OPTIMIZING VACCINE DISTRIBUTION IN DEVELOPING COUNTRIES UNDER NATURAL DISASTER RISK

Bonn Kleiford Seranilla
Luxembourg Center for Logistics and Supply Chain Management
University of Luxembourg
Luxembourg City, Luxembourg
bonn-kleiford.seranilla@uni.lu

Nils Löhndorf
Luxembourg Center for Logistics and Supply Chain Management
University of Luxembourg
Luxembourg City, Luxembourg
nils.loehndorf@uni.lu

ABSTRACT

For many developing countries, COVID-19 vaccination roll-out programs are not only slow but vaccination centers are also exposed to the risk of natural disaster, like flooding, which may slow down vaccination progress even further. Policy-makers in developing countries therefore seek to implement strategies that hedge against distribution risk in order for vaccination campaigns to run smoothly and without delays. We propose a stochastic-dynamic facility location model that allows policy-makers to choose vaccination facilities while accounting for possible facility failure. The model is a multi-stage stochastic variant of the classic facility location problem where disruption risk is modelled as a binary multivariate random process - a problem class that has not yet been studied in the extant literature. To solve the problem, we propose a novel approximate dynamic programming algorithm which trains the shadow price of opening a flood-prone facility on historical data, thereby alleviating the need to fit a stochastic model. We trained the model using rainfall data provided by the local government of several major cities in the Philippines which are exposed to multiple flooding events per year. Numerical results demonstrate that the solution approach yields approximately 30-40% lower cost than a baseline approach that does not consider the risk of flooding. Recommendations based on this model were implemented following a collaboration with two large cities in the Philippines which are exposed to multiple flooding events per year.

***K*eywords** Healthcare Facility Location Problem; Humanitarian Logistics; Stochastic Programming; Approximate Dynamic Programming

1 Introduction

The coronavirus pandemic of 2019 (COVID-19) demonstrated the vulnerability of public health safety, with approximately 578 million cases and 6.4 million deaths worldwide as of August 2022 (1). The development of vaccines against the virus substantially altered the course of the pandemic. As of August 2022, 12.3 billion doses of vaccines have been administered all over the world (1). The distribution of these vaccines posed major supply chain challenges for pharmaceutical supply chains and local governments. Since most COVID-19 vaccines are produced and distributed in developed countries, only limited and intermittent vaccine supply reached non-producing, developing countries where vaccination roll-outs are already expected to take longer than in developed countries.

The roll-out of COVID-19 vaccines was further delayed for reasons more than only a restricted supply. Other factors like natural disasters continuously daunt these countries, among them hydro-meteorological events (e.g. flooding, landslides,

and storm surges), seismic events (e.g. earthquakes), and volcanic eruptions. Particularly, the absence of adequate infrastructure, catastrophe emergency preparedness, and response systems makes developing countries more vulnerable to natural disasters (2) and puts vaccination campaigns in peril. For instance, in early December 2021, Typhoon Rai hit central Philippines with catastrophic thunderstorms missing year-end COVID-19 vaccination target due to delayed schedules and expired vaccines (3). Similarly, in January 2022, deadly floods and landslides caused by torrential rains hit the city of Sao Paulo, Brazil which cancelled the scheduled COVID-19 vaccination campaigns (4). These factors drive policy-makers in developing nations to seek decision support for planning for vaccination campaigns as well as to preparing those campaigns against the risk of natural disaster risk to prevent delays and ultimately save lives.

In this article, we report on models and methods developed as part of *Project FALCON*, which is a collaboration of the authors with the local government of Cagayan de Oro in the Philippines, for planning the roll-out of the COVID-19 vaccination campaign. Cagayan de Oro is a large city on the island of Mindanao that experiences multiple torrential rains and flooding events every year (5). Flooding poses a serious threat for successfully orchestrating large-scale vaccination campaigns. The optimal location of vaccination centers and allocation of recipients must therefore account for facility failure due to flooding as well as the resulting cost of relocation and redistribution.

The optimization problem underlying FALCON is a stochastic-dynamic facility location model whose optimal decision policy chooses opening of vaccine facilities while accounting for possible facility failure. The joint risk of facility failure is captured by a multivariate binary random process. As the model entails solution of a high-dimensional stochastic-dynamic program, FALCON uses approximate dynamic programming to optimize the decision policy. Numerical investigations show a reduction of vaccination roll-out operational cost of approximately 30-40% when compared to a greedy approach that ignores the risk of flooding.

The contribution of this work is two-fold: first, the proposed stochastic facility location model as well as the developed solution method are novel to the literature on facility location planning as we further explain below. Second, the work contributes to the practice of vaccine distribution, as Project FALCON provided the scientific background to nudge the central government to accelerate vaccination roll-out. The implementation of Project FALCON led to Cagayan de Oro City hitting the inoculation target and having one of the highest vaccination rates in the Philippines by the end of Year 2021 (6).

Our work advances the extant literature in both methodological and managerial directions. First, the multi-stage variant of the proposed facility location problem has not yet been explored in the literature, specifically the integration of disruption risks due to natural disasters. We close this gap by formulating the problem as a multi-stage stochastic mixed-integer program that captures the risk of natural disaster disruption by a binary random variable that affects the state of each facility across time. We then develop an approximate dynamic programming formulation where the value function of the underlying stochastic-dynamic program is approximated by a function that is linear in the state variables. We show that the optimal upper bound of this approximation only depends on the terms associated with facility state variables. To the best of our knowledge, this is the first time that this problem class is studied.

Furthermore, we propose an algorithm that trains the parameters of the linear value function approximation by minimizing an upper bound on the optimal objective value. Since said parameters behave like shadow prices of the non-anticipatory constraints of the multistage stochastic program, we refer to this algorithm as *shadow price approximation* (SPA). SPA closely resemble (7)'s *cost function approximation* (CFA) - a technique where a parameterized lookahead model provides a policy for solving a stochastic base model. Our approach is simpler as it does not require formulation of a lookahead model but merely adds a small set of parameters to the objective function of each subproblem. Model and solution algorithm make up the computational core of the decision support tool that has been developed for Project FALCON.

Although being motivated by COVID-19 vaccination efforts, our model is adequately general to adapt to other settings. Additionally, given the data-driven nature of the solution method, any other form of disruption can be considered, such as blackouts, hurricanes, draughts, etc. The proposed algorithm also has two major advantages over alternative solution approaches for this problem class: (1) it can be easily integrated with gradient-based and stochastic search methods that are widely used in machine learning and global optimization; (2) the algorithm does not require formulation of a stochastic model but merely needs access to independent time series that can also be real data.

The remainder of the paper is organized as follows: Section 2 presents a review of the extant literature. Section 3 introduces a multi-stage facility location problem under uncertainty of facility failure. Section 4 presents an introduction to the SPA algorithm which optimally solves the multi-stage stochastic-dynamic facility location problem. Section 5 demonstrates the application of the facility location model subsuming the details of integrating rainfall as the random parameter of the problem and solution method to large cities in the Philippines. Section 6 summarizes the results and provides an outlook of the research directions.

2 Literature Review

Choosing vaccination centers is essentially an instance of the well-known facility location problem (FLP). The classic FLP aims at selecting facility locations to meet customer demand, at minimal cost. See (8), (9), and (10) for literature reviews. FLP has also been invaluable in solving optimization problems in healthcare and humanitarian logistics, including vaccine distribution. For example, (11) discussed the location of COVID-19 mass vaccination facilities in the United States. We refer the reader to (12) for an extensive literature discussion on vaccine supply chains and to (13) for surveys on healthcare facility location.

Another important stream of literature is on emergency facility location (EFL) problems, i.e., optimization of facility locations in the face of natural disasters (14). Furthermore, EFL relates closely to another stream of literature on reliable facility location (RFL) problems. RFL deals with optimal decisions considering failure of a facility due to disruption (e.g. natural disasters, blackouts, fire, etc). A comprehensive review is presented in (15). RFL models require knowledge of the occurrence of disruptions which is a difficult undertaking in practice. Thus, the majority of works on RFL models largely assume that probability distribution of disruptions is known. Nonetheless, the emergence of robust optimization (RO) dealt with the limited information on the uncertainty of disruptions. For example, (16) introduce a model that allows disruptions to be correlated with an uncertain joint distribution and minimize the expected cost under the worst-case distribution with given marginal disruption probabilities. (17) also tackled facility location decisions with random disruptions and investigated the impact of misestimating the disruption probability using a stylized continuous location model. Additionally, recent work of (18) adopts a two-stage robust optimization method, where facility location decisions are made *here-and-now* (first stage) and *wait-and-see* (second stage) are entail reassigning customers after revealing information on facility availability.

However, most papers in the literature only consider static and two-stage cases. Furthermore, in contrast to the abundant studies on deterministic multi-period setting, there are only a limited number of works that formulate facility location problems as multi-stage stochastic programs, and all of them only consider stochastic demand but not facility disruption. For example, (19) formulate a risk-averse multistage stochastic integer program (MSIP) with random *demand* and utilize expected conditional risk measures to provide tight lower bounds. Other multi-stage formulations with stochastic demand are in (20), (21), and (22). To the best of our knowledge, there exist no prior work specifically on multistage stochastic facility location with facility disruption that accounts for the risk of facility failure.

3 Model Formulation

We formulate the facility location model with risk of disruption due to natural disaster as a multi-stage, discrete-time, stochastic-dynamic optimization problem. We decompose the planning horizon T into a sequence of time stages $t \in \{1, ..., T\}$. We define state variable $u_{it} \in \{0, 1\}$ such that $u_{it} = 1$ if facility $i \in \{1, \ldots, F\}$ is open at stage t, and $u_{it} = 0$ otherwise. The transition of the state variable from its value at the beginning of the stage to its value the end of the stage is governed by the series of constraints and relationships we detail below. In addition, we observe a realization of the random variable $\xi_{it} \in \{0, 1\}$ such that $\xi_{it} = 1$ if the vaccine facility i at stage t fails due to natural disaster and $\xi_{it} = 0$ otherwise.

