \\ \max_a \sum_g \rho_{g,a} \chi_{g,a}^{\max} & \pi = \text{ACT}, \\ \max_{g,a} \chi_{g,a}^{\max} & \pi = \text{AA}. \end{cases} \quad (\text{EC.32})$$

where the weights are:

$$\lambda_{g,a} = \frac{L_{g,a}}{\sum_{a'} L_{g,a'}}, \quad \rho_{g,a} = \frac{L_{g,a}}{\sum_{g'} L_{g',a}}, \quad \omega_{g,a} = \frac{L_{g,a}}{\sum_{g',a'} L_{g',a'}}.$$

and $L_{g,a}, K_{g,a}$ are defined in (EC.30) and $\chi_{g,a}^{\min}(\beta), \chi_{g,a}^{\max}(\beta)$ are defined in (EC.31).

(ii) Recall from (EC.31) that:

$$\chi_{g,a}^{\min}(\beta) := \frac{r_g^a - \beta \Delta_0 \Delta_1 K_{g,a}}{\beta \Delta_0 \Delta_1 L_{g,a}}, \quad \chi_{g,a}^{\max}(\beta) := \frac{r_g^a - \varepsilon \beta \Delta_0 \Delta_1 K_{g,a}}{\varepsilon \beta \Delta_0 \Delta_1 L_{g,a}}.$$

Because $L_{g,a}, K_{g,a} > 0$, both $\chi_{g,a}^{\min}(\beta)$ and $\chi_{g,a}^{\max}(\beta)$ strictly decrease in β . Moreover, each bound tends to the strictly negative value $-\frac{K_{g,a}}{L_{g,a}} < 0$ as $\beta \rightarrow +\infty$ and tends to $+\infty$ as $\beta \rightarrow 0$ (because $r_{g,a} > 0$). Because taking convex combinations, minimums or maximums of these functions preserves these properties, we can conclude with the desired result.

(iii) The inequalities between the thresholds follow from (EC.32), because a convex combination lies between the minimum and the maximum of the elements in the combination.

(iv) It suffices to prove this result for the parameter regime wherein all policies have an interior solution. This parameter regime trivially belongs to the region of interest (because outside that region, the NO-TARGET policy would not have an interior solution). For this interior parameter regime, the result follows from Proposition EC.5.

PROPOSITION EC.5. *Consider the parameter regime wherein all policies have interior optimal activity levels, i.e., $\mathbf{u}^* \in (\varepsilon, 1)^4$. Then, the optimal net welfare values are:*

$$W_{\text{NT}}^* = \frac{R^2}{B}, \quad \text{where } R := \sum_{a,g} r_g^a, \quad B = \sum_a (B_1^a + 2B_{12}^a + B_2^a) \quad (\text{EC.33})$$

$$W_{\text{ACT}}^* = \frac{1}{2} \sum_a \frac{(R^a)^2}{B^a}, \quad \text{where } R^a := \sum_g r_g^a, \quad B^a := B_1^a + 2B_{12}^a + B_2^a \quad (\text{EC.34})$$

$$W_{\text{AGE}}^* = \frac{1}{2} \cdot \frac{(R_1)^2 B_2^G + (R_2)^2 B_1^G - 2R_1 R_2 B_{12}^G}{D_B^G}, \quad \text{where} \quad (\text{EC.35})$$

$$R_g := \sum_a r_g^a, B_g^G := \sum_a B_g^a, B_{12}^G := \sum_a B_{12}^a, D_B^G := B_1^G \cdot B_2^G - (B_{12}^G)^2 \quad (\text{EC.36})$$

$$W_{AA}^* = \frac{1}{2} \sum_a \frac{r_1^{a2} B_2^a + r_2^{a2} B_1^a - 2r_1^a r_2^a B_{12}^a}{D_B^a}. \quad (\text{EC.37})$$

Moreover, the relative welfare gain of policy $\pi \in \{\text{AGE}, \text{ACT}, \text{AA}\}$ with respect to *NO-TARGET*,

$$G^\pi := \frac{W_\pi^*}{W_{\text{NT}}^*} - 1,$$

is non-monotonic with respect to the parameters $\{r_g^a\}_{g \in \mathcal{G}, a \in \mathcal{A}}, A_1, A_2$.

Proof of Proposition EC.5 The expressions for W_π^* follow from Proposition EC.1, Proposition EC.3, and Proposition EC.2, by using the expression of the optimal interior solution. In the subsequent derivations, we use the notation $R, B, R^a, B^a, R_g, B_g^G, B_{12}^G$ defined in the expressions for W_π^* above.

Change in r_g^a . We claim that

$$\frac{\partial G_{\text{ACT}}}{\partial r_g^a} \geq 0 \iff \frac{R^a}{B^a} \geq \frac{(R^1)^2/B^1 + (R^2)^2/B^2}{R^1 + R^2}. \quad (\text{EC.38})$$

Without loss of generality, take $g = 1$ and $a = 1$. Note that we can write:

$$G_{\text{ACT}} = \frac{B}{R^2} H - 1, \quad \text{where } H := \frac{(R^1)^2}{B^1} + \frac{(R^2)^2}{B^2}$$

Because only R^1 and $R = R^1 + R^2$ depend on x in G_{ACT} and the derivative of each w.r.t x is 1, we have that the derivative of G_{ACT} with respect to x has the same sign as $B^2 R^1 - B^1 R^2$, which is exactly equivalent to (EC.38) in this case. The sign therefore depends on problem parameters. An analogous argument can be used to show that G_{ACT} and G_{AA} can also either increase or decrease with r_g^a .

Change in factor A_1 . Consider a change in the factor A_1 . The change in relative gains of the ACT policy, $\frac{d}{dA_1} G^{\text{ACT}}$ has the same sign as

$$\Psi = B^{1'} \left[\frac{(R^1)^2}{(B^1)^2} (B^1 - B) + \frac{(R^2)^2}{B^2} \right] + B^{2'} \left[\frac{(R^2)^2}{(B^2)^2} (B^2 - B) + \frac{(R^1)^2}{B^1} \right],$$

where

$$B^{1'} := \partial_{A_1} B^1 = 2\kappa \sum_a (C_{11}^a I_1(0)/N_1(0) + C_{12}^a I_2(0)/N_2(0)) > 0, \quad B^{2'} = 0.$$

Then, replacing terms, we arrive at the conclusion that $\Psi > 0$ if and only if $R^2 B^1 > R^1 B^2$. A similar argument can be used for AGE and AGE-ACT, as well as with respect to A_2 . \square

Proof of Proposition 3. Recall that the main conditions in the statement of the proposition require that:

$$\exists a \in \mathcal{A} : \frac{\sum_g r_g^a}{\sum_g L_{g,a}} > \max_g \frac{\sum_a r_g^a}{\sum_a L_{g,a}} \quad \text{and} \quad \exists g \in \mathcal{G} : \frac{\sum_a K_g^a}{\sum_a L_{g,a}} < \min_{a \in \mathcal{A}} \frac{\sum_g K_{g,a}}{\sum_g L_{g,a}}.$$

where $L_{g,a}$ and $K_{g,a}$ are given by (EC.30). To simplify notation slightly, let us define the parameters:

$$\alpha_{\max} := \max_a \frac{\sum_g r_g^a}{\sum_g L_{g,a}}, \quad \gamma_{\max} := \max_g \frac{\sum_a r_g^a}{\sum_a L_{g,a}}, \quad B_{\min} := \min_a \frac{\sum_g K_{g,a}}{\sum_g L_{g,a}}, \quad D_{\min} := \min_g \frac{\sum_a K_{g,a}}{\sum_a L_{g,a}},$$

