# Strategic Optimization for Epidemic Response Planning

Jai Rajaram

June 12, 2025

**Abstract**

Epidemics caused by infectious diseases pose critical challenges to public health systems and regional economies. This work presents an optimization framework that integrates epidemiological modeling with policy planning to mitigate the impact of infectious disease outbreaks. Building on traditional SIRD models, we incorporate a suite of public health and social interventions, resulting in a complex nonlinear integer program. To address computational challenges, we propose a Lagrangian decomposition approach for lower bounds and practical heuristics for solution generation. We evaluated performance in a variety of scenarios, including those reminiscent of the COVID-19 pandemic. Our framework supports decision-making at the regional level, offering valuable insights into the trade-offs between health outcomes and intervention costs.

## 1 Introduction

The initial outbreak of the COVID-19 (SARS-CoV-2) virus occurred in Wuhan, China, in mid-November 2019 and rapidly spread across the globe by early 2020. By March of that year, the virus had reached nearly every region worldwide, quickly becoming one of the most widespread pandemics in recorded history—though not the deadliest. The disease exhibited particularly high mortality among older adults, individuals with compromised immune systems, and those with underlying health conditions. By the end of 2021, approximately 275 million people had been infected, and nearly 5.4 million deaths were attributed to the virus. Across various regions, efforts to contain the virus were frequently delayed, inadequately designed, or inconsistently enforced. Many experts, including Chakraborty, Abhijit et al., 2020 and Brodeur, Abel et al., 2021, argued that a significant number of deaths—even in affluent nations like the US and UK—could have been postponed or avoided altogether. Beyond the tragic loss of life, the pandemic imposed severe economic consequences due to the stringent policies needed to curb transmission and reduce fatality rates.

In addition to implementation delays and policy inefficiencies, there was notable disagreement—both within and between countries—about which measures should be adopted and under what circumstances. Resistance to specific US policies and mandates came from individuals, organizations, public officials, and institutions. In some cases, opposition stemmed from the economic burdens imposed on businesses, sectors, and individuals. In others, it was driven by limited understanding of viral transmission mechanisms and the necessity of coordinated, comprehensive interventions. Misconceptions about the virus and proposed policies, ideological objections, and outright denial of expert opinions further complicated the public response. Political polarization also played a role in shaping attitudes toward public health mandates.

This study proposes and tests a modeling framework designed to simulate the trajectory of an epidemic while capturing the effects of mitigation policies intended to minimize health and economic impacts. The model aims to quantify the major costs of an epidemic, including fatalities, economic disruptions, and other social effects. It is also designed to be adaptable to evolving circumstances, such as new virus variants, and to represent heterogeneous effects across population subgroups.

Our framework is built upon a discrete-time version of the classic SIRD epidemiological model Liu et al., 1987, which we extend to incorporate mitigation strategies—such as mask mandates, travel restrictions, quarantines, and lockdowns—and their impacts on infection, recovery, and mortality rates. The model also accounts for the direct implementation costs of these policies, as well as broader economic and opportunity costs. Although these effects are distributed across multiple entities, we consider total regional costs. While many disease progression and cost parameters remain uncertain, we currently use deterministic estimates; future work could incorporate stochastic elements. The model also supports sensitivity and scenario analysis to assess the implications of various parameter settings.

This paper is organized as follows. In the next section, we review the literature on epidemic modeling and policy impact analysis. Section 3 presents the base model, incorporating policy decision variables aimed at mitigating severe outcomes. This model is broadly applicable to various epidemics with minimal adjustments. In Section 4, we introduce a Lagrangian decomposition-based method to compute a lower bound for evaluating solution quality. Section 5 develops a heuristic derived from the Lagrangian approach and introduces two practical solution methods. In Section 6, we conduct computational experiments to evaluate the heuristic's performance and provide extensive numerical results on policy effectiveness. The final section summarizes key findings and outlines avenues for future research.

## 2 Literature Review

The outbreak of the COVID-19 pandemic caught governments, healthcare systems, and societies around the globe off guard, revealing substantial gaps in preparedness and crisis response strategies. As the virus spread rapidly and unpredictably, policymakers were thrust into the difficult position of crafting and implementing urgent policies without the benefit of precedent or clear evidence. This global health emergency spurred a surge in academic research aimed at understanding and managing both the health and economic impacts of the crisis.

A comprehensive survey by Brodeur, Abel et al., 2021 summarizes the economic consequences of the pandemic and examines governmental responses across various countries. Their work highlights the variation in policy effectiveness and the profound economic disruptions induced by lockdowns, business closures, and mobility restrictions. Similarly, Chakraborty, Abhijit et al., 2020 investigates the impact of different government interventions on the trajectory of the epidemic. Their study, grounded in epidemiological modeling, sheds light on the potential effectiveness of mitigation strategies such as social distancing and quarantine.

Despite these efforts, the literature reveals that many of the proposed models and strategies are tailored to specific regional or national contexts, limiting their applicability to broader epidemic settings. Most approaches lack the flexibility required to inform policy decisions for future pandemics with different characteristics or resource constraints.