After the observation of the state variable u_{it} and the random variable ξ_{it} at the beginning of each stage, the values of the decision (or control) variables need to be determined. The control variables are $x_{it} \in \{0, 1\}$ (binary variable with value 1 if vaccine facility i should be opened at stage t and 0, otherwise) and z_{ijt} (allocation of vaccinating population of district $j \in \{1, \ldots, B\}$ to vaccine facility i at stage t if vaccine facility i is open). We formulate the corresponding multistage stochastic facility location problem (MSFLP) model as follows:

$$\min \quad \sum_{t}^{T} \sum_{i}^{F} (f_{it} x_{it} + \sum_{j}^{B} d_{ij} z_{ijt}) \tag{1}$$

$$\text{s.t.} \quad u_{it} = u_{it-1}(1 - \xi_{it}) + x_{it}, \quad \forall i = 1, \ldots, F, t = 1, \ldots, T \tag{2}$$

$$\sum_{j}^{B} z_{ijt} \leq C_i u_{it}, \quad \forall i = 1, \ldots, F, t = 1, \ldots, T \tag{3}$$

$$\sum_{i}^{F} z_{ijt} = P_j, \quad \forall j = 1, \ldots, B, t = 1, \ldots, T \tag{4}$$

$$u_{it}, x_{it} \in \{0, 1\}, \quad \forall i = 1, \ldots, F, t = 1, \ldots, T \tag{5}$$

$$z_{ijt} \in \mathbb{Z}, \quad \forall i = 1, \ldots, F, j = 1, \ldots, B, t = 1, \ldots, T \tag{6}$$

The objective function (1) is a straightforward minimization of the total cost function which include the fixed cost f_{it} of opening vaccine facility i at stage t and the travel cost d_{ij} from district j to vaccine facility i. We need to keep track of the state of vaccine facility i as affected by random variable ξ_{it} across each stage through constraints (2). If at stage $t-1$ the vaccine facility i is closed [open] and we decide to open [close] it at stage t, then $x_{it}=1$ [$x_{it}=0$], the state at stage t will now become $u_{it}=1$ [$u_{it}=0$]. Constraints (3) state that the total capacity C_i will be assigned to vaccine facility i only if it is available and should at least be equal to the total number of vaccinating population z_{ijt} in vaccine facility i from district j at stage t. Constraints (4) state that the total population P_j of district j must be distributed to vaccine facilities i. Constraints (5) and (6) show the decision variable domains.

For the rest of the paper, we use bold symbols to denote vectors and matrices. Let us rewrite MSFLP in vector form,

$$\min_{\mathbf{u}_t,\mathbf{y}_t} \quad \sum_{t\in T}(v_t^\top \mathbf{u}_t + w_t^\top \mathbf{y}_t) \tag{7}$$

$$\text{s.t.} \quad \mathbf{u}_t = \mathbf{u}_{t-1}(1-\boldsymbol{\xi}_t)+\mathbf{y}_t, \quad \forall t\in 1,\dots,T \tag{8}$$

$$\mathbf{A}\mathbf{y}_t \le \mathbf{u}_t, \quad \forall t\in 1,\dots,T \tag{9}$$

$$\mathbf{B}\mathbf{y}_t = \mathbf{b}, \quad \forall t\in 1,\dots,T \tag{10}$$

$$\mathbf{u}_t\in\{0,1\}^F, \mathbf{y}_t\in\mathbb{Z}_+^{F\times B}, \quad \forall t\in 1,\dots,T \tag{11}$$

where $\mathbf{u}_t$, $\mathbf{y}_t=(\mathbf{x}_t,\mathbf{z}_t)$ are the state and control variables respectively, $v_t^\top$,$w_t^\top$ represent the cost function, and matrices $\mathbf{A}\in\mathbb{R}^{F\times B}$, $\mathbf{B}\in\mathbb{R}^{F\times B}$, and $\mathbf{b}\in\mathbb{R}^B$ correspond to the coefficients of constraints (3) and (4).

Recall that random variable $\boldsymbol{\xi}_t\in\{0,1\}$ takes the value 1 if the facility fails at stage t or 0 if it does not. To facilitate the formulation of the dynamic program, we introduce a scenario-path based notation. We take one possible path of realizations from the beginning to the end of the planning horizon, and denote it as a scenario path $\boldsymbol{\omega}$, i.e.,

$$\boldsymbol{\xi}^{\boldsymbol{\omega}} = (\boldsymbol{\xi}_1^{\boldsymbol{\omega}},\dots,\boldsymbol{\xi}_T^{\boldsymbol{\omega}}). \tag{12}$$

Accordingly, we use $\mathbf{u}_t^{\boldsymbol{\omega}}$ and $\mathbf{y}_t^{\boldsymbol{\omega}}$ to denote the state and decision vectors at stage t under scenario path $\boldsymbol{\omega}$. Since the number of realizations is finite, let $\boldsymbol{\Omega}$ be the support set of $\boldsymbol{\omega}$ with each realization $\boldsymbol{\omega}\in\boldsymbol{\Omega}$ having probability $p(\boldsymbol{\omega})$ such that

$$\sum_{\boldsymbol{\omega}\in\boldsymbol{\Omega}} p(\boldsymbol{\omega}) = 1. \tag{13}$$

See for example (23) who make similar assumptions.

3.1 Multistage Stochastic Integer Program (MSIP) Formulation

We formulate the multistage stochastic facility location problem (MSFLP) as a multistage stochastic integer program (MSIP), with the form

$$\begin{aligned}\min_{\mathbf{u_1},\mathbf{y_1}\in\mathcal{X}_1}\Big\{v_1^\top\mathbf{u_1}+w_1^\top\mathbf{y_1}+\mathbb{E}_{\boldsymbol{\xi}_{[2,T]}|\xi_{[1,1]}}\Big[\min_{\mathbf{u_2},\mathbf{y_2}\in\mathcal{X}_2(\mathbf{u}_1,\boldsymbol{\xi}_2)}\Big\{v_2^\top\mathbf{u_2}+w_2^\top\mathbf{y_2}+\dots\\ +\mathbb{E}_{\boldsymbol{\xi}_{[T,T]}|\xi_{[1,T-1]}}\Big[\min_{\mathbf{u_T},\mathbf{y_T}\in\mathcal{X}_T(\mathbf{u_{1:T-1}},\boldsymbol{\xi}_T)}\Big\{v_T^\top\mathbf{u_T}+w_T^\top\mathbf{y_T}\Big\}\Big]\dots\Big\}\Big]\Big\},\end{aligned} \tag{14}$$

where, for notation simplicity, $\mathbf{u_t}$ and $\mathbf{y_t}$ are the state variables and control (i.e. local or stage) variables respectively, $v_t^\top$ and $w_t^\top$ are the corresponding cost at time $t\in 1,\dots,T$, and $\mathcal{X}_t$ is the feasible set.

In this setting, the stochastic data process is revealed as time goes on. This stochastic data process $(\boldsymbol{\xi}_1,\dots,\boldsymbol{\xi}_T)$ is modeled where $\boldsymbol{\xi}_1$ is deterministic and $\boldsymbol{\xi}_2,\dots,\boldsymbol{\xi}_T$ will be revealed gradually in time. Thus, the history of the stochastic data process until stage t is given by

$$\boldsymbol{\xi}_{[t]} = (\boldsymbol{\xi}_1\dots\boldsymbol{\xi}_T). \tag{15}$$

This formulation allows the decision at each stage to adapt to the realized uncertainty up to that stage. Thus, by solving MSFLP, we get a policy on which facilities should be opened for each possible scenario.

3.2 Dynamic Programming Reformulation

We can further formulate (14) as a dynamic programming (DP) recursion. The optimal value function at stage t, $V_t(\mathbf{u}_{t-1},\boldsymbol{\xi}_t)$, is the optimal expected objective value given state $(\mathbf{u}_{t-1},\boldsymbol{\xi}_t)$, and assuming that optimal action will be

taken at each stage t.

$$V_t(\mathbf{u_{t-1}}, \boldsymbol{\xi}_t) = \min_{\mathbf{u}_t, \mathbf{y}_t} \{v_t^\top \mathbf{u}_t + w_t^\top \mathbf{y}_t + \mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t) : B_t \mathbf{u}_{t-1} + A_t \mathbf{u}_t + C_t \mathbf{y}_t = b_t\}, \tag{16}$$

for $t = 1, \ldots, T$ where $\mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t)$ is the expected value cost-to-go function,

$$\mathcal{V}_{t+1}(\mathbf{u}_t) := \mathbb{E}[V_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_{t+1}) | \boldsymbol{\xi}_t]. \tag{17}$$

If we assume $\boldsymbol{\xi}_t$ to be Markovian, i.e. it follows the Markov property or the *memoryless* property, the expectation depends only on $\boldsymbol{\xi}_t$ rather than the whole history of the data process, with $\mathcal{V}_T \equiv 0$. Finally, let us define the optimal policy as

$$\pi^*(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t) = \arg\min_{\mathbf{u}_t, \mathbf{y}_t} \{v_t^\top \mathbf{u}_t + w_t^\top \mathbf{y}_t + \mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t) : B_t \mathbf{u}_{t-1} + A_t \mathbf{u}_t + C_t \mathbf{y}_t = b_t\} \tag{18}$$

for $t = 1, \ldots, T$ in set $\mathbf{\Pi}$ as the *policy* which specifies the decision to make for all possible states regardless of which state at stage t.

3.3 MSIP with Binary State Variables

Each subproblem of MSFLP given realized $\boldsymbol{\xi_t}$ at stage t is a deterministic mixed-integer program. With this, we make the following assumption:

Assumption 1 *The objective function* $\mathbf{z} = \sum_{t \in T}(v_t^\top \mathbf{u}_t + w_t^\top \mathbf{y}_t)$ *in each realization* $\boldsymbol{\xi}_t$ *is an affine function in* u_t *and* y_t*, and the constraint set* $\mathcal{X}_t$ *is a nonempty compact mixed integer polyhedral set.*

Mathematically, mixed-integer programs (MIPs) are non-convex. Therefore, given integer local variables, the value functions $V_t(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t)$ and expected cost-to-go functions $\boldsymbol{\mathcal{V}}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t)$ at stage t are both non-convex with respect to the state variables. Although it is impossible to formulate convex polyhedral representations of the non-convex value functions that are tight at the evaluated state variable values, real-valued function of binary variables can be represented exactly by a convex polyhedral function. (24) also made similar assumption.

Remark 1 *The value functions,* $V_t(\mathbf{u_{t-1}}, \boldsymbol{\xi}_t)$*, are convex in state variable* $\mathbf{u}_t \in \{0, 1\}$ *if the optimal values of*

$$\min_{\mathbf{u}_t, \mathbf{y}_t} \{v_t^\top \mathbf{u}_t + w_t^\top \mathbf{y}_t + \mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi_t}) : B_t \mathbf{u}_{t-1} + A_t \mathbf{u}_t + C_t \mathbf{y}_t = b_t\} \tag{19}$$

are convex in $\mathbf{u}_t$*, which is inherent with the fact that if* $\mathbf{u}_t \in \{0, 1\}$ *and* $\mathbf{u}_t$ *is the convex combination of a set of binary vectors, then* $\mathbf{u}_t$ *is a binary vector, for* $t = 1, \ldots, T$.

3.4 Value Function Approximation

Recall that the value functions $V_t(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t)$ represent the minimal cost assuming optimal action will be taken at stage t. For tractable and small problems with a finite state space, backwards dynamic programming can be used to compute the optimal value function. However, as the number of states grows exponentially with the dimensionality of the state space, many real-world problems become computationally intractable. For this reason, a large body of research is dedicated to methods of approximate dynamic programming (ADP) which aims at using functions of lower complexity to approximate the true value function.

There exist numerous strategies for value function approximation, such as aggregation with lookup tables, basis functions, neural networks, etc. See (25) for an extensive treatment of the subject.

More sophisticated techniques, often aim at exploiting structure of the underlying optimization problem. For example, *stochastic dual dynamic integer programming* (SDDiP) - the *state-of-the-art* algorithm to solve MSIP problems advanced in (24) - approximates the convex piece-wise linear expected cost-to-go function by a set of (cutting) hyperplanes obtained from subgradients of the subproblems in stages $t = 1, \ldots, T$ of a set of sampled outcomes. However, the algorithm requires the random process to be either stagewise-independent or to follow a discrete Markov chain (26).