Then, the conditions above require that:

$$\alpha_{\max} > \gamma_{\max} \quad \text{and} \quad B_{\min} > D_{\min}. \quad (\text{EC.39})$$

Let us then recall from Proposition EC.4 the definitions of the ACT threshold $\chi_{\max}^{\text{ACT}}(\beta)$ and AGE threshold $\chi_{\max}^{\text{AGE}}(\beta)$, whose expressions we repeat for convenience:

$$\chi_{\max}^{\text{ACT}}(\beta) = \max_{a \in \mathcal{A}} \left(\frac{\alpha_a}{\beta} - \beta_a \right), \quad \text{where } \alpha_a := \frac{\sum_g r_g^a}{\varepsilon \beta \Delta_0 \Delta_1 \sum_g L_{g,a}}, \quad \beta_a := \frac{\sum_g K_{g,a}}{\sum_g L_{g,a}}$$

$$\chi_{\max}^{\text{AGE}}(\beta) = \max_{b \in \mathcal{G}} \left(\frac{\gamma_b}{\beta} - \delta_b \right), \quad \text{where } \gamma_b := \frac{\sum_a r_a^b}{\varepsilon \beta \Delta_0 \Delta_1 \sum_a L_{g,a}}, \quad \delta_b := \frac{\sum_a K_{g,a}}{\sum_a L_{g,a}}.$$

Under the premise (EC.39), as $\beta \rightarrow 0$, we have $\chi_{\max}^{\text{ACT}}(\beta) \geq \chi_{\max}^{\text{AGE}}(\beta)$, because the former has a larger term multiplying $1/\beta$. Similarly, (EC.39) also implies that as β gets sufficiently large, $\chi_{\max}^{\text{ACT}}(\beta) \geq \chi_{\max}^{\text{AGE}}(\beta)$. Because both $\chi_{\max}^{\text{ACT}}(\beta)$ and $\chi_{\max}^{\text{AGE}}(\beta)$ are continuous functions of β , there exist¹³ $\beta^1 \leq \beta^2$ such that $\chi_{\max}^{\text{ACT}}(\beta) > \chi_{\max}^{\text{AGE}}(\beta)$ for $\beta < \beta^1$ and $\chi_{\max}^{\text{ACT}}(\beta) < \chi_{\max}^{\text{AGE}}(\beta)$ for $\beta > \beta^2$. The results then readily follow from the definitions of the thresholds. The reverse argument is identical. \square

EC.4. Algorithmic Details for ROLD

In this section, we clarify the algorithmic details of the linearization-optimization procedure described in Section 5.

EC.4.1. Linearized Dynamics

We first focus on how we build a linear model given k, \mathbf{X}_k and $\hat{\mathbf{u}}_{k:T-1}$. In Step 2, our algorithm builds an approximation of the state dynamics that is linear in the controls $\mathbf{u}_k, \dots, \mathbf{u}_{T-1}$. Here, we compute the coefficients for each \mathbf{u}_t explicitly. We introduce the notation:

$$\mathbf{A}_t := \mathbb{I} + \Delta t \cdot \nabla_{\mathbf{X}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) \quad (\text{EC.40})$$

$$\mathbf{B}_t := \Delta t \cdot \nabla_{\mathbf{u}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) \quad (\text{EC.41})$$

$$\mathbf{c}_t := \Delta t \cdot [\mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) - \nabla_{\mathbf{X}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) \cdot \hat{\mathbf{X}}_t - \nabla_{\mathbf{u}} \mathbf{F}_t(\hat{\mathbf{X}}_t, \hat{\mathbf{u}}_t) \cdot \hat{\mathbf{u}}_t], \quad (\text{EC.42})$$

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

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

We can then express the state \mathbf{X}_t as¹⁴

$$\mathbf{X}_t = \left(\prod_{\tau=t-1, t-2, \dots, k} \mathbf{A}_\tau \right) \mathbf{X}_k + \sum_{\tau=k}^{t-1} \left(\prod_{i=t-1, t-2, \dots, \tau+1} \mathbf{A}_i \right) \mathbf{B}_\tau \mathbf{u}_\tau + \sum_{\tau=k}^{t-1} \left(\prod_{i=t-1, t-2, \dots, \tau+1} \mathbf{A}_i \right) \mathbf{c}_\tau. \quad (\text{EC.43})$$

It is now possible to express the objective linearly in the decisions \mathbf{u}_t .

¹³ In general, we can show that at most three intersection points are possible in this case: instead of viewing the curves as hyperbolas in $1/\beta$, one can view them as piecewise linear functions in $x = 1/\beta$ with at most two pieces, which makes it easy to argue that at most three intersection points can exist.

¹⁴ In the expressions in (EC.43) and (EC.46), we follow the convention that a product of matrices over an empty set of indices results in the identity matrix.

EC.4.2. Objective Coefficients

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

$$\sum_{t=k}^T (\mathbf{d}_t^\top \mathbf{X}_t + \mathbf{e}_t^\top \mathbf{u}_t) \quad (\text{EC.44})$$

with

$$\mathbf{d}_t^\top = \begin{cases} \hat{\mathbf{u}}_t^\top \mathbf{M} + \boldsymbol{\gamma}^\top, & \text{if } t < T \\ \boldsymbol{\eta}^\top, & \text{if } t = T, \end{cases} \quad \mathbf{e}_t^\top = \begin{cases} \hat{\mathbf{X}}_t^\top \mathbf{M}^\top, & \text{if } t < T \\ 0, & \text{if } t = T. \end{cases} \quad (\text{EC.45})$$

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

$$\hat{\mathbf{X}}_t^\top \mathbf{M}^\top + \sum_{\tau=t+1}^{T-1} (\hat{\mathbf{u}}_\tau^\top \mathbf{M} + \boldsymbol{\gamma}^\top) \left(\prod_{i=\tau-1, \tau-2, \dots, t+1} \mathbf{A}_i \right) \mathbf{B}_t + \boldsymbol{\eta}^\top \left(\prod_{i=T-1, T-2, \dots, t+1} \mathbf{A}_i \right) \mathbf{B}_t. \quad (\text{EC.46})$$

This allows calculating the coefficients recursively. The detailed function `CALCULATEOBJECTIVECOEFFICIENTS` is defined in Algorithm 1.