To address this gap, several researchers have turned to modeling tools that combine epidemiology with economics to explore the trade-offs in pandemic response strategies. For instance, Hur, Sewon, 2023 introduce a life cycle–economic epidemiology framework that evaluates when policies like stay-at-home subsidies are preferable to more stringent lockdowns. This approach provides a structured way to assess the welfare implications of different interventions across time and population segments.

Agent-based simulations have also been used to evaluate dynamic policies. Blakely, Tony et al., 2021 apply such a model to the Australian context, deriving policy insights based on simulated infection patterns and intervention outcomes. Their work demonstrates how agent-based methods can capture heterogeneous behaviors and transmission dynamics at a granular level, offering nuanced guidance for public health planning.

Fu, Yuting et al., 2022 contribute to the discourse by examining optimal lockdown strategies, particularly in the context of vaccine deployment. Their model seeks to balance the dual objectives of minimizing economic disruption and reducing mortality, offering guidance for dynamically adjusting policy intensity as vaccine coverage increases.

Parameter estimation plays a critical role in the construction of reliable epidemic models. Studies such as those by Köhler, Johannes et al., 2021, Yousefpour and Amin and Hadi Jahanshahi and Stelios Bekiros, 2020, and Tsay, Calvin et al., 2020 focus on estimating key epidemiological parameters from real-world data in Germany, China, and the US, respectively. These parameterized models are then employed within optimal control frameworks, typically to evaluate social distancing strategies. Their findings reinforce the value of using real data to calibrate models for more accurate policy analysis.

Although these contributions are valuable, they often emphasize specific interventions or localized conditions. To the best of our knowledge, our work represents the first attempt to use large-scale optimization techniques to solve a generalized epidemic mitigation problem. Our model accommodates a wide range of policies and outcomes, providing a versatile tool for policymakers seeking robust strategies under uncertain and evolving conditions.

# 3 Model and Preliminaries

This study introduces a large-scale Mixed Integer Nonlinear Programming (MINLP) model that extends the classical SIRD epidemiological framework into a discrete-time setting. The central aim of this formulation is to explore the inherent trade-offs between the costs incurred and the effectiveness achieved through the implementation of various mitigation policies.

In the context of an epidemic such as COVID-19, policies can differ significantly in both cost and impact. For instance, while a basic masking mandate represents a low-cost intervention, its effectiveness in controlling disease transmission is generally lower than that of a more comprehensive and costly lockdown. The model captures these distinctions by quantifying both the direct and indirect effects of each policy on the spread of the disease and associated economic outcomes.

The solutions obtained from our MINLP model provide optimized policy recommendations that seek to minimize the total cost burden. This includes not only the direct costs of implementing public health measures but also the indirect opportunity costs borne by the population and the expenditures incurred by the healthcare system. By evaluating different combinations of interventions, the model facilitates informed decision-making that balances economic constraints with public health objectives.

## 3.1 SIRD Model with Variable-Cost Interventions

In this subsection, we first define the parameters and variables used throughout this paper. Later, we propose a mixed-integer non-linear programming (MINLP) model incorporating the costs that change as the epidemic progresses. In subsequent sections, we modify this objective to make the analysis more tractable.

### 3.1.1 Parameters

We consider a setting where there are N total individuals, of whom I0 are initially infected. There are m interventions to consider, each with (up to) $n$ intensity levels. We build our model for a planning horizon of T days.

- Let $A_{ij}$ denote the fixed cost of implementing policy $i$ at level $j$ at time $t$.
- Let $B_{ijt}$ denote the *switching* cost of implementing policy $i$ at level $j$ at time $t$ (only incurred if policy not implemented in the previous period).
- Let $C_{ijt}$ denote the per-susceptible-individual cost of implementing policy $i$ at level $j$ at time $t$.
- Let $C_{infection}$ and $C_{death}$ denote the costs associated with an individual being infected in a given period and a single individual losing their life due to disease, respectively.
- Let $K_I$ correspond to the infection rate such that the number of new infections is proportional to $K_I$ multiplied by the number of interactions between susceptible and infected individuals, modeled as the product of the sizes of those populations.
- Let $K_R$ and $K_D$ denote the proportion of infected individuals who recover and die in each period, respectively.
- Let $P_{ijt}$ denote the factor by which new infections are decreased in period $t$ due to implementing policy $i$ at level $j$. In this model, these factors are independent of one another should multiple policies be implemented simultaneously.

### 3.1.2 Decision Variables

In our model, two binary variables are used to calculate the policy implementing costs and the cumulative efficacy of a policy.

- Let

$$y_{ijt} = \begin{cases} 1 & : \text{policy } i \text{ is implemented at level } j \text{ in time period } t \\ 0 & : \text{otherwise} \end{cases} \tag{1}$$

- Let

$$z_{ijt} = \begin{cases} 1 & : \text{policy } i \text{ is implemented at level } j \text{ in time period t, but not } t-1 \\ 0 & : \text{otherwise} \end{cases} \tag{2}$$

### 3.1.3 State Variables

Let $S_t$, $I_t$, $R_t$, $d_t$, and $D_t$ denote the population of individuals at time $t$ who are Susceptible, Infected, Recovered, dying (in the current period), and Dead (cumulatively), respectively. These values depend on the interventions applied. These variables are used to calculate the healthcare costs of a policy.