In this article, we pursue a much simpler strategy by approximating the value function by a linear function. In what follows, we will demonstrate how to find parameters of a linear value function approximation that not only finds a good solution to the MSFLP but moreover provides practitioners with actionable insights of the risks and benefits of opening a facility under risk of disruption.

3.5 Linear Value Function Approximation

We propose to approximate the convex expected value function, $\mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t)$ by a linear function,

$$\mathcal{V}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t) \approx \bar{\mathcal{V}}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t; \boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t) = \boldsymbol{\lambda}_t^{u\top}\mathbf{u}_t + \boldsymbol{\lambda}_t^{\xi\top}\boldsymbol{\xi}_t + \beta_t,\ t = 1, \dots, T-1, \tag{20}$$

with $\boldsymbol{\lambda}_t^u$, $\boldsymbol{\lambda}_t^\xi$ as slope vectors and β_t as intercept. The vector $\boldsymbol{\lambda}_t^u$ defines the marginal future cost of opening a facility in period t. We will later refer to this as the *shadow price* of opening a facility, as it accounts for the possible future price that needs to be paid later for re-opening facilities in case of failure.

In most cases, calculating the value function is just means to an end. What we are truly interested in is finding an approximation of the optimal policy π^*. The approximate policy that corresponds to the linear value function approximation (20) is given by

$$\bar{\pi}(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega; \boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t) = \underset{u,y\geq 0}{\arg\min}\{v_t^\top\mathbf{u}_t + w_t^\top\mathbf{y}_t + \bar{\mathcal{V}}_{t+1}(\mathbf{u}_t, \boldsymbol{\xi}_t^\omega) : B_t\mathbf{u}_{t-1} + A_t\mathbf{u}_t + C_t\mathbf{y}_t = b_t\}, \tag{21}$$

for $t = 1, \dots, T$, where $\omega \in \Omega$ is a scenario path of a stochastic process with its support set Ω.

With this policy, we can now define an upper bound on the optimal objective value of the MSFLP,

$$\bar{\mathbf{z}}(\bar{\pi}) = \mathbb{E}_{\boldsymbol{\omega}\in\boldsymbol{\Omega}_{[1,T]}}\left[\bar{z}(\bar{\pi})\right] \tag{22}$$

where (suppressing the dependence of $\bar{\pi}$ on $\boldsymbol{\lambda}_t^u$, $\boldsymbol{\lambda}_t^\xi$, and β_t)

$$\bar{z}(\bar{\pi}) = \sum_t^T \Big(v_t^\top\mathbf{u}_t(\bar{\pi}(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega)) + w_t^\top\mathbf{y}_t(\bar{\pi}(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega))\Big). \tag{23}$$

If all we are interested in is finding parameters $(\boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t)$ that provide the lowest upper bound, we can further simplify the policy by removing the terms $\boldsymbol{\lambda}_t^{\xi\top}\xi_t^\omega$ and β_t from the right-hand side of (20). This is supported by the following proposition.

Proposition 1 *For a linear value function approximation, $\bar{\mathcal{V}}_t(\cdot; \boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t)$ that is affine in state variables $\boldsymbol{u}_{t-1}$ and random variables $\boldsymbol{\xi}_t$, it holds that the optimal upper bound is independent of parameters $\boldsymbol{\lambda}_t^\xi$ and β_t.*

Proof. By Equations (22) and (23), we obtain an upper bound $\bar{\mathbf{z}}$ by choosing parameters of approximate policy $\bar{\pi}$, i.e., slope vectors $\boldsymbol{\lambda}_t^u$, $\boldsymbol{\lambda}_t^\xi$, and intercept β_t. Approximate policy $\bar{\pi}$ is defined by state variables $\mathbf{u_{t-1}}$ and local variables $\mathbf{y_t}$ as shown in Equation (21). Now, let us denote

$$\bar{\pi}'(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega; \boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t) = \underset{u,y}{\arg\min}\{v_t^\top\mathbf{u}_t + w_t^\top\mathbf{y}_t + \underbrace{\boldsymbol{\lambda}_t^{u\top}\mathbf{u}_t + \boldsymbol{\lambda}_t^{\xi\top}\boldsymbol{\xi}_t^\omega + \beta_t}_{\text{objective function offset}} : B_t\mathbf{u}_{t-1} + A_t\mathbf{u}_t + C_t\mathbf{y}_t = b_t\}, \tag{24}$$

for $t = 1, \dots, T$.

Random variable $\boldsymbol{\xi}_t^\omega$ enters the approximate expected cost-to-go function $\bar{\mathcal{V}}_{t+1}(\cdot)$ as an offset to the objective function. Hence, it does not effect the optimal choice of $\mathbf{u_{t-1}}$ and $\mathbf{y_t}$ in (24). Therefore, (24) can be simplified to

$$\bar{\pi}'(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega; \boldsymbol{\lambda}_t^u) = \underset{u,y}{\arg\min}\{v_t^\top\mathbf{u}_t + w_t^\top\mathbf{y}_t + \boldsymbol{\lambda}_t^{u\top}\mathbf{u}_t : B_t\mathbf{u}_{t-1} + A_t\mathbf{u}_t + C_t\mathbf{y}_t = b_t\}, \tag{25}$$

for $t = 1, \dots, T$, where the right-hand side is independent of $\boldsymbol{\lambda}_t^{\xi\top}$ and β_t. Thus, for any realization of random variable $\boldsymbol{\xi}_t^\omega$, policy $\bar{\pi}'$ yields the same values. Accordingly, it holds that

$$\bar{\pi}'(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega; \boldsymbol{\lambda}_t^u) = \bar{\pi}(\mathbf{u}_{t-1}, \boldsymbol{\xi}_t^\omega; \boldsymbol{\lambda}_t^u, \boldsymbol{\lambda}_t^\xi, \beta_t),\ \forall t = 1, \dots, T. \tag{26}$$

Hence, by equivalence of approximate policies $\bar{\pi}$ and $\bar{\pi}'$, the upper bound $\bar{\mathbf{z}}(\bar{\pi})$ remains unchanged, and it follows that upper bound $\bar{\mathbf{z}}(\bar{\pi}) \equiv \bar{\mathbf{z}}(\boldsymbol{\lambda}_t^u)$, for $t = 1, \dots, T$, is independent of parameters $\boldsymbol{\lambda}_t^\xi$ and β_t, which concludes the proof. □

Proposition 1 tells us that parameters $\boldsymbol{\lambda}_t^\xi$ and β_t do not influence the optimal upper bound, which implies that we can set them to zero, effectively removing them from the objective function. This enables us to simplify the functional form of the approximate expected value function to

$$\bar{\mathcal{V}}_{t+1}(\boldsymbol{u}_t) \equiv \boldsymbol{\lambda}_t^{u\top}\mathbf{u}_t, \tag{27}$$

which is independent of the state of the stochastic process $\boldsymbol{\xi}_t$, for $t = 1, \dots, T$.

Remark 2 *The proposed value function approximation grows linearly in the number of time-coupling control variables and the number of stages, and its parameters have a natural economic interpretation as approximations of the shadow prices (dual values) of the time-coupling balance equations that connect subproblems of successive stages.*

4 Solution Method

In this section, we describe a policy approximation algorithm to solve MSFLP. As the algorithm effectively searches for the shadow prices defined by the linear value function approximation, we refer to this method as *shadow price approximation* (SPA).

SPA trains the approximate cost-to-go function, $\bar{\mathcal{V}}_{t+1}(\mathbf{u_{t-1}}, \boldsymbol{\xi}_t),\ \forall t = 1, \ldots, T$, of the stochastic-dynamic program by directly interacting with a simulation model of the resulting policy. This is unlike many other approximate dynamic programming techniques that approximate the cost-to-go using some sort of backwards dynamic programming. The main idea of SPA is to view the parameters of a linear value function approximation as tunable parameters of a policy. Tuning parameters to improve decision-making in the face of uncertainty is widely used in practice to make deterministic optimization problems more robust against a variety of risks, for example, safety stock and buffers in production planning (27), upper and lower limits for water reservoirs to capture rainfall uncertainty (28), and slack time to buffer uncertainty in job completion (29).

An alternative to tuning the parameters of a linear value function approximation is to introduce tunable parameters into the constraints of deterministic optimization problems. (7) propose *parametric cost function approximation* (PCFA), a method that introduces such parameters to the right-hand side of the constraints of a deterministic rolling-horizon model, in which uncertainty parameters are replaced by their expected values. Naturally, such problems suffer from overly optimistic decisions and introduction of tighter constraints provides a hedge against uncertainty. In contrast to PCFA, SPA introduces tunable parameter $\boldsymbol{\lambda}_t^u$ as objective coefficient of time-coupling variable, $\boldsymbol{u}_t$.

4.1 Dynamic Programming Formulation

We start our exposition of SPA by introducing a reformulation of the MSFLP to a dynamic program using linear value function approximation (27). Following the reformulation defined in (16), the value function of the dynamic program with a linear value function approximation is given by

$$\begin{aligned} \bar{V}_t(\mathbf{u_{t-1}}, \boldsymbol{\xi}_t; \boldsymbol{\lambda}_t^u) = \min_{u_{it}, x_{it}, z_{ijt}} \quad & \underbrace{\sum_i^F (f_i x_{it} + \sum_j^B z_{ijt} d_{ij})}_{\text{Total Operations Cost}} + \underbrace{M \sum_i^F \lambda_{it}^u u_i}_{\text{Shadow Price}} \\ \text{s.t.} \quad & u_{it} = u_{i,t-1}(1 - \xi_{it}) + x_{it}, \quad \forall\, i = 1, \ldots, F, t = 1, \ldots, T, \\ & \text{Eqs.} \quad (3) - (6), \end{aligned} \tag{28}$$

for $t = 1, \ldots, T$, where $\lambda_{iT}^u \equiv 0, \forall i = 1, \ldots, F$, and M is a constant parameter (that will be discussed below).

