

# Models and Algorithms for Balancing Efficiency and Equity in Vaccine Allocation

Jeffrey Keithley<sup>1</sup>, Madeline Bonner<sup>2</sup>, Sriram V. Pemmaraju<sup>1</sup>

<sup>1</sup>University of Iowa

<sup>2</sup>University of Northern Iowa  
 sriram-pemmaraju@uiowa.edu

## Abstract

The COVID-19 pandemic was a powerful reminder that existing societal inequalities get amplified during public health emergencies. In response to the pandemic, organizations such as the CDC, WHO, and public health departments developed frameworks for equitable allocation of vaccines, using well-established ethical principles as a foundation. The overall goal of this paper is to translate these policy frameworks into a computational framework that can be used by public health departments to equitably allocate vaccines in a transparent manner during the initial stages of a pandemic, when vaccine demand far exceeds supply.

We start by developing a mathematical model for disease-spread that accounts for social vulnerability, geographic barriers to healthcare access, and differences in work constraints. On the basis of this model, we present multiple optimization formulations of a vaccine allocation problem that aims to reduce overall disease prevalence while also reducing disparity in outcomes for a given “protected class” relative to the general population. We present simple, scalable, and transparent algorithms for our optimization formulations.

Our experiments focus on allocating vaccines at the census block group granularity in Johnson County, Iowa. Our experimental test bed incorporates social vulnerability index, a hospital accessibility index, and essential worker status into CovaSim, a state-of-the-art agent-based COVID-19 model. Our experiments lead to two main takeaways. First, it is possible to substantially reduce disparity in the outcomes of the protected class (for various choices of this class) with negligible worsening in overall disease-prevalence. Second, it is critical for disparity to be considered at all stages of the computational framework, e.g., incorporating it into the optimization formulation without considering it in the modeling stage has very limited value.

## Introduction

Equitable vaccine allocation, a critical challenge in public health, aims to balance its goals of minimizing disease spread and ensuring fair outcomes resulting for various, and especially vulnerable, subpopulations. This problem is particularly important during pandemics, where vaccines are often in short supply (Srivastava and Priyadarshini 2021; Liu

and Lou 2022), and disparities in access can exacerbate existing health inequalities (Khan et al. 2023; Gu et al. 2020; Pijls et al. 2021; Tipirneni et al. 2022). Effective vaccine allocation strategies must account for both epidemiological dynamics and societal considerations, such as protecting vulnerable populations and reducing disparities across demographic groups.

Tasked by the Centers for Disease Control (CDC) and the National Institutes of Health (NIH), the National Academies of Sciences, Engineering and Medicine published a report in 2020 (NASEM 2020), describing a framework for the equitable allocation of COVID-19 vaccines. It is worth noting that this report was written at a time when no COVID-19 vaccines had yet been approved by the US FDA. Thus, the recommendations of the report were primarily meant to be used in the early stages of the pandemic, when the demand for vaccines was far in excess of the supply. The report described 3 *ethical* principles that form the foundation of any allocation policy:

- **Maximum benefit:** the overall goal of vaccine allocation, namely the obligation to protect and promote public health and socioeconomic well-being in the short and long term.
- **Equal concern:** the requirement that every individual be treated as having “equal dignity, worth, and value.”
- **Mitigation of health inequities:** the requirement to explicitly take into account the higher burden of COVID-19 experienced by some subpopulations, including the elderly, rural, racial minorities, etc., due to greater susceptibility, exposure and worse access to healthcare.

In some settings, these principles may appear to be in tension with each other. For example, the report elaborates that the principle of equal concern does not prevent prioritizing people’s social roles (for example, emergency responder) or mitigating the greater burden of diseases experienced by certain subpopulations. To ensure that the public has an understanding of how these principles are applied and such tensions are addressed, the report also describes 3 *procedural* principles.

**Fairness:** requires engagement with all stakeholders, especially those most affected by the pandemic.

**Transparency:** requires clear, accurate, and timely communication about allocation policies, even as they are de-

veloped.

**Evidence-based:** requires that the allocation framework and policies are based on the best and most up-to-date scientific evidence.