Calculation of \mathbf{M} , $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$. Expanding the objective (4), we have:

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

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

$$\mathbf{M}[\ell_g^{\text{work}}, S_g] = \mathbf{M}[\ell_g^{\text{work}}, E_g] = \mathbf{M}[\ell_g^{\text{work}}, I_g] = \mathbf{M}[\ell_g^{\text{work}}, R_g] = \begin{cases} -w_g \nu^{\text{work}}, & \text{if } g = 0-19, 20-64 \text{ y.o.} \\ 0, & \text{otherwise.} \end{cases}$$

$$\begin{aligned} \mathbf{M}[\ell_h^a, S_g] &= \mathbf{M}[\ell_h^a, E_g] = \mathbf{M}[\ell_h^a, I_g] = \mathbf{M}[\ell_h^a, R_g] \\ &= \begin{cases} \frac{-w_g \nu^{\text{other activities}}}{3|\mathcal{G}|}, & \text{if } g = 0-19, 20-64 \text{ y.o., } h \in \mathcal{G}, a \in \{\text{transport, leisure, other}\} \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

$$\mathbf{M}[\ell_g^{\text{school}}, S_g] = \mathbf{M}[\ell_g^{\text{school}}, E_g] = \mathbf{M}[\ell_g^{\text{school}}, I_g] = \mathbf{M}[\ell_g^{\text{school}}, R_g] = \begin{cases} -\theta f_g \delta_g v_{20-64 \text{ y.o.}}^{\text{employment}}(\mathbf{1}), & \text{if } g = 0-19 \text{ y.o.} \\ 0, & \text{otherwise.} \end{cases}$$

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

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

$$\boldsymbol{\gamma}[S_g] = \boldsymbol{\gamma}[E_g] = \boldsymbol{\gamma}[I_g] = \boldsymbol{\gamma}[R_g] = \begin{cases} -w_g \nu^{\text{fixed}}, & \text{if } g = 0-19, 20-64 \text{ y.o.} \\ 0, & \text{otherwise.} \end{cases} \quad (\text{EC.52})$$

$$\boldsymbol{\gamma}[\cdot] = 0, \quad \text{otherwise.} \quad (\text{EC.53})$$

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

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

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

Algorithm 1 CALCULATEOBJECTIVECOEFFICIENTS

-
- 1: Initialization: $\bar{\mathbf{X}}_k := \mathbf{X}_k$
 - 2: **for** $t = k, k + 1, \dots, T - 1$ **do**
 - 3: Calculate coefficients for \mathbf{u}_τ in the period t summand of (EC.44) as:

$$\begin{cases} \mathbf{e}_t^\top, & \text{for } \tau = t \\ \mathbf{d}_t^\top \bar{\mathbf{A}}_{t,\tau} \mathbf{B}_\tau, & \text{for } \tau = t - 1, t - 2, \dots, k \end{cases} \quad (\text{EC.47})$$

where

$$\bar{\mathbf{A}}_{t,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = t - 1 \\ \bar{\mathbf{A}}_{t,\tau+1} \mathbf{A}_{\tau+1}, & \text{for } \tau = t - 2, t - 3, \dots, k. \end{cases} \quad (\text{EC.48})$$

- 4: **end for**
- 5:
- 6: **for** $t = T$ **do**
- 7: Calculate coefficients for \mathbf{u}_τ in the period T summand of (EC.44) as:

$$\mathbf{d}_T^\top \bar{\mathbf{A}}_{T,\tau} \mathbf{B}_\tau, \quad \text{for } \tau = T - 1, T - 2, \dots, k \quad (\text{EC.49})$$

where

$$\bar{\mathbf{A}}_{T,\tau} := \begin{cases} \mathbb{I}, & \text{for } \tau = T - 1 \\ \bar{\mathbf{A}}_{T,\tau+1} \mathbf{A}_{\tau+1}, & \text{for } \tau = T - 2, T - 3, \dots, k. \end{cases} \quad (\text{EC.50})$$

- 8: **end for**
- 9:
- 10: Sum up the terms attributed to a common \mathbf{u}_τ :

$$\mathbf{e}_\tau^\top + \sum_{t=\tau+1}^T \mathbf{d}_t^\top \bar{\mathbf{A}}_{t,\tau} \mathbf{B}_\tau. \quad (\text{EC.51})$$

- 11: **return** Coefficients for \mathbf{u}_τ for each $\tau = k, k + 1, \dots, T - 1$
-

EC.4.3. Specifics of the Iterative Linearization-Optimization Procedure

Having defined the function CALCULATEOBJECTIVECOEFFICIENTS, the LINEARIZATION-OPTIMIZATION function, which is the main subroutine of ROLD, is described in Algorithm 2. 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 $\hat{\mathbf{u}}_{k:T-1}^{(k)}$. We denote that ϵ -ball by $B_\epsilon^\infty(\hat{\mathbf{u}}_{k:T-1}^{(k)})$.

Algorithm 2 LINEARIZATION-OPTIMIZATION

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

- 1: Calculate nominal trajectory $\hat{\mathbf{X}}_{k:T}$ from \mathbf{X}_k and $\hat{\mathbf{u}}_{k:T-1}^{(k)}$
 - 2: Run CALCULATEOBJECTIVECOEFFICIENTS to compute the objective coefficients for decisions \mathbf{u}_t , $t = k, \dots, T - 1$
 - 3: Solve resulting linear program with additional trust region constraints $\mathbf{u}_{k:T-1} \in B_\epsilon^\infty(\hat{\mathbf{u}}_{k:T-1}^{(k)})$ to obtain controls $\hat{\zeta}_{k:T-1}$
 - 4: **return** $\hat{\zeta}_{k:T-1}$
-

The full ROLD procedure is given in Algorithm 3. 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 $\hat{\mathbf{u}}_{0:T-1}^{(0)}$; in our experiments we initialize this with a control sequence obtained by a random seeding procedure, which we describe below.

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 $\mathbf{u}_{k:T-1}$ 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 $mult \cdot \lceil \frac{1}{\epsilon} \rceil$.

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.

Algorithm 3 ROLD

Require: initial state \mathbf{X}_0 , initial nominal control $\hat{\mathbf{u}}_{0:T-1}^{(0)}$

- 1: **for** $k = 0, \dots, T - 1$ **do**
- 2: $i := 0$ and $\mathbf{u}_{k:T-1}^{(k,0)} := \hat{\mathbf{u}}_{k:T-1}^{(k)}$
- 3: **while** $i \leq mult \cdot \lceil \frac{1}{\epsilon} \rceil$ and $\mathbf{u}_{k:T-1}^{(k,i)} \neq \mathbf{u}_{k:T-1}^{(k,j)}, \forall 0 \leq j < i$ **do**
- 4: $\mathbf{u}_{k:T-1}^{(k,i+1)} = \text{LINEARIZATION-OPTIMIZATION}(k, \mathbf{X}_k, \mathbf{u}_{k:T-1}^{(k,i)})$
- 5: $i = i + 1$
- 6: **end while**
- 7: Set $\mathbf{u}_{k:T-1}^{*(k)} = \arg \min_{\mathbf{u}_{k:T-1} \in \{\mathbf{u}_{k:T-1}^{(k,0)}, \dots, \mathbf{u}_{k:T-1}^{(k,i-1)}\}} J([\mathbf{u}_{0:k-1}^*, \mathbf{u}_{k:T-1}])$
- 8: Set the nominal control sequence to $\hat{\mathbf{u}}_{k+1:T-1}^{(k+1)} = \mathbf{u}_{k+1:T-1}^{*(k)}$
- 9: Set $\mathbf{u}_k^* = \mathbf{u}_k^{*(k)}$ and update the system state one step:

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \mathbf{F}_k(\mathbf{X}_k, \mathbf{u}_k^*) \tag{EC.54}$$

10: **end for**

EC.4.4. Initialization for ROLD

ROLD relies on a seeding for the nominal control initialization $\hat{\mathbf{u}}_{0:T-1}^{(0)}$. In this section, we describe how we do this seeding. In order to prevent ROLD from getting stuck at a bad local minimum, we seed it with several initializations and choose the best performing one at the end. We organize the seeding in two passes.