Let $P_t$ denote the cumulative factor by which new infections are decreased between periods $t-1$ and $t$. That is,

$$P_t = \prod_{\substack{i,j \text{ s.t.} \\ \text{policy } i \text{ used} \\ \text{at level } j \\ \text{in period } t}} P_{ijt} \tag{3}$$

### 3.1.4 Model Formulation: Epidemic Mitigation Problem (EMP)

The objective function, as described in (4), comprises the summation of costs associated with policy interventions across all time periods, as well as the costs attributed to the disease's impact, including lost productivity and resources, throughout the entire time horizon.

To model the progression of the disease, the constraints (5), (6), (7), (8), and (9) delineate the SIRD compartment subpopulations while considering infection-reduction factors denoted as $P_t$ for each time period $t = 1, \cdots, T$. Additionally, constraint (10) captures the multiplicative effects arising from concurrent interventions applied within the same period, as outlined in (3). The logical constraint that, in each period, at most one level from each policy is employed is enforced by equation (11), and (12) ensures that the variables yijt align with the binary definition given in (15).

$$\min_{y,P,S,I,R,D,d} \quad \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t=1}^{T} A_{ijt}y_{ijt} + B_{ijt}z_{ijt} + C_{ijt}S_t y_{ijt} + \sum_{t=1}^{T} C_{infection} I_t + C_{death} d_t \tag{4}$$

$$\begin{aligned}
\text{s.t} \quad & S_t && = S_{t-1} - K_I \cdot P_t \cdot S_{t-1} \cdot I_{t-1} && t \in \{2,\ldots,T\} && (5)\\
& I_t && = I_{t-1} + K_I \cdot P_t \cdot S_{t-1} \cdot I_{t-1} - K_R \cdot I_{t-1} - K_D \cdot I_{t-1} && t \in \{2,\ldots,T\} && (6)\\
& R_t && = R_{t-1} + K_R \cdot I_{t-1} && t \in \{2,\ldots,T\} && (7)\\
& d_t && = K_D \cdot I_{t-1} && t \in \{2,\ldots,T\} && (8)\\
& D_t && = D_{t-1} + d_t && t \in \{2,\ldots,T\} && (9)\\
& P_t && = \prod_{i=1}^{m}\prod_{j=1}^{n}(1 - y_{ijt} + P_{ijt} \cdot y_{ijt}) && t \in \{1,\ldots,T\} && (10)\\
& \sum_{j=1}^{n} y_{ijt} \leq 1 && && i,t && (11)\\
& y_{ijt} && \in \{0,1\} && i,j,t && (12)
\end{aligned}$$

$$\begin{aligned}
& z_{ijt} && \geq y_{ij(t)} - y_{ij(t-1)} \quad (\text{let } y_{ij0} = 0) && i,j,t && (13)\\
& 0 \leq z_{ijt} \leq 1 && && i,j,t \geq 1 && (14)\\
& I_1 && = I_0 \\
& S_1 && = N - I_0 \\
& D_1 && = 0 \\
& R_1 && = 0 \\
& d_1 && = 0
\end{aligned}$$

**Theorem 1. The (EMP)** is NP-*Complete when* $K_I \cdot S_0 \cdot I_0 < 1$ (Appendix - A)

*Proof.* Let's consider a very simple instance of our original problem with no per-susceptible-individual cost of implementing policy ($C_{ijt} = 0$) nor switching cost ($B_{ijt} = 0$) and only one possible intervention ($m = 1$), one level of intensity ($n = 1$). Let $p$ be the factor by which new infections are decreased in period $t$ due to implementing this policy (i.e. $P_{ijt} = p$).

Let, the decision variable $y_t$ be defined as follows:

$$y_t = \begin{cases} 1 & : \text{Policy is implemented in time period } t \\ 0 & : \text{otherwise} \end{cases} \tag{15}$$

The **(EMO)** reduces to a standard knapsack problem of the form :

$$\begin{aligned} \underset{y}{\text{minimize}} \quad & \sum\nolimits_{t=1}^{T} a_t \cdot y_t + \sum\nolimits_{t=1}^{T} c_t \cdot I_t \\ \text{s.t.} \quad & S = K_S \cdot Y \\ & I = K_I \cdot Y \\ \textit{where} \quad & \begin{bmatrix} S_1 \\ S_2 \\ \vdots \\ S_T \end{bmatrix}, \quad I = \begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_T \end{bmatrix}, \quad Y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_T \end{bmatrix} \end{aligned}$$

$K_S$ and $K_I$ are $T \times T$ matrices with entries that are functions of $p$, $K_I$ , $K_R$, and $K_D$

The idea behind this reduction is to express the number of infected ($I_t$) and susceptible ($S_t$) individuals at each period in terms of the rate of infection ($K_I$) and the initial state of the epidemic ($S_0$ and $I_0$) and notice that all the non-linear terms have negligible coefficients.