Other agencies (e.g., WHO and Johns Hopkins University) developed similar frameworks (World Health Organization 2020; Toner et al. 2020) and subsequently states and localities in the US (NMDOH 2020; CDPH 2020; MichiganD-HHS 2020; TexasDSHS 2020) developed more specific and concrete versions better suited to their regional context.

The overall goal of this paper is to translate these ethical and procedural principles into a computational framework that can be used by regional public health departments to equitably allocate vaccines during the initial stages of a pandemic. Our specific contributions are as follows:

- (a) We present a fine-grained, data-driven model in which all individuals in a region (e.g., a county) are modeled as agents and their demographic features, household membership and location, job type and location, etc., are estimated from the American Community Survey Public Use Microdata Sample (PUMS) dataset from the US Census Bureau (USCensus 2025b). Using shared household and workplace membership along with techniques described in (Tulchinsky et al. 2024), we then construct an agent-based *contact network*. Using geographic locations of individuals' households, the model additionally incorporates individual social vulnerability and healthcare access – all of which are estimated from data sources (CDC 2021b; HRSA 2025; USCensus 2025b) – into a state-of-the-art agent-based model for the spread of COVID-19.
- (b) We focus on *outcome disparity* (as opposed to allocation disparity) and formalize this notion for a given *protected* class. In our context, outcomes refer to becoming infected, becoming severely infected and requiring hospitalization, or becoming critically ill and possibly dying from the infection. Focusing on outcome disparity is technically more challenging than focusing on disparity in allocations, but outcomes are what the main stakeholders, i.e., the public being vaccinated, care most about. Both our model and the optimization formulations (see below) take the protected class as given and are agnostic to this choice. The protected class can be defined in many different ways, e.g., using demographic features, geography, type of work, etc.
- (c) Our goal is to find an allocation of a limited supply of vaccines such that their overall effectiveness is maximized while disease-related outcomes for the protected class are not much worse than the outcomes for the overall population.

We present two alternate combinatorial optimization formulations of this problem. In the first formulation, we aim to maximize the benefit of vaccine allocation – defined as a linear combination of bad outcomes averted and outcome equity achieved for the protected class – while satisfying a budget constraint on the number of allocated vaccines. This formulation is inspired by the problem of maximizing a submodular function subject to a cardinality constraint (Nemhauser, Wolsey, and Fisher

1978). In the second formulation, we aim to minimize the cost of vaccine allocation – defined as a linear combination of outcome disparity achieved for the protected class and number of vaccines allocated – while satisfying a coverage constraint that ensures that sufficiently many bad outcomes are averted. This formulation is inspired by the submodular cost submodular cover (SCSC) problem (Wan et al. 2010; Crawford, Kuhnle, and Thai 2019)

- (d) We present simple, scalable greedy algorithms for both problem formulations. We also present a simulated annealing algorithm for our first formulation, which we evaluate in comparison with the greedy algorithm.
- (e) We experimentally evaluate our modeling approach, problem formulations, and algorithms for disparity-aware vaccine allocation in Johnson County, Iowa. Our experiments lead to two main takeaways. First, it is possible to substantially reduce disparity in the outcomes of the protected class (for various choices of this class) with negligible worsening in overall disease-prevalence. In fact, this is possible to achieve using the simple, greedy algorithms mentioned above. Second, it is critical for disparity to be considered at all stages of the computational framework. More specifically, incorporating disparity into the optimization formulation without considering it in the modeling stage has very limited value.

We have designed all components of our computational framework, i.e., modeling, optimization formulations, and algorithms, to be **i)** transparent, **ii)** flexible, and **iii)** extendable. More specifically, our framework can easily incorporate various geographic scales, different definitions of the protected class and disparity, arbitrary disease models, other optimization formulations, and other algorithms. All of this allows for easy adaptation of our framework to disparate public health scenarios. The framework is also designed to be extended easily for the addition of entirely new components, such as new vaccine allocation approaches or a more explicit incorporation of vaccine distribution channels. Our framework only uses publicly available data and all our code, along with detailed documentation, is included in the Supplementary Material<sup>1</sup>.