First pass. For all targeting strategies, we run ROLD by seeding it with the fully open and fully closed policies, as well as 20 random starting points. We sample these random seeds in different ways, depending on the targeting strategy.

- ROLD NO-TARGET: We sample i.i.d. uniformly at random an integer τ from the time steps at which confinement decisions are made. We set the activity level to 0 for $t \leq \tau$, and to 1 for $t > \tau$.
- ROLD AGE: We use the same sampling procedure as for NO-TARGET but separately for each age group. In addition, we also initialize at the best ROLD NO-TARGET solution of the first pass.
- ROLD ACT: We use the same sampling procedure as for NO-TARGET but separately for each activity other than home. In addition, we also initialize at the best ROLD NO-TARGET solution of the first pass.
- ROLD AGE-ACT: We use the same sampling procedure as for NO-TARGET but separately for each (age group, activity) pair such that the activity is relevant for that group, and with activity other than home. In addition, we also initialize at the best ROLD AGE and ROLD ACT solutions of the first pass.

In each case, we select the solution that attains the best objective.

Second pass. Here we initialize ROLD at additional solutions: the solutions of ROLD runs on parameter values that are adjacent to the focal ones in the experimental setup. For example, if in an experiment the tested values of R_0 are $\{0.01, 0.5, 1, 1.5, 2, 2.5, \dots, 9.5, 10\}$, then to run ROLD with $R_0 = 3$ we initialize it also at the ROLD solutions for $R_0 = 2.5$ and $R_0 = 3.5$.

- ROLD NO-TARGET. Initialize at two or four additional solutions: the best ROLD NO-TARGET solution from the first pass for parameters adjacent to the focal ones.¹⁵
- ROLD AGE: Initialize at three or five additional solutions: the best ROLD NO-TARGET solution from the second pass, and the best ROLD AGE solution from the first pass for parameters adjacent to the focal ones.
- ROLD ACT: Initialize at three or five additional solutions: the best ROLD NO-TARGET solution from the second pass, and the best ROLD ACT solution from the first pass for parameters adjacent to the focal ones.
- ROLD AGE-ACT: Initialize at four or six additional solutions: the best ROLD AGE and ROLD ACT solutions from the second pass, and the best ROLD AGE-ACT solution from the first pass for parameters adjacent to the focal ones.

In each case, we select the solution that attains the best objective.

For the robust runs of Section 7, the initialization happens for each of the three scenarios, with the same random starting points used across the three scenarios.

EC.5. Robust ROLD

Here, we describe the algorithmic details for the robust version of ROLD which we use to solve the robust formulation from Section 7. To extend these ideas to the robust model in (16), we first note that the key requirement in solving (16) is to determine the static decisions \mathbf{u}_t to follow in periods $t = 0, \dots, T^{\text{new}} - 1$, before the period T^{new} when the parameter values θ are revealed. Once we reach period T^{new} and we know that scenario $\omega \in \Omega$ materializes, we can determine the decisions \mathbf{u}_t^ω to follow in subsequent periods $t \in \{T^{\text{new}}, \dots, T - 1\}$ by applying the “plain-vanilla” ROLD procedure exactly as described above for the case with known parameters, implementing the first-step decision and re-solving with a shrinking horizon, which improves performance. Our robust implementation will mirror this logic.

¹⁵ In the runs for Figure 2, we have two additional solutions, each corresponding to the solution of ROLD when run with a value of χ that is adjacent to the focal value of χ among the set of tested χ values. In the runs for the heatmaps in Section 6.3 with χ and R_0 , we have four additional solutions: two solutions correspond to the solution of ROLD when run with the same χ , and with the two values of R_0 that are adjacent to the focal R_0 among the set of tested R_0 values; another two solutions correspond to the solution of ROLD when run with the same R_0 , and with the two values of χ that are adjacent to the focal χ among the set of tested χ values. Similarly for the experiments for the heatmaps with χ and p_g^D .

To that end, we first describe how to solve the robust problem at time $t = 0$ to find the decisions \mathbf{u}_t to follow in periods $t = 0, \dots, T^{\text{new}} - 1$.¹⁶ The main ideas are the same as in a “plain-vanilla” ROLD: we iteratively solve a linearized version of the problem within a trust region, changing the linearization at each step to deal with the nonlinear dynamics. The main building block requires solving a linearized version of the problem, which we obtain by linearizing the states under *each* scenario $\omega \in \Omega$. Formally, given a nominal control sequence $\hat{\mathbf{u}}_{0:T-1}^\omega$ and its associated nominal system trajectory $\hat{\mathbf{X}}_{0:T}^\omega$ for each scenario $\omega \in \Omega$, we linearize the states $\mathbf{X}_{0:T}^\omega$ with evolution described by (15) around the nominal trajectories. Paralleling the notation from Appendix EC.4.1, we use $\mathbf{A}_t^\omega, \mathbf{B}_t^\omega, \mathbf{c}_t^\omega$ to denote the system matrices and $\mathbf{e}_t^\omega, \mathbf{d}_t^\omega$ to denote the objective coefficients for period t . Then, the main building block at $t = 0$ requires solving the following linear program:

$$\begin{aligned}
& \min_{\{\mathbf{u}_{0:T-1}^\omega\}_{\omega \in \Omega}, U} U \\
& \text{s.t. } U \geq \sum_{\omega \in \Omega} p_\omega \cdot \left[\sum_{t=0}^T \left((\mathbf{d}_t^\omega)^\top \mathbf{X}_t^\omega + (\mathbf{e}_t^\omega)^\top \mathbf{u}_t^\omega \right) \right] \quad \forall \mathbf{p} \in \{\mathbb{P}^1, \mathbb{P}^2, \dots, \mathbb{P}^m\} \\
& \mathbf{X}_t^\omega = \left(\prod_{\tau=t-1, t-2, \dots, 0} \mathbf{A}_\tau^\omega \right) \mathbf{X}_0^\omega + \sum_{\tau=0}^{t-1} \left(\prod_{i=t-1, t-2, \dots, \tau+1} \mathbf{A}_i^\omega \right) \mathbf{B}_\tau^\omega \mathbf{u}_\tau^\omega \\
& \quad + \sum_{\tau=0}^{t-1} \left(\prod_{i=t-1, t-2, \dots, \tau+1} \mathbf{A}_i^\omega \right) \mathbf{c}_\tau^\omega \quad \forall t \in \{0, \dots, T-1\}, \forall \omega \in \Omega \\
& \mathbf{u}_t^\omega \in [0, 1]^{|\mathcal{G}||\mathcal{U}|} \quad \forall t \in \{0, 1, \dots, T-1\}, \forall \omega \in \Omega \\
& \mathbf{u}_t^\omega = \mathbf{u}_t^{\omega'} \quad \forall t \in \{0, 1, \dots, T^{\text{new}} - 1\}, \forall \omega, \omega' \in \Omega.
\end{aligned} \tag{EC.55}$$

Note that this is an epigraph reformulation of the robust model (16) for an SEIR system that follows the approximate linearized dynamics in each scenario $\omega \in \Omega$. Notably, although our robust model in (16) allowed the adversary to choose any distribution $\mathbf{p} \in \mathcal{P}$, the formulation above restricts attention to choices $\mathbf{p} \in \mathbb{P}^1, \dots, \mathbb{P}^m$; this is without loss of optimality, because the adversary’s objective function is linear in \mathbf{p} (the expectation operator is linear in the probabilities) and the extreme beliefs are the extreme points of \mathcal{P} . The optimal value of U will therefore be the smallest worst-case expected loss obtained based on linearizing the dynamics around $\hat{\mathbf{u}}_{0:T-1}^\omega$ and $\hat{\mathbf{X}}_{0:T}^\omega$.