The resulting parametric policy that uses this reformulation is given by

$$\begin{aligned} \bar{\pi}(\mathbf{u_{t-1}}, \boldsymbol{\xi}_t; \boldsymbol{\lambda}_t^u) = \arg\min_{u_{it}, x_{it}, z_{ijt}} \quad & \sum_i^F (f_{it} x_{it} + \sum_j^B z_{ijt} d_{ij}) + M \sum_i^F \lambda_{it}^u u_{it} \\ \text{s.t.} \quad & u_{it} = u_{i,t-1}(1 - \xi_{it}) + x_{it}, \quad \forall\, i = 1, \ldots, F, t = 1, \ldots, T, \\ & \text{Eqs.} \quad (3) - (6). \end{aligned} \tag{29}$$

Given a set of slope vectors, $\lambda^u = (\lambda_{i,1}^u, \ldots, \lambda_{i,T-1}^u), \forall i = 1, \ldots, F$, and a process of facility site failures, $\xi^\omega = (\xi_{i,1}^\omega, \ldots, \xi_{i,T}^\omega), \forall i = 1, \ldots, F$, where $\omega \in \Omega$ is a scenario path of a stochastic process with its support set Ω, we can now use this policy to obtain an estimate of the upper bound $\bar{\mathbf{z}}(\lambda^u)$. By sampling N scenarios of possible site failures, $\hat{\xi}^n = (\hat{\xi}_{i,1}^n, \ldots, \hat{\xi}_{i,T}^n), \forall i = 1, \ldots, F$, and assuming time-coupling variable $u_{i,0}^n \equiv 0$ where $i = 1, \ldots, F$ and $n = 1, \ldots, N$, we obtain

$$\bar{\mathbf{z}}(\lambda^u) = N^{-1} \sum_{n=1}^N \hat{\bar{z}}^n(\lambda^u), \tag{30}$$

where

$$\hat{\bar{z}}^n(\lambda^u) = \sum_{t \in T} \left[v_t^\top u_t(\bar{\pi}(u_{t-1}^n, \hat{\xi}_t^n; \lambda_t^u)) + w_t^\top y_t(\bar{\pi}(u_{t-1}^n, \hat{\xi}_t^n; \lambda_t^u)) \right] \tag{31}$$

is a realization of the total cost for given slope vectors λ^u.

4.2 Shadow Price Approximation

As slope vector λ^u can be viewed as the shadow price of opening facilities, it accounts for the possible future price that needs to be paid later for re-opening facilities after failure due to natural disaster.

The objective of SPA is to choose a slope vector that minimizes the upper bound,

$$\lambda^{u*} = \arg\min_{\lambda^u} \bar{\mathbf{z}}(\lambda^u). \tag{32}$$

A problem of practical importance is that the mapping $\lambda^u \mapsto \hat{\bar{z}}$ is not convex and noisy, which means that we must effectively solve a non-convex, stochastic optimization problem that is known to be computationally intractable in its general form. In the following, we therefore cast SPA as a policy search strategy that can be supported by any method that is suitable for unconstrained (derivative-free) global optimization. In Section 5.2, we briefly present a few global optimization methods that were tested to aide SPA with solving the MSFLP.

The shadow price approximation algorithm works as follows: first, we initialize shadow prices before the first iteration i.e. $\lambda_t^{u,0} = 0,\ \forall t = 1, \ldots, T$. Then, at each iteration n, we randomly select a scenario path $(\hat{\xi}_1^n, \ldots, \hat{\xi}_T^n)$, and solve Problem (29) forward in time. Then, the algorithm passes parameters $\lambda^{u,n-1}$ and realized upper bound $\hat{\bar{z}}^n$ to some update function $\boldsymbol{U}^n(\cdot,\cdot)$, supported by a global optimization method, that returns a new set of trial shadow prices $\lambda^{u,n}$. At the end of iteration N, the algorithm returns the trial $\lambda^{u,N}$ that minimize $\bar{z}$. This procedure is summarized in Algorithm 1.

Algorithm 1: Shadow Price Approximation

Initialize $\lambda_t^{u,0} = 0,\ \forall t = 1, \ldots, T$
Set initial state $u_0 \equiv 0$
Set M and $n = 1$
while $n < N$ **do**
  Draw scenario $(\hat{\xi}_1^n, \ldots, \hat{\xi}_T^n)$ from the stochastic process
  for $t \in \{1, \ldots, T\}$ **do**
    $(u_t^n, x_t^n, z_t^n) \leftarrow \arg\min \left\{ \sum_i^F (f_{it}x_{it} + \sum_j^B z_{ijt}d_{ijt}) + M \sum_i^F \lambda_{it}^{u,n-1} u_{it} : \text{Eqs (2)-(6)} \right\}$
  end for
  $\hat{\bar{z}}^n \leftarrow \sum_{t \in T} \left[\sum_i^F (f_{it}x_{it}^n + \sum_j^B z_{ijt}^n d_{ijt}) \right]$
  $\lambda^{u,n} \leftarrow \boldsymbol{U}^n(\hat{\bar{z}}^n, \lambda^{u,n-1})$
  $n \leftarrow n + 1$
end while
Return $\lambda^{u,N}$

Note that parameter M is optional, but as we will see in Appendix B (6), changing the choice of M can accelerate the convergence of the algorithm. A good *rule-of-thumb* in choosing initial value of parameter M is also presented in Appendix B (6).

An advantage of SPA over conventional stochastic and dynamic programming methods is that generating scenarios can be as simple as drawing a bootstrap sample from historical data. Bootstrapping alleviates the need to develop a stochastic model and effectively turns SPA into a data-driven approximate dynamic programming algorithm. When scenarios are generated by drawing from the empirical distribution of historical data, cross-validation and out-of-sample testing follow naturally. As such techniques are well-known to data scientists and practitioners of machine learning, SPA may be particularly attractive for the large pool of practitioners that are already familiar with such paradigms.

5 Case Study

This section presents the construction of the rainfall-to-flood susceptibility mapping to incorporate the variability of rainfall as a random disruption parameter of the overall MSFLP model. This also briefly introduces the geographical and climactic characteristics of Cagayan de Oro City, a large city in the Philippines, which serves as the study case for this paper. Section 5.2 briefly introduces the three different global optimization methods chosen to aid the SPA algorithm. Finally, we discuss the numerical results in Section 5.3 and cross-validation in Section 5.4. We programmed our solution approaches as single threaded applications in Python 3.7 with GUROBI 9.1.1 as the solver of the mixed integer linear program. All computations were carried out on a Linux 4.15 server with an Intel Xeon Gold 2.10GHz processor and 256 GB RAM.

5.1 Rainfall-to-Flood Susceptibility Mapping

In most developing countries, like the Philippines, one of the most detrimental natural disasters is heavy rainfall resulting in flooding. During the height of the COVID-19 pandemic, in addition to immediate hazards when heavy rainfall and flooding occurs, the probability of vaccination facilities becoming unavailable and inaccessible increases. It is for this reason that we chose flooding, which cause the facility failure, as the natural disaster risk to hedge from for the MSFLP.

In this paper, the risk of failure of a vaccine facility i at time t is highly dependent upon its state u_{it} as triggered by the random disruption variable ξ_{it}. Thus, it is beneficial to know flooding event information of vaccine facility i. We constructed rainfall-to-flood susceptibility (RTFS) mapping of the different vaccination facilities F, where we combine real rainfall data with geo-spatial flood risk map data to infer whether a site is accessible or otherwise. The RTFS mapping provides the information for when the vaccination facility i becomes at risk of failure i.e. closed and inaccessible, due to flooding at time t. Figure 1 shows the overall flowchart for the construction of the RTFS mapping.

5.1.1 Case Study: Cagayan de Oro City, Philippines

We performed numerical experiments on the data from a large city in the Philippines, Cagayan de Oro City, which experience heavy rainfall multiple times a year. Cagayan de Oro City is located along the northern central coast of Mindanao Island. As the capital city of the Province of Misamis Oriental, it serves as a regional center and hub for Northern Mindanao region. The city has an estimated population of more than 728,402 as of 2020 census and is the 10th most populated city in the Philippines, divided into eighty (80) districts (locally, *barangays*).

Dividing the city traverses the Cagayan de Oro River, one of the major rivers in Mindanao. This catchment discharges huge amount of water during a heavy downpour, even more so since the city is classified with a tropical monsoon climate. The city's flat slope and swallowing of the channel as it approaches the delta poses flood risks to the residents of the city. In the last twenty years, the city suffered heavy losses due to flooding (30). Thus, appropriate mitigation measure is imperative especially for long-term projects like the COVID-19 vaccination campaigns.

As of July 2021, the local government unit of Cagayan de Oro opens three (3) vaccination facility sites. These three vaccine facilities currently cater to the first few delegations of the vaccination. As the COVID-19 vaccination campaign ramps up, the city's local government plans to open more vaccine facilities to expedite the herd immunity goal of inoculating at least 70% of the total population. The roster of candidate vaccine facilities F were pre-selected to adhere to the standards of Philippines' Department of Health (DOH). The specifics of the RTFS mapping construction take into consideration this Cagayan de Oro City as an example.

5.1.2 RTFS map construction

The Philippine Atmospheric, Geophysical, and Astronomical Services Administration (PAGASA), the national meteorological and hydrological services agency of the Philippines, provided us with twenty-one (21) years of daily rainfall (in mm) data from year 2000 until 2020. The twenty-one years of data induce the different scenario cases S for the rainfall-to-flood susceptibility (RTFS) mapping of the vaccination facilities F.

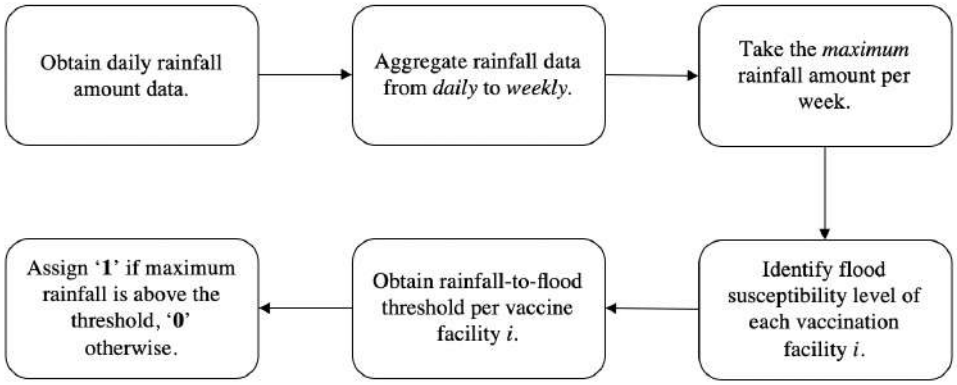

Figure 1: Rainfall-to-Flood Susceptibility (RTFS) Mapping Flowchart

One general assumption of this paper is the aggregation of rainfall on a *weekly* basis. Given the twenty-one years of daily rainfall data from PAGASA, we aggregated one week of data following the calendar week, following the

fifty-two-week calendar year. Thus, the decision epoch t considered for this paper is *weekly* over an entire year of planning horizon $T = 52$ decision weeks. It then follows that the decision to open and close facility i dynamically changes on a weekly basis. This assumption is carefully escalated to the local government unit and policy-makers noting that any vaccine facility may be instantaneously opened or closed in a chosen week.

In order to aggregate the daily precipitations into weekly periods, the *maximum* rainfall amount of week t is taken to represent the aggregated rainfall amount of week t. This implies that if the *maximum* amount of rain on week t surpasses the rainfall-to-flood threshold, the vaccine facility i will be closed and inaccessible for the entire week t. The rainfall-to-flood threshold is identified through a series of circumspect investigation using disaster risk geo-spatial information system (GIS) and the official guidelines of PAGASA on rainfall advisory and classifications.

GeoRisk Philippines led by the Philippine Institute of Volcanology and Seismology (PHIVOLCS) created *Hazard-HunterPH* - an online GIS tool which generates indicative hazard assessment reports on specified locations in the Philippines (31). *HazardHunterPH* maps the entire Philippines and shows different layers of geographical hazards e.g. seismic (earthquakes), volcanic, and hydro-meteorological (flood, landslides, and storm surges).