## Mathematical Modeling

We focus on allocating vaccines equitably in a given *region* of focus (e.g., a state or a county) that consists of a collection of smaller *subregions* (e.g., census tracts, census block groups<sup>2</sup>). Within our region of focus, each subregion refers to a geographical area along with its residents, infrastructure, homes, schools, workplaces, and hospitals.

We start by proposing a disparity-aware, agent-based, computational framework for disease-spread, vaccine delivery, and response. Our model consists of the following 4

<sup>1</sup>Our code is available at <https://zenodo.org/records/16783299>

<sup>2</sup>The United States is organized into a hierarchical system of geographic units: each state is divided into counties, which are subdivided into census tracts; these tracts are further divided into census block groups, which are composed of individual blocks (ProximityOne 2024).

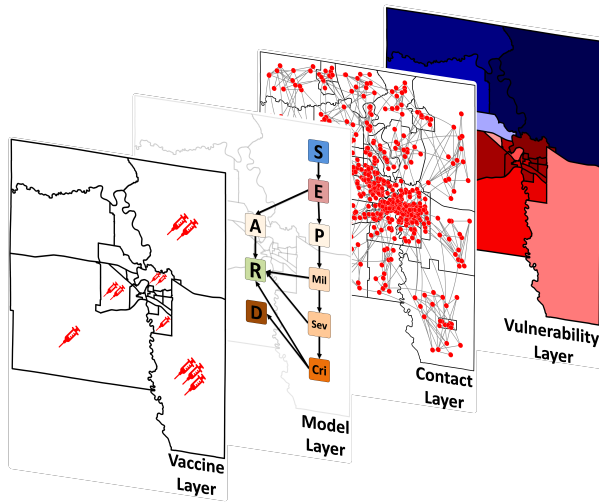


Figure 1: The four layers of disparity-aware, agent-based, computational framework for disease-spread, vaccine delivery, and response.

layers (see Figure 1): (1) indices of health-associated vulnerability at the subregion level, (2) individual-level contact network for the region’s population, (3) agent-based disease-spread model for COVID-19 incorporating health-associated vulnerability indices, and (4) a vaccination delivery and response model. We describe each of these layers separately below.

### Health-associated Vulnerability Indices

We use two indices in our work.

- **Socioeconomic Social Vulnerability Index:** The *Social Vulnerability Index (SVI)* is a widely used measure developed by the Agency for Toxic Substances and Disease Registry within the CDC that indicates the relative vulnerability of each census tract in the United States (CDC 2021b). The SVI is used to identify communities that are particularly susceptible to be disproportionately damaged in events such as natural disasters and disease outbreaks. SVI values are calculated using a wide range of community characteristics, from data on socioeconomic status, household characteristics, racial and ethnic minority status, and housing type and transportation. A relationship between SVI values and population-level health outcomes in clinical, surgical, mortality, or health promotion areas is well-established (Higginbotham et al. 2025). Previous research has also demonstrated that SVI is positively correlated with COVID-19 prevalence (Biggs et al. 2021; Karaye and Horney 2020). Figure 2A shows SVI values at the census tract granularity for Johnson County, Iowa. The figure illustrates that central and southern regions have a higher relative SVI (corresponding to worse health outcomes), and we note that the southwest corner of Johnson County is a particularly rural part of the county.
- **Treatment Accessibility Index:** The *Treatment Accessibility Index (TAI)* is an index that measures accessibil-

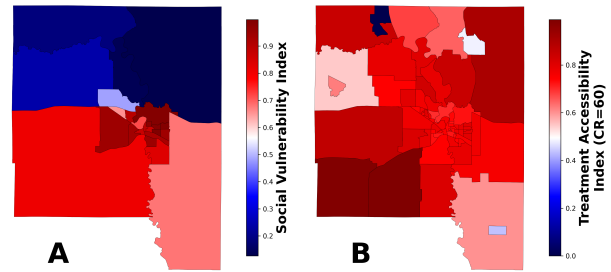


Figure 2: Vulnerability index values for Johnson County, Iowa, where a higher value indicates a higher relative vulnerability. The subplots show (A) SVI values at a census tract granularity (24 census tracts), and (B) TAI values at a census block group granularity (72 census block groups) with a catchment radius of 60 minutes.