The main building block at time $t = 0$ will solve the program above, constraining the decisions using the same trust-region logic as the “plain-vanilla” ROLD. We summarize this building block in Algorithm 4.

Algorithm 4 ROBUST LINEARIZATION-OPTIMIZATION AT $t = 0$

Require: initial state \mathbf{X}_0 , nominal control initialization $\hat{\mathbf{u}}_{0:T-1}^\omega$ for every scenario $\omega \in \Omega$

- 1: Calculate nominal trajectory $\hat{\mathbf{X}}_{0:T}^\omega$ from \mathbf{X}_0 and $\hat{\mathbf{u}}_{0:T-1}^\omega$ for every scenario $\omega \in \Omega$
 - 2: For every $\omega \in \Omega$, run CALCULATEOBJECTIVECOEFFICIENTS to compute the objective coefficients for decisions \mathbf{u}_t^ω for $t \in \{0, \dots, T-1\}$ when the parameters take value $\boldsymbol{\theta} = \boldsymbol{\theta}^\omega$.
 - 3: Solve the linear program (EC.55) with additional trust region constraints $\mathbf{u}_{0:T-1}^\omega \in B_\epsilon^\infty(\hat{\mathbf{u}}_{0:T-1}^\omega)$ for every scenario $\omega \in \Omega$, to obtain controls $\{\boldsymbol{\zeta}_{0:T-1}^\omega\}_{\omega \in \Omega}$
 - 4: **return** $\{\boldsymbol{\zeta}_{0:T-1}^\omega\}_{\omega \in \Omega}$
-

Given this building block, we summarize the main Robust ROLD procedure in Algorithm 5.

¹⁶ We do not allow any re-solving of the ROLD heuristic for periods $t = 1, 2, \dots, T^{\text{new}} - 1$ because no new information is revealed during those periods. Moreover, to formulate a robust model with re-solving over a shrinking horizon while maintaining dynamic consistency in the formulations and decisions is nontrivial.

Algorithm 5 Robust ROLD

Require: initial state \mathbf{X}_0 , initial nominal controls $\hat{\mathbf{u}}_{0:T-1}^{\omega,(0)}$ for every scenario $\omega \in \Omega$

```

1: #Calculate decisions  $\mathbf{u}_t$  for periods  $t \in \{0, \dots, T^{\text{new}} - 1\}$ 
2:  $i := 0$  and  $\mathbf{u}_{0:T-1}^{\omega,(0,0)} := \hat{\mathbf{u}}_{0:T-1}^{\omega,(0)} \forall \omega \in \Omega$ 
3: while  $i \leq \text{mult} \cdot \lceil \frac{1}{\epsilon} \rceil$  and  $\mathbf{u}_{0:T-1}^{\omega,(0,i)} \neq \mathbf{u}_{0:T-1}^{\omega,(0,j)}, \forall 0 \leq j < i, \forall \omega \in \Omega$  do
4:    $\{\mathbf{u}_{0:T-1}^{\omega,(0,i+1)}\}_{\omega \in \Omega} = \text{ROBUST LINEARIZATION-OPTIMIZATION AT } t = 0 \left( \mathbf{X}_0, \{\mathbf{u}_{0:T-1}^{\omega,(0,i)}\}_{\omega \in \Omega} \right)$ 
5:    $i = i + 1$ 
6: end while
7: Set  $\{\mathbf{u}_{0:T-1}^{*,\omega,(0)}\}_{\omega \in \Omega} = \underset{\{\mathbf{u}_{0:T-1}^{\omega,(0,0)}\}_{\omega \in \Omega}, \dots, \{\mathbf{u}_{0:T-1}^{\omega,(0,i-1)}\}_{\omega \in \Omega}}{\text{arg min}} \max_{p \in \mathcal{P}} \sum_{\omega \in \Omega} p_{\omega} \cdot J(\mathbf{u}_{0:T-1}^{\omega}, \boldsymbol{\theta}^{\omega})$ 

8: #Calculate decisions  $\mathbf{u}_t^{\omega}$  for periods  $t \in \{T^{\text{new}}, \dots, T - 1\}$  separately, for each scenario  $\omega \in \Omega$ 
9: for  $\omega \in \Omega$  do
10:   Set  $\mathbf{u}_{0:T^{\text{new}}-1}^{*,\omega} = \mathbf{u}_{0:T^{\text{new}}-1}^{*,\omega,(0)}$ 

11:   # Calculate the state at time  $T^{\text{new}}$ 
12:   for  $t = 0, \dots, T^{\text{new}} - 1$  do
13:      $\mathbf{X}_{t+1}^{\omega} = \mathbf{X}_t^{\omega} + \mathbf{F}_t(\mathbf{X}_t^{\omega}, \mathbf{u}_t^{*,\omega}, \boldsymbol{\theta}^{\omega})$  (EC.56)
14:   end for

15:   # Initialize the nominal control sequence for times  $T^{\text{new}}, \dots, T - 1$ 
16:   Set  $\hat{\mathbf{u}}_{T^{\text{new}}:T-1}^{\omega,(T^{\text{new}})} = \mathbf{u}_{T^{\text{new}}:T-1}^{*,\omega,(0)}$ 
17:   for  $k = T^{\text{new}}, \dots, T - 1$  do
18:      $i := 0$  and  $\mathbf{u}_{k:T-1}^{\omega,(k,0)} := \hat{\mathbf{u}}_{k:T-1}^{\omega,(k)}$ 
19:     while  $i \leq \text{mult} \cdot \lceil \frac{1}{\epsilon} \rceil$  and  $\mathbf{u}_{k:T-1}^{\omega,(k,i)} \neq \mathbf{u}_{k:T-1}^{\omega,(k,j)}, \forall 0 \leq j < i$  do
20:        $\mathbf{u}_{k:T-1}^{\omega,(k,i+1)} = \text{LINEARIZATION-OPTIMIZATION}(k, \mathbf{X}_k^{\omega}, \mathbf{u}_{k:T-1}^{\omega,(k,i)})$ 
21:        $i = i + 1$ 
22:     end while
23:     Set  $\mathbf{u}_{k:T-1}^{*,\omega,(k)} = \underset{\mathbf{u}_{k:T-1}^{\omega} \in \{\mathbf{u}_{k:T-1}^{\omega,(k,0)}, \dots, \mathbf{u}_{k:T-1}^{\omega,(k,i-1)}\}}{\text{arg min}} J([\mathbf{u}_{0:k-1}^{*,\omega}, \mathbf{u}_{k:T-1}^{\omega}], \boldsymbol{\theta}^{\omega})$ 
24:     Set the nominal control sequence to  $\hat{\mathbf{u}}_{k+1:T-1}^{\omega,(k+1)} = \mathbf{u}_{k+1:T-1}^{*,\omega,(k)}$ 
25:     Set  $\mathbf{u}_k^{*,\omega} = \mathbf{u}_k^{*,\omega,(k)}$  and update the system state one step:
26:      $\mathbf{X}_{k+1}^{\omega} = \mathbf{X}_k^{\omega} + \mathbf{F}_k(\mathbf{X}_k^{\omega}, \mathbf{u}_k^{*,\omega}, \boldsymbol{\theta}^{\omega})$  (EC.57)
27:   end for
28: end for

```

EC.6. Details on Parametrization of the Model for Île-de-France**EC.6.1. Basic SEIR Model Parameters**

The SEIR model parameters that are constant across age groups are summarized in Table EC.2. The age-group specific parameters are reported in Table EC.3. We use the parameters as reported in Salje et al. (2020).¹⁷

¹⁷ We retrieve the parameter values as reported before Salje et al. (2020) updated them on July 8, 2020.