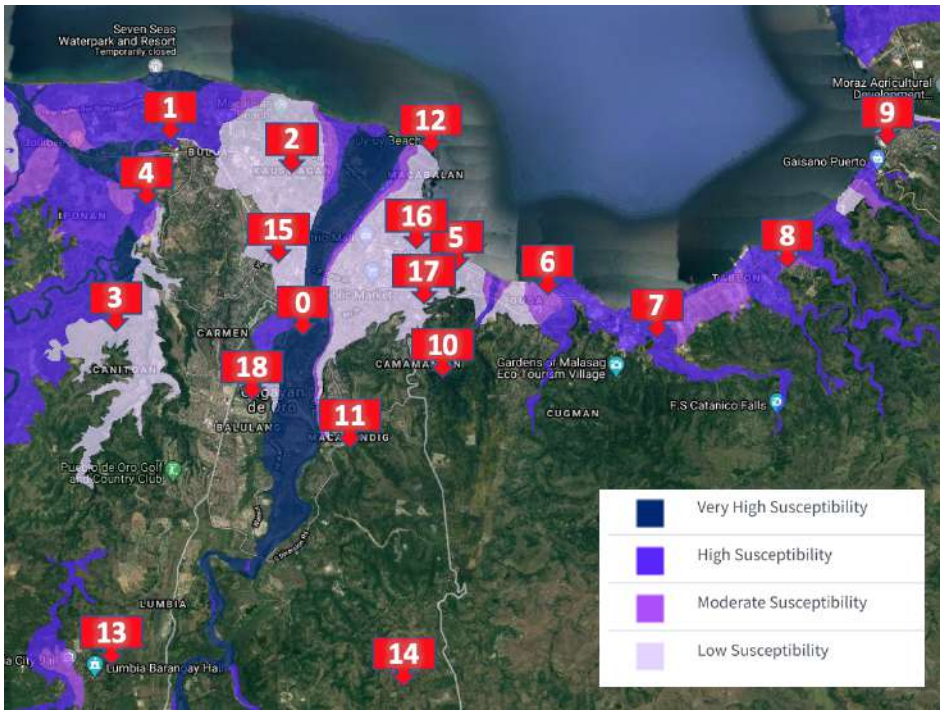

Figure 2: Candidate vaccine facility sites of Cagayan de Oro City, Philippines with overlaid flood hazard map

Figure 2 shows a map of Cagayan de Oro City with the overlaid flood hazard map. The different hues of violet represent the flood susceptibility of an area ranging from *Very High Susceptibility* in darker hue to *Low Susceptibility* in lighter hue. This signifies that if a vaccine facility i is located in a *Very High Susceptibility* zone, vaccine facility i is easily flooded more than vaccine facilities located in *Low Susceptibility* zones.

We also mapped the set of candidate vaccine facility sites F, as provided by the local government unit, to distinguish the flood susceptibility of each site location shown in Figure 2. For a more thorough investigation, a further examination is conducted with the Engineering Resource Center of Xavier University - Ateneo de Cagayan, a leading university of Cagayan de Oro City, on the flood susceptibility level of the routes leading to the different facilities. The assumption on this further examination insinuates that if *routes* leading to the vaccine facility i has a higher flood susceptibility level, the chosen susceptibility level of facility i will be lifted to the next (higher) susceptibility level. This alludes logically that even though vaccine facility i is not flooded, it will be inaccessible due to the routes leading to it being flooded.

After assigning the prominent flood susceptibility level of each vaccine facility site i, the rainfall-to-flood threshold is identified for each site. Nonetheless, the rainfall-to-flood susceptibility (RTFS) mapping leaves room for modeling error. Thus, we have created three conservativity scenario cases that vary the translation of precipitation amounts to whether a facility is flooded or not. These scenario cases are based on the prominent flood susceptibility level and the official rainfall advisories, classification, and measurement of PAGASA (32). The cases $\mathbf{S} = \{Low, Medium, High\}$

are instituted from the different ranges of rainfall measurement. Each case $\mathbf{s}$ take the level of conservativity of the ranges of rain measurement. $High$ conservativity case takes the minimum value, the $Medium$ conservativity case takes the median value, and the Low conservativity case takes the maximum value of the rainfall measurement range. The inclusion of the different cases $\mathbf{S}$ parlays more robustness in the decision-making process of the study. This would also allow a more encompassing solution to understand in depth the different plausibilities played by natural disasters.

The different RTFS mapping cases were transformed into zeros and ones to show the value of the random variable ξ_{it} of each vaccine facility i at time t. Recall that if a vaccine facility site i flooded at time t i.e. $\xi_{it} = 1$, then the facility is disrupted and unavailable, otherwise, $\xi_{it} = 0$. Each of these three different RTFS mapping conservativity cases contain the twenty-one years scenarios with fifty-two weeks, one for each decision epoch t of the entire planning horizon T.

Finally, it should be noted that the steps carried out on the construction of the RTFS mapping for Cagayan de Oro City can be utilized to create RTFS mapping for other cities in other developing countries. Thus, distinct differences on geophysical and geographical characteristics of the city should be assessed thoroughly to provide more conclusive inference on the decision planning. In Section 5.4, another large city in the Philippines, General Santos City, served as a validation test data set for the SPA algorithm and the construction of the RTFS mapping.

5.2 Global Optimization Methods

In this section, we briefly introduce three different global optimization methods used for the numerical experiments of solving the multistage stochastic facility location with risk of failure due to a natural disaster using the SPA algorithm. Although, in all fairness, the SPA algorithm may be aided by any global optimization methods, these three methods and algorithms were chosen as they seem to perform well in the extant literature in comparison to other methods. For the MSFLP, the global optimization method tunes the shadow prices of flood-prone facilities, λ^u_{it}, iteratively through evaluations of the objective function.

5.2.1 Stochastic Gradient Descent

Stochastic gradient descent (SGD) is an iterative method to optimize a differentiable or subdifferentiable objective function (33). At large, it is a stochastic approximation of gradient descent algorithm, as it reinstates the actual gradient i.e. computed from the total set of data, by an estimate computed from a subset of the data as randomly chosen. SGD aids notably in optimization problems with high dimension in reducing the computational strain which achieves faster iterations as offset for a lower convergence rate.

5.2.2 Bayesian Optimization

Bayesian Optimization (BayOpt) is an algorithm that optimize objective functions which are intensive and expensive to evaluate (34). It is befitting for optimization problems with continuous domains of around 20 dimensions, and stochasticity-tolerant in function evaluations. BayOpt constructs a surrogate for the objective function and quantifies the stochasticity in that surrogate using a Bayesian machine learning technique, Gaussian process regression, and then uses an acquisition function defined from this surrogate where decision to sample is made.

5.2.3 Covariance Matrix Adaptation Evolution Strategy (CMA-ES)

The *covariance matrix adaptation evolution strategy* (CMA-ES) is a powerful stochastic, derivative-free, and population-based search method to solve difficult numerical optimization problems, e.g. non-convex, non-smooth, noisy problems, etc. in continuous domain. It is considered as the *state-of-the-art* in evolutionary computation and adopted in multiple applications and fields. See (35) for a detailed review.

CMA-ES optimizes a continuous black-box function, i.e. the problem's objective function, $f : \mathbb{R}^n \longrightarrow \mathbb{R}$ by sampling from a non-stationary multivariate normal search distribution $\mathcal{N}(m^{(g)}, {\sigma^{(g)}}^2, C^{(g)})$, with mean $m^{(g)}$ (center), step-size ${\sigma^{(g)}}^2$ (scale), and covariance matrix $C^{(g)}$ (shape).

5.3 Numerical Results

In this section, we show the results of the numerical experiments of employing SPA with different global optimization methods to solve MSFLP with failure under risk of flooding. We have provided the local government unit with a spreadsheet tool containing vaccine facility information sheet, population information sheet, and distance information sheet which we used for the numerical experiment. Appendix A shows these three different spreadsheets. The RTFS mapping for Cagayan de Oro City, discussed in Section 5.1, is also adopted.

Initialized with shadow prices $\lambda^{u,0} \equiv 0$ and best parameter M tuned iteratively through the SPA algorithm, Figure 3 shows the evolution of the mean objective function value of the three susceptibility cases discussed in Section 5.1 using three global optimization methods.

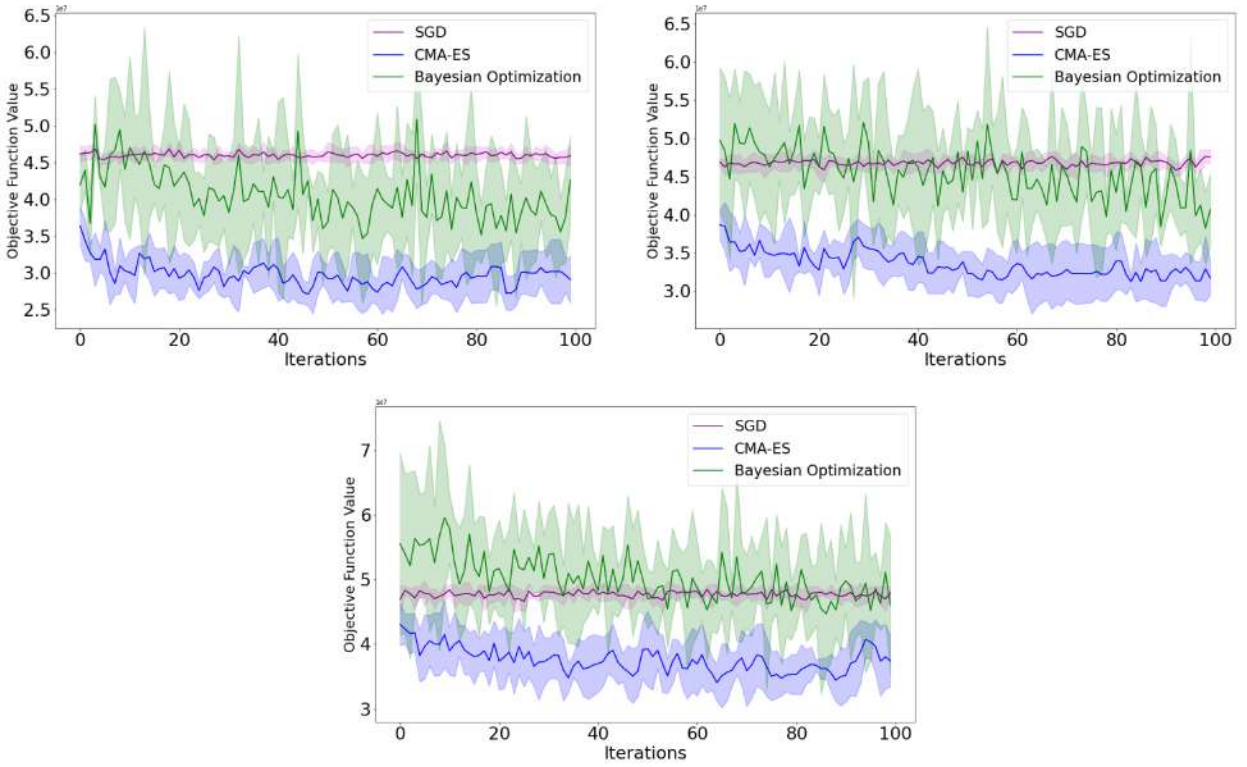

Figure 3: Evolution of the objective function value in three different cases - *Low* (top-left), *Medium* (top-right), and *High* (bottom) conservativity case for Cagayan de Oro City using three global optimization methods.