ity of healthcare for each census block group. We use a Gaussian two step floating catchment area (Ga2SFCA) (Ye et al. 2024) method, which combines hospital capacity and population demand within travel-time catchment areas, weighted by a Gaussian decay function to reflect distance (the derivation of which may be found in the full paper (Keithley, Bonner, and Pemmaraju 2025)). The score for each census block group then takes all the hospitals within a distance threshold (including hospitals outside the region of focus, lying in neighboring regions) into account, weighted by their distance to the centroid of the census block group. Here, we use hospital location data from (HRSA 2025), which includes precise location and capacity. The “distance” in this context is the time it takes to drive from one point to another, which is determined from (OSRM 2010; Luxen and Vetter 2011). From Figure 2B, it is clear that the value of TAI, like that of SVI, is relatively high in the southwest corner of the county. The north central part of Johnson County has relatively easy access to hospitals in Ceder Rapids, Iowa, which is just to the north of Johnson county (USCensus 2021).

### Contact Network

A high-quality representation of the contact patterns within a region is crucial to effectively model the spread of disease within that region with high-granularity (Tulchinsky et al. 2024). We model the contact patterns of Johnson County, Iowa, as a contact network, which we refer to as  $G = (V, E)$ . Each node in  $G$  represents an individual with an associated set of attributes, such as age, sex, industry type, and grade level in school. The region of focus contains locations for entities such as homes, schools, workplaces, and hospitals. Each individual is associated with a set of these entities, such as the home they reside in or the school they go to. In essence, a contact network must reflect the following:

- Population level statistics and geographic distributions of attributes such as age, size of household, occupation, and race/ethnicity.
- Relative levels of contact within and between demo-

graphic groups.

We use GreasyPop (Tulchinsky et al. 2024), which is an open-source package designed to construct realistic individual-level contact networks. GreasyPop takes the following steps to reflect the required elements of the contact network (described above):

- (1) The first step in creating the contact network is to construct a synthetic population which fulfills the requirements of (i). GreasyPop achieves this by using simulated annealing (Kirkpatrick, Gelatt, and Vecchi 1983) to match the joint distributions of the aforementioned population attributes, the specific aggregated statistics of which are provided by Public Use Microdata Sample (PUMS) data (USCensus 2025b).
- (2) The synthetic population is then used to build the set of pairwise interactions between nodes (individuals) to form the *edges* of the contact network. GreasyPop constructs these edges conditioned on previously established levels of contact between individuals specific to age (Prem, Cook, and Jit 2017; Prem et al. 2021), school settings (Mossong et al. 2008), and varying types of workplaces (Gaffney et al. 2023).

The result is a contact network that reflects the specific population distributions and levels of contact in a geographical region. For this paper, the contact network is static (unchanging over time), undirected (if node A has contact with node B, then the reverse is also true), and unweighted (no values are associated with edges).

Property	Value
Number of nodes	145,984
Number of edges	920,034
Average degree	6.3
Number with high (> 0.8) SVI	15,898
Number with high (> 0.8) TAI	23,141
Number of essential workers	34,579

Table 1: Key properties of the synthetic population and contact network

In addition, we amplify the relative level of contact within workplaces classified as “essential”. In general, industries are classified using the North American Industry Classification System (NAICS) (USCensus 2025a), where a two-digit NAICS code represents a general industry, such as mining (21) or information services (51). For this paper, we classify the following groups as high-exposure essential workers according to (CDC 2021a; OSHA 2025): health care and social assistance (NAICS code 62), accommodation and food services (NAICS code 72), retail trade (NAICS codes 44,45), and transportation, warehousing, and utilities (NAICS codes 48, 49, 22). We increase the relative contact rates within essential workplaces through the following steps: **i**) identify subgraphs corresponding to each essential industry in each census block group, and **ii**) add contacts within each of those subgraphs to increase their relative contact rates.