The details of the reduction can be found in the Appendix.

□

# 4 Problem Decomposition and Lower Bound

Lagrangian relaxation can be applied to tackle the entire **(EMP)** problem by relaxing the constraints defined in (7) and, instead, by introducing a penalty into the objective function using dual multipliers. This relaxation partitions the problem into two computationally more manageable subproblems. We progressively enhance the lower bound by iteratively updating the multipliers through gradient ascent. To begin with, our focus is on a model variant in which interventions are characterized by a fixed cost that is directly proportional to the total population size ($N$), as depicted in (16). In essence, the costs are now time invariant. This characteristic facilitates partitioning the Lagrangian minimization problem into two distinct subproblems.

Replace the objective (4) in the mathematical program formulation of the **(EMP)** model with

$$\min_{y,P,S,I,R,D,d} \quad \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t=1}^{T} A_{ijt}y_{ijt} + B_{ijt}z_{ijt} + C_{ijt}\cdot S_t \cdot y_{ijt} + \sum_{t=1}^{T} C_{infection}\cdot I_t + C_{death}\cdot d_{t_{(4)}}$$

$$\Downarrow$$

$$\min_{y,P,S,I,R,D,d} \quad \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t=1}^{T} A_{ijt}y_{ijt} + B_{ijt}z_{ijt} + C_{ijt}\cdot N \cdot y_{ijt} + \sum_{t=1}^{T} C_{infection}\cdot I_t + C_{death}\cdot d_t \tag{16}$$

Then, we transform constraint (10) using logarithms:

$$\ln P_t = \sum_{i=1}^{m}\sum_{j=1}^{n} \ln(1 - y_{ijt} + P_{ijt}y_{ijt}), \quad t = 1,\ldots,T \tag{10-log}$$

Next, we remove this constraint from **(EMO)** and relaxed the objective (16) via multipliers $\lambda_t, t = 1,\cdots,T$ to obtain a relaxed minimization problem:

$$\min_{y,P,S,I,R,D,d} \quad [\text{Objective (16)}] + \sum_{t=1}^{T}\lambda_t\left(\sum_{i=1}^{m}\sum_{j=1}^{n}\ln(1 - y_{ijt} + P_{ijt}y_{ijt}) - \ln P_t\right) \tag{17}$$

$$\text{s.t.} \quad 0 \le P_t \le 1 \tag{18}$$

Constraints from **(EMO)** except for (10).

An optimal solution to the lagrangian relaxation (17) serves as a lower bound for the optimal value of the full problem **(EMP)**. To enforce the logical constraints that the policy effectiveness factors are between 0 and 1; we introduce an additional constraint (18). It is worth noting that in a numerical implementation, it may actually be preferable to pre-compute a reasonable lower bound on $\underline{P_t} \in (0,1)$ and constraint $\underline{P_t} \le P_t \le 1$ because the logarithm in (17) is undefined for $P_t = 0$. By iteratively solving the augmented problem (17) and then using subgradient ascent to update $\lambda_t$ for all $t = 1,\cdots,T$, we obtain increasingly tighter lower bounds on the optimal value for the full problem **(EMP)**.

Note that the augmented problem (17) can be decomposed into two minimization prob-

lems with optimal values $L_1(\lambda)$ and $L_2(\lambda)$:
$L_1(\lambda)$ is the solution to

$$\begin{aligned}
\min_{y} \quad & \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t=1}^{T}\left[A_{ijt}\cdot y_{ijt} + B_{ijt}z_{ijt} + C_{ijt}\cdot N\cdot y_{ijt} + \lambda_t \ln(1 - y_{ijt} + P_{ijt}y_{ijt})\right] \\
\text{s.t.} \quad & \sum_{j=1}^{n} y_{ijt} \le 1 && i,t \quad (11)\\
& y_{ijt} \in \{0,1\} && i,j,t \quad (12)\\
& z_{ijt} \ge y_{ij(t)} - y_{ij(t-1)} && i,j,t \quad (13)\\
& 0 \le z_{ijt} \le 1 && i,j,t \ge 1 \quad (14)
\end{aligned}$$

$L_2(\lambda)$ is the solution to

$$\begin{aligned}
\min_{P,S,I,R,D,d} \quad & \sum_{t=1}^{T}\left[C_{\text{infection}}\cdot I_t + C_{\text{death}}\cdot d_t - \lambda_t \ln P_t\right] \\
\text{s.t.} \quad & S_t = S_{t-1} - K_I\cdot P_t\cdot S_{t-1}\cdot I_{t-1} && t\in\{2,\dots,T\} \quad (5)\\
& I_t = I_{t-1} + K_I\cdot P_t\cdot S_{t-1}\cdot I_{t-1} - K_R\cdot I_{t-1} - K_D\cdot I_{t-1} && t\in\{2,\dots,T\} \quad (6)\\
& R_t = R_{t-1} + K_R\cdot I_{t-1} && t\in\{2,\dots,T\} \quad (7)\\
& d_t = K_D\cdot I_{t-1} && t\in\{2,\dots,T\} \quad (8)\\
& D_t = D_{t-1} + d_t && t\in\{2,\dots,T\} \quad (9)\\
& P_t \le 1 \quad (18)\\
& I_1 = I_0\\
& S_1 = N - I_0\\
& D_1 = 0\\
& R_1 = 0\\
& d_1 = 0
\end{aligned}$$

This problem has no integer constraints and can be solved by any nonlinear programming software.