Interestingly, the MSFLP solved by SPA with SGD does not constitute any decrease of the objective function value across the three conservativity cases. The objective function value maintains rather smoothly around a local optimum even with the use of *adaptive gradient algorithm* (AdaGrad) step-sizes, a popular addition to stochastic gradient descent to re-compute step-sizes on the fly with the given gradients (36). Nonetheless, it is imperative to show how SGD performs in solving MSFLP as SGD is generally inexpensive and rather popular in the extant literature.

Furthermore, it is notable that both global and derivative-free optimization methods, CMA-ES and BayOpt, perform better than SGD. Both yield approximately 30-40% lower cost than a baseline approach that does not consider the risk of flooding (deterministic approach) i.e. where $\lambda^u_{it} = 0, \forall t = 1, \ldots, T, i = 1, \ldots, F$, across the different conservativity cases. This intently reduces, not only the total cost of operation of the vaccination campaign of the city, but also of its duration. Without the inclusion of the risk of failure due to flooding, the chosen facilities will be the same from the start until the end of the planning horizon.

Different parameters for both CMA-ES and BayOpt were investigated - different M-values, varied number of iterations, various population size for CMA-ES, and changing lower- and upper-bounds for BayOpt. The best parameter configurations for both global optimization methods were tested multiple times to set lower and upper 95% confidence interval as shown by the filled gaps in Figure 3 with the darker line as the mean.

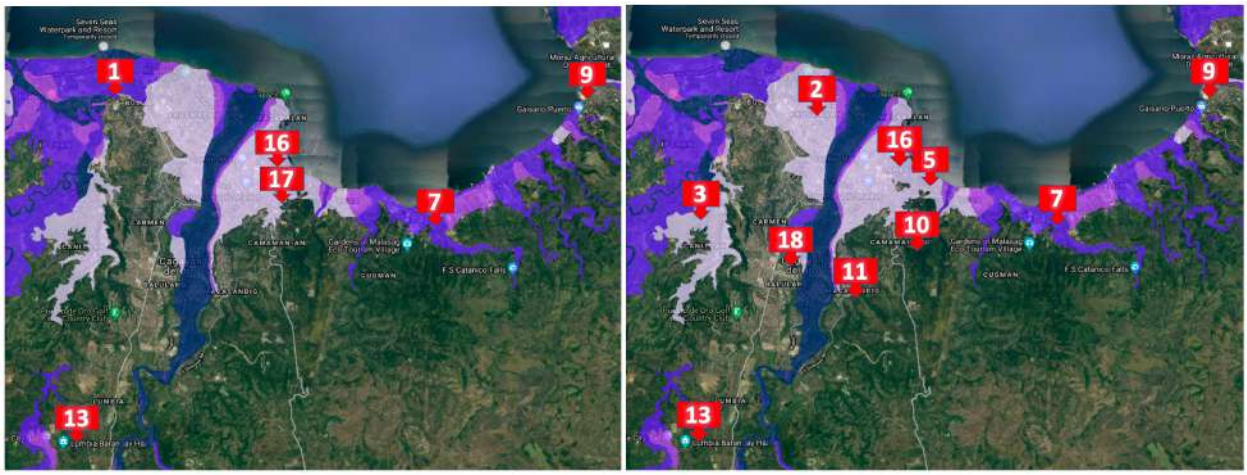

Figure 4: Chosen facilities when a) $\lambda^u = 0$ (left) and when b) $\lambda^u = \lambda^{u,N}$ (right) for Cagayan de Oro City.

Figure 4 shows the geographical maps of chosen optimal facilities with and without hedging for the risk of disruption due to flooding of *High* conservativity case. As demonstrated, our solution approach lead to a reduction of 30-40% of the overall vaccination roll-out operation cost in contrast to ignoring the risk of natural disaster occurrence. Figure 4 shows that opening more facilities reduces the total cost of the vaccination program of the city. The choice of opening facilities, without hedging for the risk of flooding, may incur additional costs in the year-long vaccination campaign since these facilities will possibly be flooded. Thus, they need to be re-opened multiple times across the planning horizon.

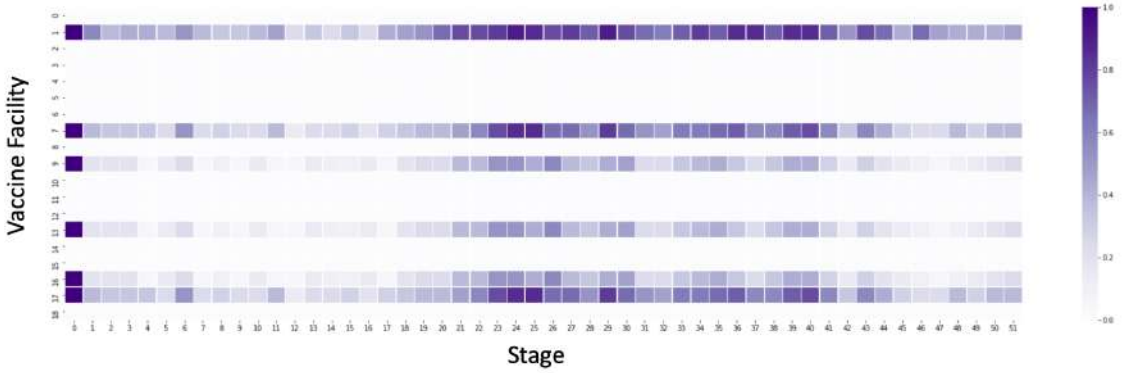

[a]

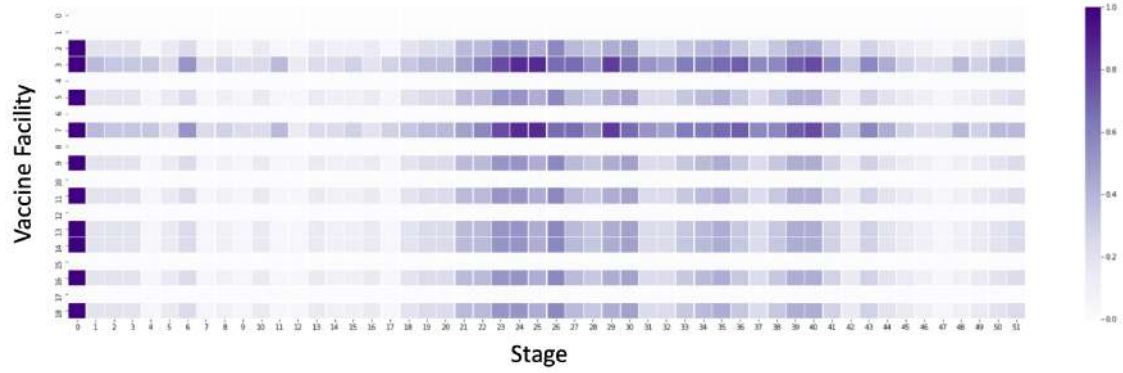

[b]

Figure 5: Heatmap of chosen facilities when a) $\lambda^u = 0$ and when b) $\lambda^u = \lambda^{u,N}$ for Cagayan de Oro City.

In Figure 4a (left), two out of six facilities are in the very high susceptibility areas and two other facilities are in the moderate susceptibility areas. In contrast to that, Figure 4b (right) shows that only one out of ten facilities are in the very high susceptibility area and four are in the moderate to low susceptibility areas, and the rest are in the safe areas.

Additionally, Figure 5 shows the heatmap of the average re-opening of each facility. As observed, chosen facilities for policy when $\lambda_{it}^{u} = 0,\ \forall t = 1, \ldots, T, i = 1, \ldots, F$, Figure 5a tend to be re-opened more than the policy when $\lambda_{it}^{u} = \lambda_{it}^{u,N},\ \forall t = 1, \ldots, T, i = 1, \ldots, F$, Figure 5b. Every time a facility is flooded and re-opened, a corresponding cost is incurred. Thus, if there are more dark hues of violet in the heatmap for a policy, then more facilities are to be re-opened leading higher costs.

Moreover, given the data-driven characteristic of SPA, we used cross-validation to avoid over-fitting. First, we executed 7-fold cross validation where the 21-year RTFS data set was split into seven training sets. Then, we held out one subset at a time and trained the algorithm on the remaining set to obtain $\lambda_{it}^{u,N},\ \forall t = 1, \ldots, T, i = 1, \ldots, F$. We then evaluated the model on the held-out subset using the obtained $\lambda_{it}^{u,N},\ \forall t = 1, \ldots, T, i = 1, \ldots, F$. This step was repeated across all the different subset folds. Furthermore, we were interested in how well our approach would perform on unseen data. To achieve this, we tested our approach on data for another city in the Philippines, General Santos City.

5.4 Test/Validation Data Set: General Santos City (Philippines) Instance

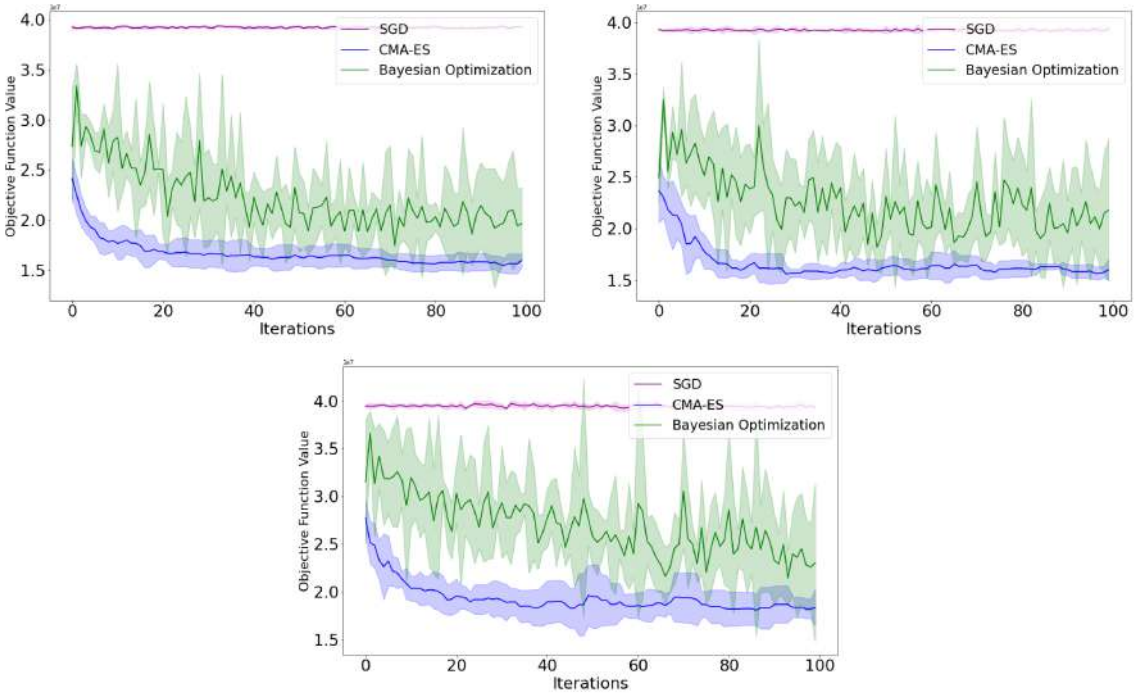

Figure 6: Evolution of the objective function value in three different cases - *Low* (top-left), *Medium* (top-right), and *High* (bottom) conservativity case for General Santos City using three global optimization methods.