Figure 3: Disease state transitions and their probabilities. The possible states for each individual are Susceptible, Exposed, Presymptomatic, Asymptomatic, Mild infection, Severe infection, Critical infection, Recovered, and Dead. (inspired by (Kerr et al. 2021) Figure 2)

## Disparity-Aware Disease-Spread Model

One of the most commonly used methods for modeling the spread of disease is the *homogeneous-mixing compartmental model* (Kermack and McKendrick 1927), which works by partitioning a population into disease states (called compartments), and defining rules for the rate at which the population changes between disease states. For instance, the rate at which susceptible individuals become infected at a given point in time is determined by a parameter governing the transmissibility of the disease and the concentration of already infected individuals in the population. While this classical model assumes *homogeneous* mixing, we use a *compartmental agent-based model (ABM)* that assumes that individuals transition between different disease states and “mixing” within the population occurs according to the contact network described earlier (Kiss, Miller, and Simon 2017; Germann et al. 2006; Bissett et al. 2021).

**Model Specification.** We employ CovaSim (Kerr et al. 2021), which is a sophisticated COVID-19 ABM that has been used extensively to support policy decisions (Cohen et al. 2020; Panovska-Griffiths et al. 2020; Scott et al. 2020). This model assumes that each individual who is initially *susceptible* may get infected, at which point they are considered *exposed* for an amount of time sampled from a log-normal distribution (known as the *latent* period). At the end of the latent period, they become *mildly symptomatic* with a parameterized probability, otherwise they proceed to a *asymptomatic infection* state.

**Infectivity Dynamics.** CovaSim uses a base infectivity value  $\beta$  as the probability that an infected individual will successfully infect a susceptible individual given a contact between them. In addition, the viral load of each individual scales the infectiousness. We denote the viral load of individual  $v$  at  $\tau$  days since infection as  $\xi_v^\tau$ , drawn from a negative binomial distribution with mean 1.0 and shape parameter 0.45, implying that the viral load is initially high before falling (Kerr et al. 2021). CovaSim scales the probability each susceptible individual  $v$  will get infected by an individual-specific value referred to as the relative susceptibility, denoted  $r_v^{\text{sus}}$ . By default, the age of each individual is one attribute that contributes to their relative susceptibility. For example, individuals who are 80+ years of age are 1.47 times more likely to get infected given a contact with an infected individual than their 20-59 year old counterparts. In summary, an infected agent  $u$  will successfully infect sus-

ceptible agent  $v$  given a contact between them with probability  $\beta \cdot \xi_u^\tau \cdot r_v^{\text{sus}}$ . Note that here “infected” refers to any individual who enters the Exposed compartment, independent of whether they show symptoms.

In addition to susceptibility, individual *infection severity* can be adjusted as well. Individuals who are symptomatic may progress to a severe infection state, then a critical infection state, and finally death if they do not recover. Each mildly infected individual will become *severely infected* with probability  $p_{\text{sev}}$ , each severely infected individual will become *critically infected* with probability  $p_{\text{cri}}$ , and each critically infected individual will die with probability  $p_{\text{dea}}$ . The durations between those four states is sampled from a log-normal distribution, and each individual in an infected state can transition to recovery if they do not proceed to each subsequent infection state. These transitions are shown in Figure 3.

**Incorporating Vulnerability into CovaSim.** In this section, we describe how the SVI and TAI definitions of vulnerability (described above) are incorporated into the CovaSim disease model.