To increase the tightness of the bound in the gradient-ascent step for the multipliers $\lambda$, where $\lambda^+$ the vector of multipliers at a subsequent iteration, we use the updating rule:

$$\lambda^+ = \lambda + \gamma(\nabla L_1(\lambda) + \nabla L_2(\lambda)),$$

i.e.

$$\lambda^+ = \lambda_t + \gamma \cdot \left( \sum_{i=1}^{m} \sum_{j=1}^{n} \ln(1 - y_{ijt} + P_{ijt} y_{ijt}) - \ln P_t \right). \tag{19}$$

# 5 Heuristics and Upper Bounds

The previous section has established that a significantly simplified version of our overarching problem (DMO) is NP-Hard. This makes it evident that commercial solvers are inadequate for providing optimal solutions to real-sized instances of the problem. We confirm this in our computational experiments. Given this computational complexity, we explored explore heuristic approaches to solve the (EMP). These heuristics aim to provide practical solutions to the (EMP) while simultaneously serving as upper bounds for the optimal solution the (EMP).

## 5.1 Restriction Heuristic

Within our comprehensive DMO problem, we require a large set of decision variables, which grow with the increase in the planning horizon, giving us decision variables of the order O(nmT). In most practical scenarios, policymakers plan for a minimum horizon of 30 days or more. This results in an overwhelming number of decision variables surpassing commercial optimization solvers' capabilities.

Additionally, from a practical perspective, it is highly unlikely that policymakers would implement a new policy daily rather, they change a policy once in a while and continue to implement the same policy till the next change is considered. Therefore, we consider reducing the complexity of our problem by confining decision-making to a subset of days within the planning horizon, which we denote as $T^*$.

Our heuristic shares the same constraints as the (EMP), but differs in its objective function, which is as follows:

$$\min_{y,P,S,I,R,D,d} \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{t\in T} A_{ijt}y_{ijt} + B_{ijt}z_{ijt} + C_{ijt}\cdot y_{ijt} + \sum_{t\in T} C_{infection}\cdot I_t + C_{death}\cdot d_t \tag{20}$$

$$\begin{aligned} &\text{Constraints from } (\textbf{EMP}) \\ &y_{ijt} \in \{0,1\} && i,j,t \in T^* \subset T && (21) \\ &y_{ijt} = y_{ij(t-1)} && i,j,t \notin T^* \subset T && (22) \end{aligned}$$

The performance of our heuristic, concerning both its optimality and runtime, is closely tied to the frequency of decision-making over the planning horizon. A higher decision frequency results in an increased number of decision variables, leading to longer computation times, but it also brings us closer to reaching an optimal solution. In practice, information about decision frequency is typically sourced from surveys and insights shared by policymakers. Section 6 discusses the performance of our heuristics for various decision-making frequencies.

# 6 Computational Results

To understand the computational difficulty in solving our problem exactly using commercial software, we tried solving our **(EMP)** with different parameters using several software, of which the BARON solver was the most efficient. Table 1 shows the time BARON took to give an optimal solution to our MINLP for different problem instances for a planning horizon of 30 days, 60 days, and 90 days. The details of the parameter choices are provided in the Appendix. From Table 1, we can see that the time BARON takes to solve our MINLP increases exponentially with the duration of the planning horizon. In a few instances, BARON does not terminate even for as short as a 60-day planning horizon. These computational results further warrant the need for a heuristic to solve our **(EMP)** in a reasonable amount of time.

[1] [2] [3]

As discussed in the previous section, our heuristic's performance depends on the frequency at which new policy decisions are made during the planning horizon. Table 2 shows the objective value and the heuristic's run time for various decision-making frequencies for an

[1] All values reported in this table are in minutes.

[2] "NA" represents that BARON did not terminate even after 24 hours of runtime.