Akin to the first city, MSFLP solved with SPA algorithm yields lower cost, approximately 30-40% lower, across the different cases i.e. $\lambda_{it}^{u} = 0,\ \forall t = 1, \ldots, T, i = 1, \ldots, F$, for the second city. The evolution of the objective function value is shown in Figure 6. Furthermore, Figure 7 shows the geographical map of chosen optimal facilities with and without hedging for the risk of failure due to flooding of *High* conservativity case.

Looking at Figure 7a (left), three vaccine facilities (Facilities 7, 11, and 25) are chosen when the risk of flooding is not considered, i.e. $\lambda_{it}^{u} = 0,\ \forall t = 1, \ldots, T, i = 1, \ldots, F$. However, when hedging from flooding risks, eight new vaccine facilities (Facilities 1, 3, 5, 8, 14, 18, 20, and 24) are chosen, shown in at Figure 7b (right). Although this is counter-intuitive, the eight new vaccine facilities with policy $\lambda_{it}^{u} = \lambda_{it}^{u,N},\ \forall t = 1, \ldots, T, i = 1, \ldots, F$ are less risky of flooding in contrast to the policy without hedging from natural disaster risk. Thus, lesser cost is incurred following the risk-averse policy.

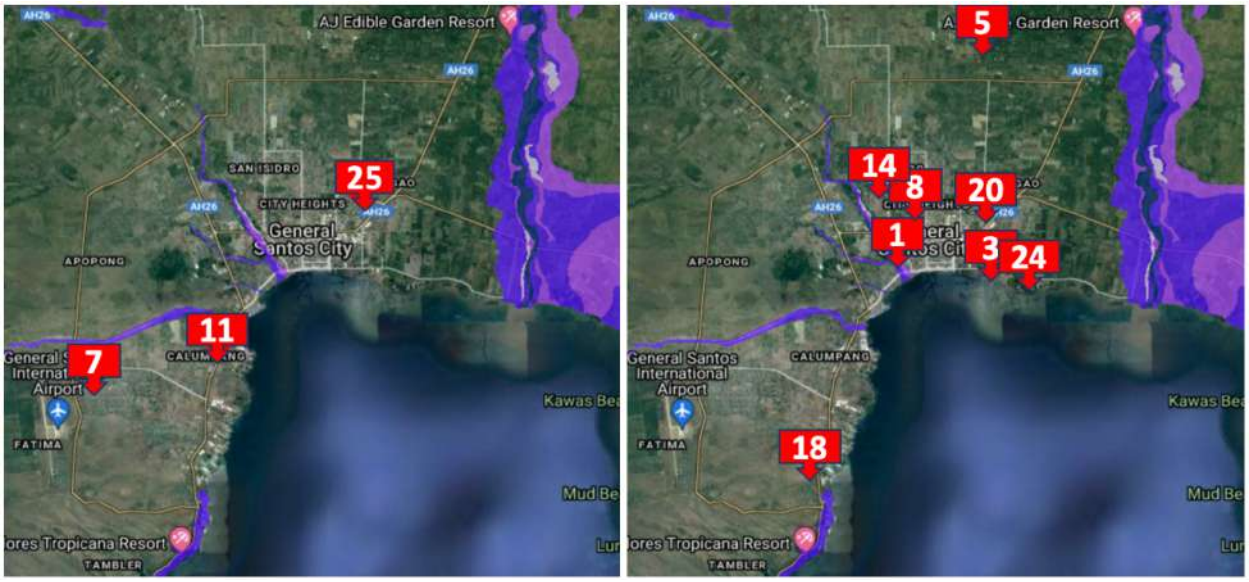

Figure 7: Chosen facilities when a) $\lambda^u = 0$ (left) and when b) $\lambda^u = \lambda^{u,N}$ (right) for General Santos City.

6 Conclusion and Outlook

In this paper, we consider a multistage facility location model for vaccine distribution which integrates the risk of disruption or failure of a vaccine facility due to natural disaster, i.e. flooding. Through an introduction of shadow price, a new cost parameter which approximates the cost of opening flood-prone vaccine facilities, we solve the multistage stochastic facility location problem (MSFLP) through a proposed algorithm, *shadow price approximation* (SPA), and with the aid of a global optimization method e.g. *stochastic gradient descent* (SGD), *covariance matrix adaptation evolution strategy* (CMA-ES), and *Bayesian optimization* (BO). The SPA algorithm, together with the global optimization method, effectively tune the shadow cost of opening a flood-prone facilities and are trained with real weather data, thereby alleviating the need to fit a stochastic model.

SPA algorithm paired with CMA-ES or BayOpt algorithm yield approximately 30-40% lower cost than a baseline approach that does not consider the risk of flooding, across the different rainfall-to-flood susceptibility cases. An interesting and counter-intuitive insight from the results show that not only are vaccine sites opened in less flood-prone areas but also more vaccine sites are being opened by the risk-averse solution. This intently reduces, not only the total cost of operation of the vaccination campaign of the city, but also of its duration. Without the inclusion of the risk of failure due to flooding, the chosen facilities will be the same from the start until the end of the planning horizon, which is more expensive.

Cagayan de Oro, a large city in the Philippines, which is prone to multiple flooding events per year, provided the data instance for the vaccination campaign costs. Weather and rainfall data were provided by a government agency and an industrial weather company. After having presented the results to the local government of the aforementioned city, our study was used as a basis to request for the independent handling of the inoculation process of the city and be provided with more vaccine supply in order to achieve herd immunity as fast as possible. The request was approved by the Chair of the COVID-19 Vaccine Cluster of the Philippines and by the end of Year 2021, Cagayan de Oro City posted one of the highest vaccination rates in the country and the highest in the region. To test the validity of the model and the solution method, another large city in the Philippines was tested.

Finally, there are several interesting directions to investigate for future research. Flooding is not the only natural disaster that occur in different nations; tornadoes, storm surges, landslides, earthquakes, and volcanic eruptions daunt all countries and can disrupt long-term vaccination campaigns and other projects. We may extend the formulation of the model to include other sources of natural disasters to expand the scope further. Other gradient- and gradient-free optimization methods may also be investigated further as the black-box tool aiding the *shadow price approximation* algorithm.

References

[1] WHO. Who covid-19 dashboard, 2022.

[2] Hannah Ritchie and Max Roser. Natural disasters. *Our World in Data*, 2014. https://ourworldindata.org/natural-disasters.

[3] Philstar. Galvez: 'odette' made gov't miss yearend vaccination target. *Philstar Global*, 2021.

[4] Newspaper The Economist. The world in brief, january 31st 2022. *The Economist*, 2022.

[5] About CDO. The recent floods in cagayan de oro has precedence that goes way back to the 1970's, 2022.

[6] DOH-10. Region 10 covid-19 situational report, 2022.

[7] Saeed Ghadimi, Raymond T Perkins, and Warren B Powell. Reinforcement learning via parametric cost function approximation for multistage stochastic programming. *arXiv preprint arXiv:2001.00831*, 2020.

[8] Susan Hesse Owen and Mark S Daskin. Strategic facility location: A review. *European journal of operational research*, 111(3):423–447, 1998.

[9] M Teresa Melo, Stefan Nickel, and Francisco Saldanha-Da-Gama. Facility location and supply chain management–a review. *European journal of operational research*, 196(2):401–412, 2009.

[10] Lawrence V Snyder. Facility location under uncertainty: a review. *IIE transactions*, 38(7):547–564, 2006.

[11] Dimitris Bertsimas, Vassilis Digalakis Jr, Alexander Jacquillat, Michael Lingzhi Li, and Alessandro Previero. Where to locate covid-19 mass vaccination facilities? *Naval Research Logistics (NRL)*, 69(2):179–200, 2022.

[12] Lotty Evertje Duijzer, Willem Van Jaarsveld, and Rommert Dekker. Literature review: The vaccine supply chain. *European Journal of Operational Research*, 268(1):174–192, 2018.

[13] Amir Ahmadi-Javid, Pardis Seyedi, and Siddhartha S Syam. A survey of healthcare facility location. *Computers & Operations Research*, 79:223–263, 2017.

[14] Xueping Li, Zhaoxia Zhao, Xiaoyan Zhu, and Tami Wyatt. Covering models and optimization techniques for emergency response facility location and planning: a review. *Mathematical Methods of Operations Research*, 74(3):281–310, 2011.

[15] Lawrence V Snyder, Zümbül Atan, Peng Peng, Ying Rong, Amanda J Schmitt, and Burcu Sinsoysal. Or/ms models for supply chain disruptions: A review. *Iie Transactions*, 48(2):89–109, 2016.

[16] Mengshi Lu, Lun Ran, and Zuo-Jun Max Shen. Reliable facility location design under uncertain correlated disruptions. *Manufacturing & Service Operations Management*, 17(4):445–455, 2015.

[17] Michael K Lim, Achal Bassamboo, Sunil Chopra, and Mark S Daskin. Facility location decisions with random disruptions and imperfect estimation. *Manufacturing & Service Operations Management*, 15(2):239–249, 2013.

[18] Chun Cheng, Yossiri Adulyasak, and Louis-Martin Rousseau. Robust facility location under disruptions. *INFORMS Journal on Optimization*, 3(3):298–314, 2021.

[19] Xian Yu, Siqian Shen, and Shabbir Ahmed. On the value of multistage stochastic facility location with risk aversion. *arXiv preprint arXiv:2105.11005*, 2021.

[20] Stefan Nickel, Francisco Saldanha-da Gama, and Hans-Peter Ziegler. A multi-stage stochastic supply network design problem with financial decisions and risk management. *Omega*, 40(5):511–524, 2012.

[21] Patricio Hernandez, Antonio Alonso-Ayuso, Fernanda Bravo, Laureano F Escudero, Monique Guignard, Vladimir Marianov, and Andres Weintraub. A branch-and-cluster coordination scheme for selecting prison facility sites under uncertainty. *Computers & operations research*, 39(9):2232–2241, 2012.

[22] Maria Albareda-Sambola, Antonio Alonso-Ayuso, Laureano F Escudero, Elena Fernández, and Celeste Pizarro. Fix-and-relax-coordination for a multi-period location–allocation problem under uncertainty. *Computers & operations research*, 40(12):2878–2892, 2013.