Table EC.2 SEIR model parameters

Variable	Description	Value	Reference
		[Uncertainty Range]	
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)

For R_0 , the reported uncertainty range is a 95% confidence interval. For σ^{-1} (i.e., the mean stay in compartment E), the uncertainty range is calculated as $4 \pm 0.8 \cdot 0.6$ 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).¹⁸

Table EC.3 Age-group specific SEIR model probability parameters

Age group g (y.o.)	p_g^D Prob. of Death given infection
0-19	0.000012 [0.000003 - 0.000021]
20-64	0.001466 [0.000808 - 0.002123]
65+	0.041862 [0.024186 - 0.059538]

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:

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

$$\frac{\partial I_g}{\partial t} \approx \sigma E_g(t).$$

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

$$\mathbf{\Phi} = \beta \cdot \begin{bmatrix} \mathbf{0} & [c_{gh}]_{g,h \in G} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (\text{EC.58})$$

¹⁸ 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.

and

$$\Lambda = \left[\begin{array}{c|c} -\text{diag}(\boldsymbol{\sigma}) & \mathbf{0} \\ \hline \text{diag}(\boldsymbol{\sigma}) & -\text{diag}(\boldsymbol{\mu}) \end{array} \right]. \quad (\text{EC.59})$$

Then R_0 can be identified as the spectral radius (i.e., the largest absolute value of the eigenvalues) of the matrix $-\Phi\Lambda^{-1}$ (Diekmann et al. 2010, Perasso 2018). Since the eigenvalues of a matrix $\beta \cdot \mathbf{A}$ are simply β multiples of the eigenvalues of \mathbf{A} , we can therefore determine β as R_0 divided by the spectral radius of the matrix $(-\Phi/\beta)\Lambda^{-1}$.

EC.6.2. 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.15) corresponds to the employment value for a member of group g , under normal conditions. To estimate w_g , we use two datasets from INSEE:

- Yearly full time equivalent (FTE)¹⁹ 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 different from 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.

When doing this, we treat the 0-19 y.o. age group specially by assuming the employment rate is reported only with respect to the work-eligible population in that bucket (15-19 year olds). We also set the work-eligible population for the 65+ age bucket to 0. The formula we use is

$$\text{average_wage}_g = \text{employment_rate}_g \times \text{fraction_work_eligible_population}_g \times \text{employed_average_wage}_g. \quad (\text{EC.60})$$

¹⁹This is a normalization to account for employees doing part time work.

- 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 10.88%. We scale all wages average_wage_g proportionally so as to retrieve the real total wage amount w_g .

Table EC.4 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.4 Year-based economic value parameters under normal activity, per age group (in €)

Age Group g (y.o.)	$w_g \times 365$ (€)
0-19	972.90
20-64	25 236.13
65+	0.00

Estimation of ν^{work} , $\nu^{\text{other activities}}$, ν^{fixed} . We move on to the estimation of parameters ν^{work} , $\nu^{\text{other activities}}$, ν^{fixed} in (EC.15). These measure the sensitivity of economic value to the confinement pattern $\ell(t)$. 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). To simplify the estimation, we first set $\nu^{\text{other activities}} = 0$ and determine parameters ν^{work} , ν^{fixed} solving the system of equations

$$\nu^{\text{work}} + \nu^{\text{fixed}} = 1 \quad (\text{EC.61})$$

$$\nu^{\text{work}} \ell_{\text{April}}^{\text{work}} + \nu^{\text{fixed}} = 0.5851, \quad (\text{EC.62})$$

where $\ell_{\text{April}}^{\text{work}} = 0.213$ corresponds to the average value of $\ell^{\text{work}}(t)$ through the month of April 2020. Then in our experiments, we take ν^{work} to be 95% of the solution from the above system; keep the value for ν^{fixed} from the system; and set $\nu^{\text{other activities}} = 1 - \nu^{\text{work}} - \nu^{\text{fixed}}$.

EC.6.3. Parameter Values for the Experimental and Optimization setup

Table EC.5 summarizes values for the parameters in our experimental and optimization setup.

EC.7. Robustness Checks and Additional Results

EC.7.1. The Gains from Targeting for Other Pandemics and Geographies

Here we provide some additional robustness to the Pareto dominance of AGE-ACT which we observe in Section 6. We do this by testing ROLD with other pandemics and geographies with

Table EC.5 Parameter values for experimental and optimization setup. The parameters ν^{work} , $\nu^{\text{other activities}}$, ν^{fixed} , r , f_g and θ related to our economic model are defined in Appendix EC.2.

Parameter Description	Notation	Values in Experiments
Social mixing parameter	α	0.3874
Cost of death	χ	30 values in $[0, 1000] \times$ GDP per capita of France
GDP per capita of France		€37199.03
Sensitivity of econ. value on confinement	ν^{work}	0.5009
Sensitivity of econ. value on confinement	$\nu^{\text{other activities}}$	0.0264
Sensitivity of econ. value on confinement	ν^{fixed}	0.4727
Discount rate (used to calculate $v_g(\ell)$ and v_g^{life})	r	0.03
Fraction going to school	$f_{0-19 \text{ y.o.}}$	0.9543
Mult. factor for value of schooling	θ	0.5
Date of patient zero		December 20 2019
Starting time for optimization		March 17 2020
Optimization horizon	T'	90 days
Frequency of confinement decisions		14 days

different social contact and population structures than we consider in our COVID-19 in Île-de-France focal case-study.

Other pandemics. We fit our SEIR model to two other known pandemics beyond COVID-19. The first is Ebola, which has a smaller reproduction number (1.95) than COVID-19 but significantly higher death probability (81% for the 65+ age group). The second is a strain of Seasonal Influenza which has a smaller reproduction number (1.28) than COVID-19 and smaller death probability (under 1%). The parameters we use for Ebola and Seasonal Influenza are found in Table EC.6. We use these parameters using the studies of Muzembo et al. (2024), Garske et al. (2017), Chowell et al. (2004) in the case of Ebola, and Biggerstaff et al. (2014), World Health Organization (2025), McDonald et al. (2023) for Seasonal Influenza.