[3] Description of each instance is given in the Appendix.

| Table 1 | | | |
|---|---|---|---|
| | 30 days | 60 days | 90 days |
| Instance – 1 | 2 | 12 | 34 |
| Instance – 2 | 5 | 39 | 116 |
| Instance – 3 | 1 | 11 | 149 |
| Instance – 4 | 10 | 109 | 451 |
| Instance – 5 | 5 | NA | NA |
| Instance – 6 | 30 | NA | NA |
| Instance – 7 | 2 | 55 | NA |

illustrative problem instance (instance-3 in Appendix) of a 90-day planning horizon. A frequency of 4 represents that the policymakers can change the existing policy only on days 1, 16,31, and 46 of the 60-day planning horizon. The decision on each of these four days will be fixed for the next 15 days until the next decision is made.

| Table 2 | | |
|---|---|---|
| Number of decisions over a 90 day planning horizon | Cost | Runtime |
| 3 | 11273811 | 4 |
| 5 | 11273811 | 15 |
| 9 | 11273811 | 49 |
| 10 | 11273811 | 67 |
| 30 | 11121990 | 101 |
| 90 | 11021981 | 147 |

[4] [5]

From Table 2, we can see that the runtime increases as the frequency of decision-making increases along with the accuracy of the heuristic solution. This table suggests that we can maintain the quality of the heuristic by selecting a decision-making frequency with lesser runtime.

Tables 3, 4, and 5 summarize the comparative performance of our heuristic with the optimal solution obtained from BARON and the Lagrangian lower bound, as outlined in Section 4, across planning horizons of 30, 60, and 90 days, respectively. We tried our heuristic at different frequencies and picked one with a reasonable solution in an acceptable run time. A description of each instance is given in the Appendix.

In our experiment results, our heuristic solution is mostly within 5 percent of the lagrangian lower bounds. In the case of a 30-day planning horizon, the worst-case performance of our heuristic was 3.63 percent relative to BARON's optimal solution. In the case of a 60-day planning horizon, the worst-case performance of our heuristic was 5.22 percent of the Lagrangian lower bound, while the worst-case performance for a 90-day planning horizon was 3.80 percent of the Lagrangian lower bound.

| Table 3 - 30 days planning horizon | | | | | |
|---|---|---|---|---|---|
| Instances | Heuristic Upper Bound (HUB) | BARON's Optimal (BO) | Lagrangian Lower Bound (LLB) | 100*(HUB - BO)/BO | 100*(HUB - LLB)/LLB |
| Instance – 1 | 3691380 | 3591371 | 3514188 | 2.78 | 2.20 |
| Instance – 2 | 4078436 | 3935556 | 3865842 | 3.63 | 1.80 |
| Instance – 3 | 3804124 | 3704115 | 3632046 | 2.70 | 1.98 |
| Instance – 4 | 3216273 | 3116264 | 3049704 | 3.21 | 2.18 |
| Instance – 5 | 10274402 | 10072261 | 9755596 | 2.01 | 3.24 |
| Instance – 6 | 24814380 | 24784185 | 24513650 | 0.12 | 1.10 |
| Instance – 7 | 3969552 | 3869543 | 3826080 | 2.59 | 1.14 |

[4] All values reported in this Table are in minutes

[5] A description of Instance - 3 is given in the Appendix

**Table 4 - 60 days planning horizon**

| Instances | Heuristic Upper Bound (HUB) | BARON's Optimal (BO) | Lagrangian Lower Bound (LLB) | 100*(HUB - BO)/BO | 100*(HUB - LLB)/LLB |
|---|---|---|---|---|---|
| Instance – 1 | 7350017 | 7065437 | 6985137 | 4.02 | 5.22 |
| Instance – 2 | 8366226 | 8215604 | 8213950 | 1.83 | 1.85 |
| Instance – 3 | 7668017 | 7436231 | 7363505 | 3.12 | 4.14 |
| Instance – 4 | 6039953 | 5905291 | 5856708 | 2.28 | 3.13 |
| Instance – 5 | 15969367 | NA | 15951365 | NA | 0.11 |
| Instance – 6 | 477484128 | NA | 477435249 | NA | 0.01 |
| Instance – 7 | 8226272 | 8044481 | 7985381 | 2.26 | 3.02 |

[6]

**Table 5 - 90 days planning horizon**

| Instances | Heuristic Upper Bound (HUB) | BARON's Optimal (BO) | Lagrangian Lower Bound (LLB) | 100*(HUB - BO)/BO | 100*(HUB - LLB)/LLB |
|---|---|---|---|---|---|
| Instance – 1 | 10684305 | 10351742 | 10293480 | 3.21 | 3.80 |
| Instance – 2 | 12843333 | 12700593 | 12656582 | 1.12 | 1.48 |
| Instance – 3 | 11273811 | 11021981 | 10955420 | 2.28 | 2.91 |
| Instance – 4 | 8712383 | 8660430 | 8597906 | 0.60 | 1.33 |
| Instance – 5 | 20241880 | NA | 20206575 | NA | 0.18 |
| Instance – 6 | 264647207 | NA | 264617312 | NA | 0.01 |
| Instance – 7 | 12499440 | 12328581 | 12468464 | 1.39 | 0.25 |

[7]

# 7 Summary and Future Work

In this paper, we considered the intricate problem of assisting policymakers in making informed binary decisions regarding the implementation of various epidemic control policies over a specified planning horizon. To approach this challenge, we proposed a comprehensive Mixed Integer Nonlinear Programming (MINLP) framework that models the trade-offs between minimizing total implementation costs and adhering to constraints that govern the dynamics of epidemic transmission.