[23] Kai Huang and Shabbir Ahmed. The value of multistage stochastic programming in capacity planning under uncertainty. *Operations Research*, 57(4):893–904, 2009.

[24] Jikai Zou, Shabbir Ahmed, and Xu Andy Sun. Stochastic dual dynamic integer programming. *Mathematical Programming*, 175(1):461–502, 2019.

[25] Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

[26] Nils Löhndorf and Alexander Shapiro. Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research*, 273(2):650–661, 2019.

[27] Karl Inderfurth. Multistage safety stock planning with item demands correlated across products and through time. *Production and Operations Management*, 4(2):127–144, 1995.

[28] Fabricio Nápoles-Rivera, Ma Guadalupe Rojas-Torres, José María Ponce-Ortega, Medardo Serna-González, and Mahmoud M El-Halwagi. Optimal design of macroscopic water networks under parametric uncertainty. *Journal of Cleaner Production*, 88:172–184, 2015.

[29] Robert L Burdett and Erhan Kozan. Techniques to effectively buffer schedules in the face of uncertainties. *Computers & Industrial Engineering*, 87:16–29, 2015.

[30] Kelvin Mabao and Ruth G Cabahug. Assessment and analysis of t he floodplain of cagayan de oro river basin. *Mindanao Journal of Science and Technology*, 12, 2014.

[31] PHIVOLCS and DOST. Hazardhunterph, 2019.

[32] DOST-PAGASA. Rainfall advisories, classification, and measurement, 2019.

[33] Léon Bottou and Olivier Bousquet. 13 the tradeoffs of large-scale learning. *Optimization for machine learning*, page 351, 2011.

[34] Martin Pelikan, David E Goldberg, Erick Cantú-Paz, et al. Boa: The bayesian optimization algorithm. In *Proceedings of the genetic and evolutionary computation conference GECCO-99*, volume 1, pages 525–532. Citeseer, 1999.

[35] Nikolaus Hansen. The cma evolution strategy: a comparing review. *Towards a new evolutionary computation*, pages 75–102, 2006.

[36] Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes, from any initialization. *arXiv preprint arXiv:1806.01811*, 2, 2018.

APPENDICES

Appendix A: Spreadsheet Tool for Data Collection

Through a spreadsheet tool we have created as a data template, the local governments of the two cities provided us with the necessary cost parameters (e.g. fixed cost f_i to open vaccine facility i, cost s of hiring a vaccine administrator, distance cost d_{ij} from vaccine facility i to district j), capacity C_i of each vaccine facility i, and the vaccinating population P_j of each district j. Appendix A shows screenshots of the spreadsheet tool given to the local government for data collection.

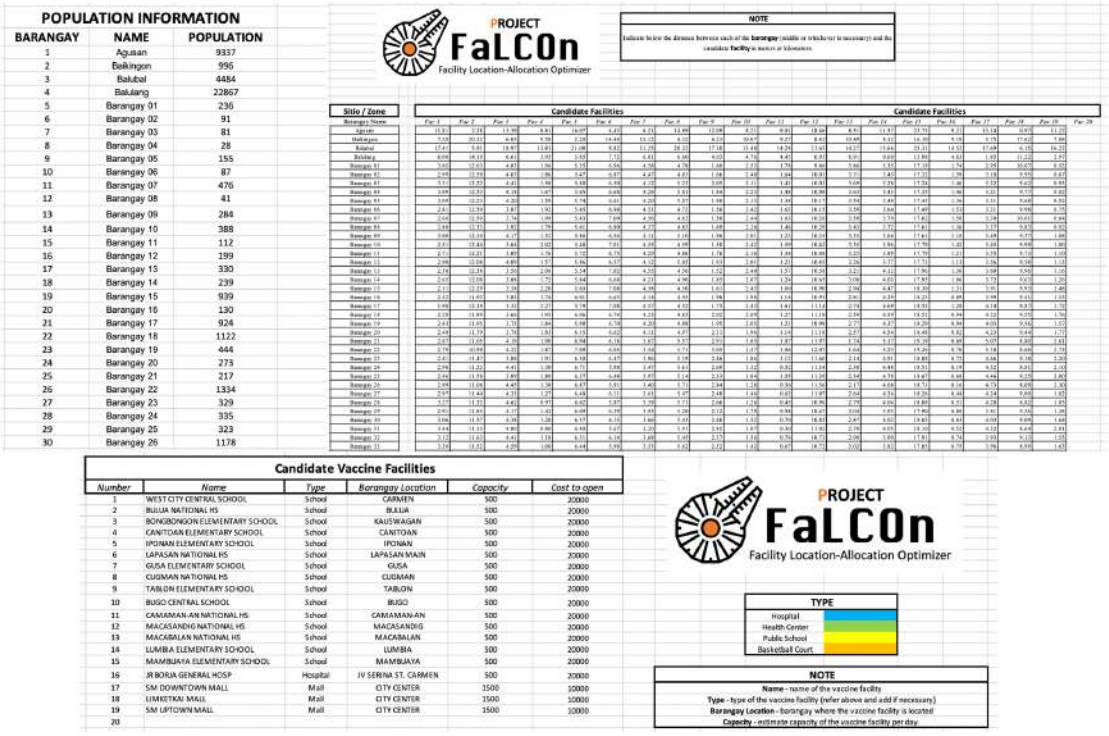

Figure 8: Spreadsheet Tool for Data Collection

Appendix B: How different values of the M-coefficient affects the objective function value

We tested different choices of tunable parameter M, starting with $M = 1$ and increasing the value incrementally up to an upper bound of the shadow price which equals the total expected cost when $\lambda_{it}^u = 0$. The following figures show how different values of the M-coefficient (x-axis) affects the objective function value (y-axis). As seen, there are optimum values of the M-coefficient which coincide to f_{it}, the fixed cost of opening a vaccine facility i at time t, or higher.

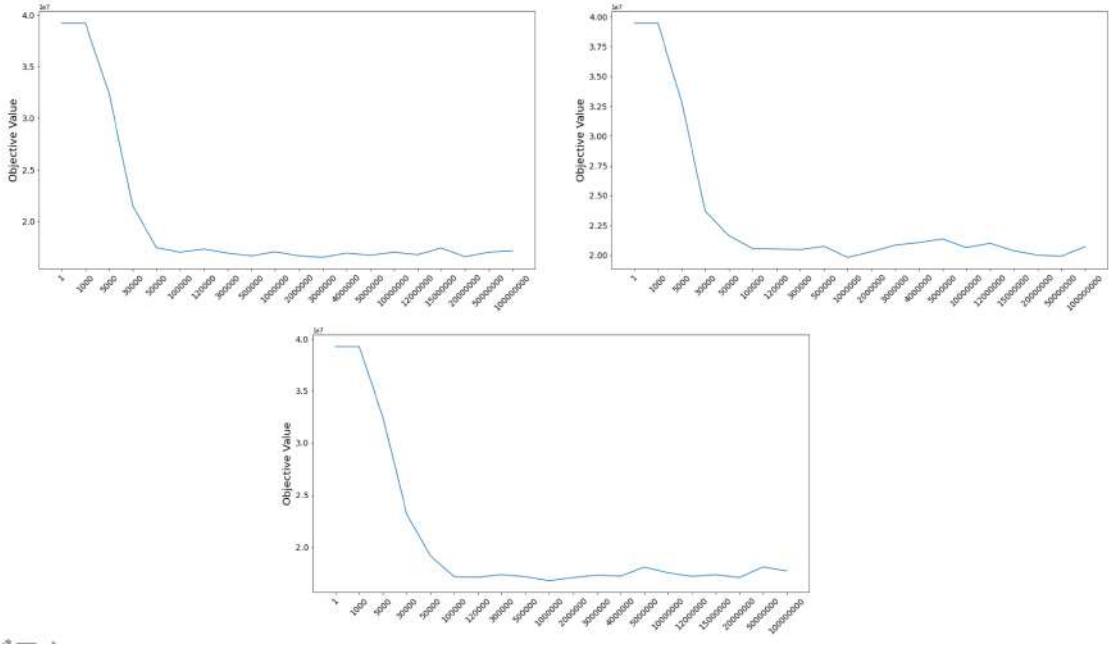

Figure 10: Different values of the *M*-coefficient in relation to the objective function value for City 2

A *rule-of-thumb* and starting point where the best *M*-coefficient values which yield the lowest possible cost starts when *M*-coefficient is at least as high as f_i, the cost of opening facility i, or higher. When *M*-coefficient value is less than the opening cost f_i, there is less, even possibly, no change from when $\lambda_{it}^u = 0$. The *M*-coefficient value, for obvious reasons, affects the tuning of λ_{it} as *M*-coefficient serves as the mean value of cost of failure of facility i due to a natural disaster.

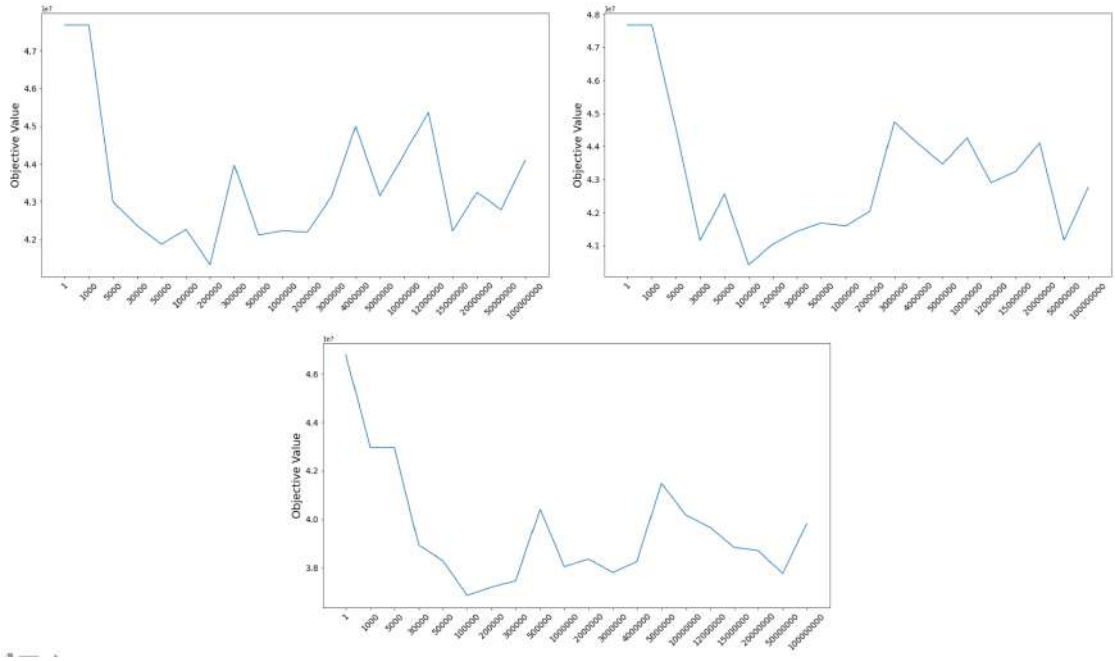

Figure 9: Different values of the *M*-coefficient in relation to the objective function value for City 1