Table EC.6 Epidemiological parameters for Ebola and Influenza by age group

Age Group g (y.o.)	Ebola					Influenza				
	R_0	μ	σ	p_{death}	p_{recovery}	R_0	μ	σ	p_{death}	p_{recovery}
0–19	1.95	0.17825	0.15873	0.56472	0.43528	1.28	0.14286	0.50	0.00010	0.99990
20–64	1.95	0.17825	0.15873	0.63848	0.36152	1.28	0.14286	0.50	0.00159	0.99841
65+	1.95	0.17825	0.15873	0.81494	0.18506	1.28	0.14286	0.50	0.00540	0.99460

Other geographies. We also consider two other geographies, Hong Kong and South Africa. We calibrate these geographies to be different from Île-de-France across two important problem parameters:

- **Social contacts.** We use the social contact matrices calibrated for these geographies, which we retrieve from the data tool of Wille et al. (2020), but rescale them in the following way. The contacts of South Africa, are much lower than France’s; this makes the robustness check uninteresting since the optimal policies are full open at all χ levels. We thus rescale all three contact matrices so that the overall number of contacts is the same and set to that of France. However, the heterogeneity across activities and population groups remains specific to the geography.
- **Population buckets and wage information.** Each geography is calibrated to a different size and bucketing of population, as well as different economic contributions from wages, which we calibrate from Census and Statistics Department, The Government of the Hong

Kong Special Administrative Region (2025a,c,b) for Hong Kong and Department: Statistics, Republic of South Africa (2019, 2024, 2025) for South Africa and we report in Table EC.7. All other economic model characteristics are replicated from our baseline model for Île-de-France.

Table EC.7 Population and wage data for Hong Kong and South Africa geographies

Age Group g (y.o.)	Hong Kong		South Africa	
	Population	$w_g \times 365$	Population	$w_g \times 365$
0–19	1,039,200	253.20	21,559,289	950.40
20–64	4,766,000	15,057.28	33,694,956	3170.22
65+	1,718,900	2,133.89	3,520,776	0.00

In order to see whether the Pareto dominance of AGE-ACT replicates across all these settings, we run a factorial 3×3 experiment where we construct a problem instance using each pandemic and geography pair in $\{\text{COVID-19, Ebola, Seasonal Influenza}\} \times \{\text{Hong Kong, Île-de-France, South Africa}\}$. We report the results in Figure EC.1. We observe that across all instances, all targeting variants not only generate gains over NO-TARGET, but Pareto dominate it. As we observed in Section 6, it is not readily clear that AGE or ACT Pareto dominate each other, but they are dominated by AGE-ACT.

EC.7.2. Sensitivity Analyses

We analyze additional problem instances by changing the value of each of six estimated parameters within a sensitivity range, as shown in Table EC.8. We sample 1000 samples i.i.d. from the joint uniform distribution specified by the sensitivity ranges indicated in the table.

Table EC.8 Sensitivity analysis: parameters and sensitivity ranges

Parameter	Estimated Value	Sensitivity Range
R_0	2.9	2.7-3.1
σ^{-1}	4.0	3.0-5.0
μ^{-1}	4.0	3.0-5.0
Multiplier of p_g^D	1	0.5-1.5
Social mixing α	0.3874	0.1874-0.5874
Sensitivity of econ. value ν^{work}	0.5009	0.4009 - 0.6009

Figure EC.2 shows sensitivity analysis 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 1000 problem instances, for each value of χ . These results reinforce our findings from Section 6 on the gains of targeting.

EC.7.3. Additional Plots on Sensitivity of Gains from Targeting on Key Parameters

Figure EC.3 breaks down deaths and economic loss of optimized ROLD policies, showing the dependence on R_0 and χ .

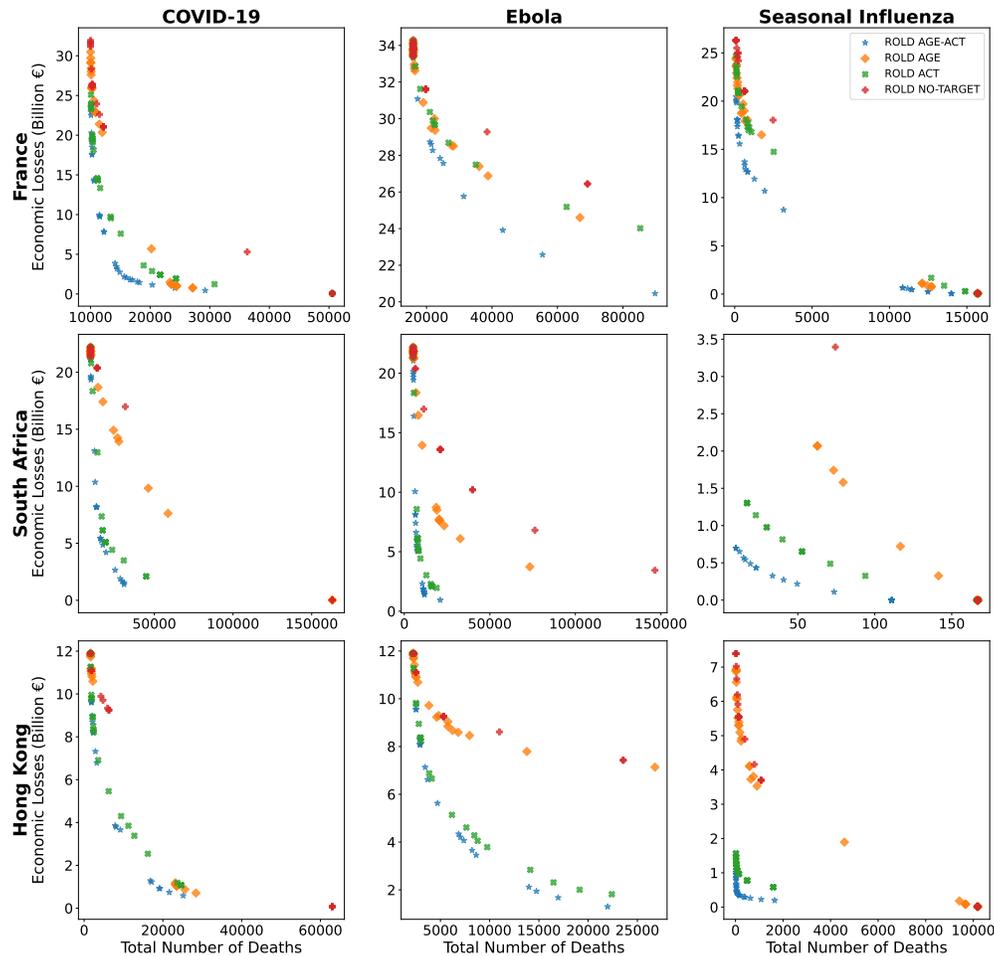


Figure EC.1 The total number of deaths and the economic losses generated by targeted ROLD policies in each problem instance of the 3×3 experiment for the gains of targeting for other pandemic and geographies. We construct a problem instance using each pandemic and geography pair in $\{\text{COVID-19, Ebola, Seasonal Influenza}\} \times \{\text{Hong Kong, Île-de-France, South Africa}\}$.

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.
- 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).
- Biggerstaff M, Cauchemez S, Reed C, Gambhir M, Finelli L (2014) Estimates of the reproduction number for seasonal, pandemic, and zoonotic influenza: a systematic review of the literature. *BMC infectious diseases* 14(1):1–20.
- Census and Statistics Department, The Government of the Hong Kong Special Administrative Region (2025a) Table 110-01001 : Population by sex and age group. https://www.censtatd.gov.hk/en/web_table.html?id=110-01001, accessed June 13, 2025.
- Census and Statistics Department, The Government of the Hong Kong Special Administrative Region (2025b) Table 210-06301:employed persons by age and sex. https://www.censtatd.gov.hk/en/web_table.html?id=210-06301, accessed June 13, 2025.
- Census and Statistics Department, The Government of the Hong Kong Special Administrative Region (2025c) Table 220-23012 : Monthly wage level and distribution by employment nature and age group. https://www.censtatd.gov.hk/en/web_table.html?id=220-23012, accessed June 13, 2025.