Our contributions include both theoretical insights and computational results. We established that the resulting MINLP formulation is computationally intractable due to its large scale and combinatorial complexity. To address this, we employed a decomposition-based strategy that partitions the original problem into two more tractable subproblems. Solving these subproblems to optimality allowed us to generate a Lagrangian lower bound for the original MINLP. In addition, we developed a practical and intuitive heuristic approach whose

[6] "NA" represents that BARON did not terminate even after 24 hours of runtime

[7] "NA" represents that BARON did not terminate even after 24 hours of runtime

performance was benchmarked against the Lagrangian lower bound and the optimal solution produced by a commercial solver.

Looking forward, our research agenda includes several key extensions. One priority is to refine and enhance our heuristic techniques to narrow the optimality gap relative to the Lagrangian bound. In particular, we are exploring dynamic heuristics that iteratively construct the set $T^*$ based on the evolving trajectory of the epidemic. This dynamic extension aims to improve the responsiveness and adaptability of our approach, enabling policymakers to revise decisions in real time as new information becomes available.

Another important direction is the integration of vaccination strategies into our model. This enhancement will allow us to evaluate both the timing and allocation of vaccines, and to determine optimal proactive deployment policies that maximize epidemiological and economic benefits.

We also plan to expand the scope of our model to address operational decisions in healthcare systems, such as optimal staffing levels and timely procurement of medical equipment. These applications are critical for strengthening the resilience and preparedness of hospitals and public health infrastructure during epidemics.

Finally, we are in the process of collecting and curating detailed COVID-19 data from Los Angeles County. Applying our model to this dataset will enable a series of "what-if" simulations, providing actionable insights and deepening our understanding of how different intervention strategies might influence epidemic outcomes under varying conditions.

# Appendix

# A Proof of Theorem 1

Below is the detailed proof of Theorem 1.

*Proof.* We first express the number of susceptible individuals $S_t$ and infected individuals $I_t$ in terms of $K_I$, $S_0$, and $I_0$.

$$S_2 = S_1 - K_I \cdot (1 - y_2 + p \cdot y_2) \cdot S_1 \cdot I_1 = K_{21} + K_{22} \cdot y_2$$

*where,*

$$K_{21} = S_0 - K_I \cdot S_0 \cdot I_0;$$

$$K_{22} = (1 - p) \cdot K_I \cdot S_0 \cdot I_0$$

*Note,*

$$K_{21} < S_0 \quad \text{and} \quad K_{22} < K_I \cdot S_0 \cdot I_0$$

$$I_2 = I_1 + K_I \cdot (1 - y_2 + p \cdot y_2) \cdot S_1 \cdot I_1 - K_R \cdot I_1 - K_D \cdot I_1 = K'_{21} + K'_{22} \cdot y_2$$

*where,*

$$K'_{21} = (1 - K_R - K_D) \cdot I_0 + K_I \cdot S_0 \cdot I_0;$$

$$K'_{22} = (p - 1) \cdot K_I \cdot S_0 \cdot I_0$$

*Note,*

$$K'_{21} \approxeq K_I \cdot I_0 \cdot S_0 \quad \text{and} \quad K'_{22} < K_I \cdot S_0 \cdot I_0$$

$$S_3 = S_2 - K_I \cdot (1 - y_3 + p \cdot y_3) \cdot S_2 \cdot I_2$$

$$= K_{31} + K_{32} \cdot y_2 + K_{33} \cdot y_3 + K_{34} \cdot y_2 \cdot y_3 + K_{35} \cdot y_2^2 + K_{36} \cdot y_2^2 \cdot y_3$$

*where,*

$$K_{31} = K_{21} - K_I \cdot K_{21} \cdot K'_{21}$$

$$K_{32} = K_{22} - K_I \cdot (K_{22} \cdot K'_{21} + K_{21} \cdot K'_{22})$$

$$K_{33} = (1 - p) \cdot K_I \cdot K_{21} \cdot K'_{21}$$

$$K_{34} = (1 - p) \cdot K_I \cdot (K_{22} \cdot K'_{21} + K_{21} \cdot K'_{22})$$

$K_{35} = K_I \cdot K_{22} \cdot K'_{22}$

$K_{36} = (1 - p) \cdot K_I \cdot K_{22} \cdot K'_{22}$

*Note,*

$K_{31} < S_0$

$K_{32} < K_I \cdot S_0 \cdot I_0$

$K_{33} < K_I \cdot S_0 \cdot I_0$

$K_{34} \ll 1$

$K_{35} \ll 1$

$K_{36} \ll 1$

$I_3 = I_2 + K_I \cdot (1 - y_3 + p \cdot y_3) \cdot S_2 \cdot I_2 - K_R \cdot I_2 - K_D \cdot I_2 = K'_{31} + K'_{32} \cdot y_2 + K'_{33} \cdot y_3 + K'_{34} \cdot y_2 \cdot y_3$
$+ K'_{35} \cdot y_2^2 + K'_{36} \cdot y_2^2 \cdot y_3$