- (1) First, each individual is assigned a *vulnerability score* according to the vulnerability index value associated with the subregion where they reside. When using SVI as a basis for the protected class definition, we define  $s_v^{\text{SVI}}$  to be the normalized SVI value for  $v$ ’s home census tract<sup>3</sup>. Similarly, we define  $s_v^{\text{TAI}}$  to be the normalized TAI value for  $v$ ’s home census block group when using TAI as a basis for the protected class.
- (2) We then define a function to transform a vulnerability score into a value that can be incorporated into the disease model.  $\mathcal{L}(\cdot)$  is a parameterized sigmoid function, which increases above 1 with an increasing vulnerability score, and decreases below 1 once vulnerability dips below a certain threshold. The motivation behind this vulnerability scaling function is to incorporate disparity into the disease model without altering the overall prevalence as a result. A more comprehensive description of the function  $\mathcal{L}(\cdot)$  may be found in the full paper (Keithley, Bonner, and Pemmaraju 2025).
- (3) When using SVI as the basis for protected class, we scale the relative susceptibility  $r_v^{\text{sus}}$  (described above) of each individual  $v \in V$  by  $\mathcal{L}(\cdot)$ :  $r_v^{\text{sus}}$  as  $r_v^{\text{sus}} \leftarrow r_v^{\text{sus}} \cdot \mathcal{L}(s_v^{\text{SVI}})$ , where  $\mathcal{L}(\cdot)$  is described above.
- (4) Similarly, we scale the severity of disease for each individual  $v$  according the same function. CovaSim denotes the probability of each individual progressing from symptomatic to severe, severe to critical, and from critical to death as  $p_v^{\text{sev}}$ ,  $p_v^{\text{cri}}$ , and  $p_v^{\text{dea}}$ , respectively. Thus, we scale each of these values in the same manner as above, as  $p_v^* \leftarrow p_v^* \cdot \mathcal{L}(s_v^{\text{TAI}})$ .

## Vaccination Model

As mentioned earlier, this paper considers vaccine allocation at the very early stages of an outbreak. In other words, a small number of individuals in the population are initially

infected and this is known to the vaccine allocation algorithm. We assume that the region’s public health department is tasked with distributing vaccines to the regions’ population. We allow for two possible settings: (i) either the region’s public health department receives a fixed number of vaccines, specified by a *vaccine budget*  $D$  or (ii) there is no hard constraint on the number of available vaccines, but the public health department incurs a cost for each vaccine it receives. We assume that region’s population is partitioned into a collection  $\mathcal{U}$  of subpopulations. Let  $\mathbf{x} \in \mathbb{Z}_+^{\mathcal{U}}$  be a vaccine allocation, where each element of the allocation vector  $\mathbf{x}$  encodes the number of shipments of vaccines to allocated to each subpopulation. This partition  $\mathcal{U}$  of a region’s population into subpopulations can be based on geography, e.g., the population of a state can be partitioned into subpopulations, one per county in the state. But this partition could also be based on other characteristics of the individuals in the population, e.g., work type, age, etc. We assume that a random subset of each subpopulation takes advantage of the vaccines allocated to that subpopulation. This setting, in which a random subset of a subpopulation (instead of specific targeted individuals) is vaccinated has been called the *group vaccine allocation problem* (Zhang et al. 2016). We refer to the specific policy used to allocate vaccines within each subpopulation as a *local allocation policy*. Thus, a group vaccine allocation policy is composed of a number of local allocation policies, one for each subpopulation. In its simplest form a local allocation policy is just uniform, i.e., every individual in the subpopulation is equally likely to receive a vaccine. But, we may also consider non-uniform local allocation policies. For example, within each subpopulation we might require that the probability that individuals 65+ years of age receive a vaccine is 5 times the probability that younger individuals receive a vaccine. To simplify our setting, we assume that (i) vaccine uptake is 100%, i.e., all allocated vaccines are used up and (ii) immunity acquired from vaccination is perfect. Both of these assumptions are relatively simple to relax.

## Problem Formulations

One goal of this paper is to examine methods that allocate a limited number of vaccines to subpopulations within a region to decrease overall disease prevalence while ensuring that a protected group is not disproportionately affected.

Our approach to achieving this goal is through two mathematical formulations. First, we maximize the proportion of infections averted minus a weighted disparity subject to a given vaccine budget. In this formulation, the goal is to maximize the number of infections averted, but without making disparity too high. Our second formulation minimizes the “cost” of a vaccine allocation, which is defined as a linear combination of the disparity and number of vaccines allocated, while keeping the fraction of infections averted above a threshold. In this formulation, we don’t have a “hard” budget constraint; instead we view both disparity and the number of vaccines allocated as different types of cost. To further discuss these formulations, we need some definitions.