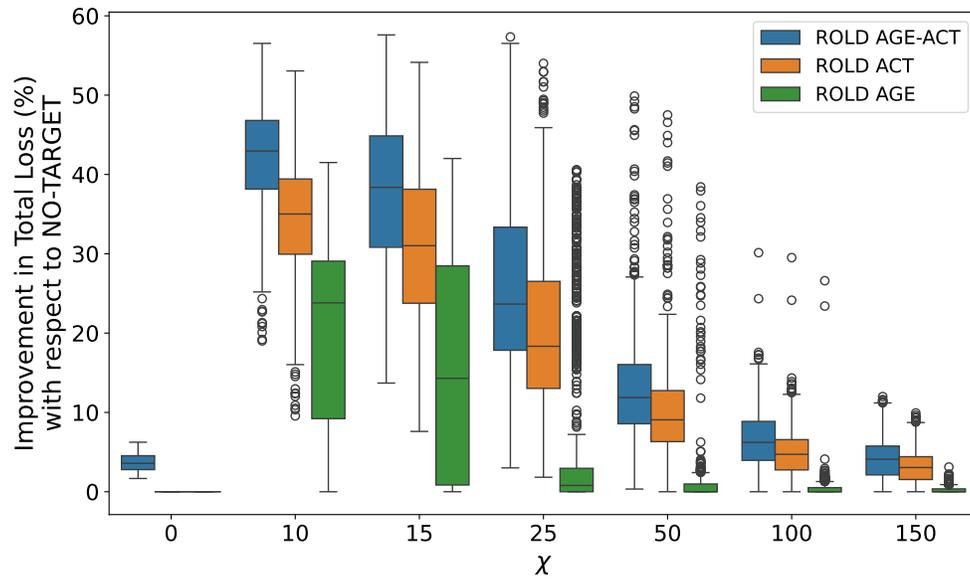
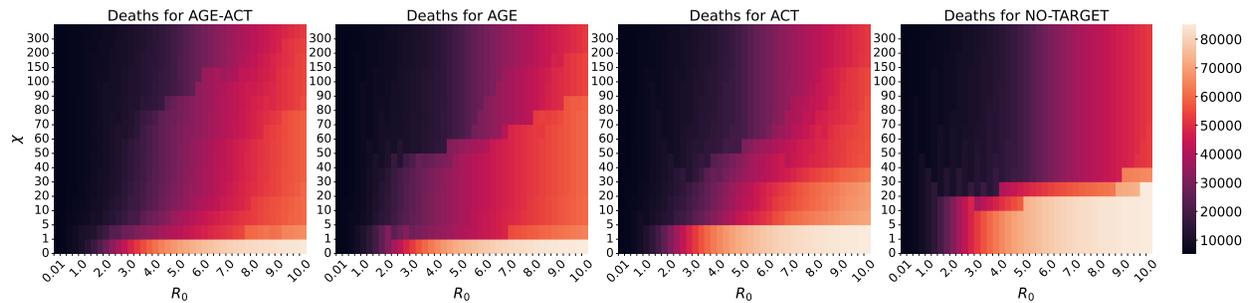
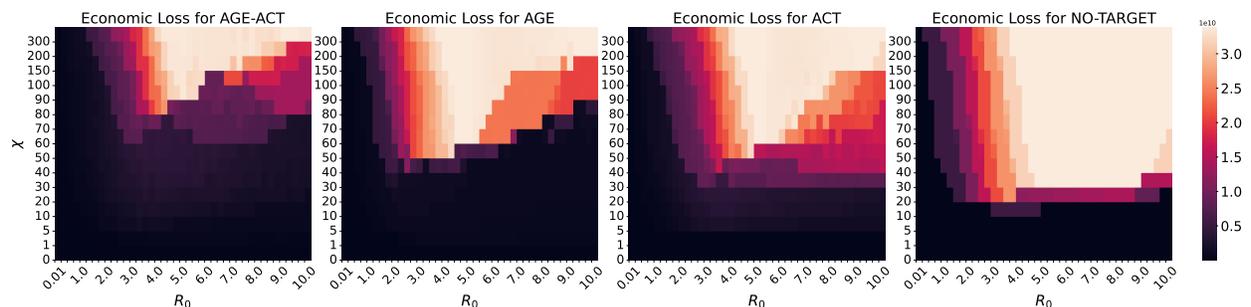


Figure EC.2 Sensitivity analyses showing the superiority of ROLD AGE-ACT over ROLD policies with less granular targeting, for different values for the cost of death χ , over a wide set of problem instances. All improvements are with respect to ROLD NO-TARGET. For each value of χ , the boxplots summarize results over 1000 different problem instances. (The figure is best viewed in color.)



(a) Total Deaths



(b) Economic Loss

Figure EC.3 Heatmaps of the Total Deaths and Economic Loss under optimal ROLD policies as a function of the basic reproduction number R_0 and the cost of death χ . (The figure is best viewed in color.)

- Department: Statistics, Republic of South Africa (2019) Mid-year population estimates 2019. <https://www.statssa.gov.za/publications/P0302/P03022019.pdf>, accessed June 13, 2025.
- Department: Statistics, Republic of South Africa (2024) Monthly earnings in south africa 2017 - 2022. <https://www.statssa.gov.za/publications/02-11-20/02-11-202022.pdf>, accessed June 13, 2025.
- Department: Statistics, Republic of South Africa (2025) Quarterly labour force survey quarter 4: 2024. <https://www.statssa.gov.za/publications/P0211/P02114thQuarter2024.pdf>, accessed June 13, 2025.
- 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).
- Garske T, Cori A, Ariyaratna A, Blake IM, Dorigatti I, Eckmanns T, Fraser C, Hinsley W, Jombart T, Mills HL, et al. (2017) Heterogeneities in the case fatality ratio in the west african ebola outbreak 2013–2016. *Philosophical Transactions of the Royal Society B: Biological Sciences* 372(1721):20160308.
- 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 quatrième trimestre 2019, le taux de chômage passe de 8,5 % à 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.
- McDonald SA, Teirlinck AC, Hooiveld M, van Asten L, Meijer A, de Lange M, van Gageldonk-Lafeber AB, Wallinga J (2023) Inference of age-dependent case-fatality ratios for seasonal influenza virus subtypes a (h3n2) and a (h1n1) pdm09 and b lineages using data from the netherlands. *Influenza and Other Respiratory Viruses* 17(6):e13146.
- 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.
- Muzembo BA, Kitahara K, Mitra D, Ntontolo NP, Ngatu NR, Ohno A, Khatiwada J, Dutta S, Miyoshi SI (2024) The basic reproduction number (r_0) of ebola virus disease: A systematic review and meta-analysis. *Travel Medicine and Infectious Disease* 57:102685.
- 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.
- 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).
- World Health Organization (2025) Influenza (seasonal). [https://www.who.int/news-room/fact-sheets/detail/influenza-\(seasonal\)](https://www.who.int/news-room/fact-sheets/detail/influenza-(seasonal)), accessed June 13, 2025.