*where,*

$K'_{31} = K'_{21} + K_I \cdot K_{21} \cdot K'_{21}$

$K'_{32} = K'_{22} + K_I \cdot (K_{22} \cdot K'_{21} + K_{21} \cdot K'_{22})$

$K'_{33} = -(1 - p) \cdot K_I \cdot K_{21} \cdot K'_{21}$

$K'_{34} = -(1 - p) \cdot K_I \cdot (K_{21} \cdot K'_{22} + K_{22} \cdot K'_{21})$

$K'_{35} = K_I \cdot K_{22} \cdot K'_{22}$

$K'_{36} = -(1 - p) \cdot K_I \cdot K_{22} \cdot K'_{22}$

*Note,*

$K'_{31} \approxeq K_I \cdot I_0 \cdot S_0$

$K'_{32} < K_I \cdot S_0 \cdot I_0$

$K'_{33} < K_I \cdot S_0 \cdot I_0$

$K'_{34} \ll 1$

$K'_{35} \ll 1$

$K'_{36} \ll 1$

We notice that the coefficients of all the non-linear terms are negligible, which allows us to write the number of infected ($I_t$) and susceptible ($S_t$) individuals at each period as a linear function of the decision variables $y_t$. Replacing $S_t$ and $I_t$ in our original problem will give us a knap-sack problem which is NP-Complete. □

# B Time independent parameters

In this work, we used time-independent parameters. The values of the parameters reported are the same for the entirety of the planning horizon.

Below is the table for the different costs of each of the interventions

| **Cost and Efficacy Parameters** | | | | | | |
|---|---|---|---|---|---|---|
| Cost and Efficacy | policy-1, level-LOW | policy-1, level-HIGH | policy-2, level-LOW | policy-2, level-HIGH | policy-3, level-LOW | policy-3, level-HIGH |
| Setup Cost | 9 | 18 | 90 | 180 | 900 | 1800 |
| Policy Cost | 1 | 2 | 10 | 20 | 100 | 200 |
| Switching Cost | 2 | 4 | 20 | 40 | 200 | 400 |
| Infection Cost | 3000 | 3000 | 3000 | 3000 | 3000 | 3000 |
| Death Cost | 8000 | 8000 | 8000 | 8000 | 8000 | 8000 |
| P | 0.05 | 0.01 | 0.0005 | 0.0001 | 0.000005 | 0.000001 |
| P* | 0.07 | 0.03 | 0.0007 | 0.0003 | 0.000007 | 0.000003 |

Below is the table for the parameters for each of the instances used in our computational experiments.

| **Parameter values for each instance** | | | | | | |
|---|---|---|---|---|---|---|
| Instance | Infection Rate (Ki) | Recovery Rate (Kr) | Death Rate (Kd) | Initial Infections (I0) | Population Size (N) | Policy-Efficacy |
| Instance–1 | 0.000009 | 0.01 | 0.05 | 10 | 100000 | P |
| Instance–2 | 0.000013 | 0.01 | 0.05 | 10 | 100000 | P |
| Instance–3 | 0.000009 | 0.002 | 0.05 | 10 | 100000 | P |
| Instance–4 | 0.000009 | 0.01 | 0.1 | 10 | 100000 | P |
| Instance–5 | 0.000009 | 0.01 | 0.05 | 100 | 100000 | P |
| Instance–6 | 0.000009 | 0.01 | 0.05 | 10 | 500000 | P |
| Instance–7 | 0.000009 | 0.01 | 0.05 | 10 | 100000 | P* |

$P$ and $P^*$ are as per the previous table

# References


Blakely, Tony et al. (2021). Determining the optimal COVID-19 policy response using agent-based modelling linked to health and cost modelling: Case study for Victoria, Australia. *MedRxiv*, 2021–01.

Brodeur, Abel et al. (2021). A literature review of the economics of COVID-19. *Journal of Economic Surveys*, *35*(4), 1007–1044.

Chakraborty, Abhijit et al. (2020). Analyzing COVID-19 data using SIRD models. *medRxiv*, 2020–05.

Fu, Yuting et al. (2022). Optimal lockdown policy for vaccination during COVID-19 pandemic. *Finance Research Letters*, *45*, 102123.

Hur, Sewon. (2023). The distributional effects of COVID-19 and optimal mitigation policies. *International Economic Review*, *64*(1), 261–294.

Köhler, Johannes et al. (2021). Robust and optimal predictive control of the COVID-19 outbreak. *Annual Reviews in Control*, *51*, 525–539.

Liu, Wei-min, Hethcote, H. W., & Levin, S. A. (1987). Dynamical behavior of epidemiological models with nonlinear incidence rates. *Journal of Mathematical Biology*, *25*, 359–380.

Tsay, Calvin et al. (2020). Modeling, state estimation, and optimal control for the US COVID-19 outbreak. *Scientific Reports*, *10*(1), 10711.

Yousefpour and Amin and Hadi Jahanshahi and Stelios Bekiros. (2020). Optimal policies for control of the novel coronavirus disease (COVID-19) outbreak. *Chaos, Solitons & Fractals*, *136*, 109883.