In addition, a subset  $P \subseteq V$  of the population is designated as the *protected class*. The choice of  $P$  is independent

<sup>3</sup>SVI data is only available at a census tract level granularity

dent of the partition of the population  $V$  into subpopulations. For example, the subpopulations may be defined by geography, while  $P$  may be defined by age or by type of work, e.g.,  $P$  may be the set of all essential workers in the population. We use  $f^I(\mathbf{x})$  to denote the *expected* proportion of infected individuals in the population, as a function of  $\mathbf{x}$ , the vaccine allocation vector<sup>4</sup>. In other words,  $f^I(\mathbf{x})$  counts (in expectation) all individuals who enter compartments *Mil* (Mild infection) or *A* (Asymptomatic infection). Note that the expectation here is over the stochasticity of the disease-spread model. Similarly, we define  $f^S(\mathbf{x})$ ,  $f^C(\mathbf{x})$ , and  $f^D(\mathbf{x})$  respectively as the expected proportion of the population that becomes severely infected, critically infected, or dies, given the vaccine allocation vector  $\mathbf{x}$ . Let  $g^I(\mathbf{x}) = 1 - f^I(\mathbf{x})/f^I(\mathbf{0})$  be the expected proportion of infections averted by vaccine allocation  $\mathbf{x}$ , where  $f^I(\mathbf{0})$  is the proportion of individuals infected with no vaccine allocation. Similarly, corresponding to the other versions of the function  $f$ , i.e.,  $f^S(\mathbf{x})$ ,  $f^C(\mathbf{x})$ , and  $f^D(\mathbf{x})$ , we define  $g^S(\mathbf{x})$ ,  $g^C(\mathbf{x})$ , and  $g^D(\mathbf{x})$  as the expected proportions of severe infections averted, critical infections averted, and deaths averted (respectively) by the vaccine allocation vector  $\mathbf{x}$ . Note that it is not too difficult to incorporate more sophisticated versions of these notions into our formulation. For example, in literature on optimizing vaccine allocation (Medlock and Galvani 2009; Emanuel and Wertheimer 2006), authors have considered assigning a weight to each individual based on the productive years cut short by death. For notational convenience, we drop the superscripts from the functions  $f(\mathbf{x})$  and  $g(\mathbf{x})$ , except when the specific bad outcome (e.g., critical infection, death, etc.) matters for our discussion.

Additionally, for some  $P \subseteq V$ , let  $f_P(\mathbf{x})$  and  $g_P(\mathbf{x})$  denote the infection prevalence and proportion of infections averted, respectively, in  $S$ . We are also given a *vaccine budget*  $D \in \mathbb{Z}_+$  that denotes the number of shipments of vaccines we are allowed to allocate in total.

**Definition of Disparity.** Recall that our focus is on outcome disparity; we aim to ensure, via appropriate vaccine allocation, that bad outcomes for the protected class are not much worse than these outcomes for the overall population. This motivates either considering a difference measure of disparity or a ratio measure of disparity. More specifically, one can define disparity as a difference in some metric of bad outcomes between the protected class and the general population. One can alternately, use a ratio measure instead of a difference measure. Motivated by prior research (Yeh et al. 2024; Ye et al. 2024; Yi and Marathe 2015) that a ratio definition of fairness provides scale-invariant properties that can help prevent misleading outcomes that can occur with difference definitions of fairness, we use a ratio definition in this paper. More precisely, we define disparity as the ratio of infection prevalence in the protected population to infection

prevalence in the remainder of the population.

$$d_P(\mathbf{x}) = \max\left(1, \frac{f_P(\mathbf{x})}{f_{V \setminus P}(\mathbf{x})}\right) \quad (1)$$

For example, a disparity value of 2.0 indicates that the protected population has twice the infection prevalence of the rest of the population, while a value of 1.5 represents a 50% higher infection rate. The  $\max(1, \cdot)$  elements ensures that when the protected population has lower or equal infection prevalence compared to the rest of the populated, the disparity is capped at 1 to prevent over-correction. Recall that earlier, multiple variants of the function  $f$  have been defined, i.e.,  $f^I(\mathbf{x})$ ,  $f^S(\mathbf{x})$ ,  $f^C(\mathbf{x})$ , and  $f^D(\mathbf{x})$  corresponding to mild infections, severe infections, critical infections, and death. Each of these induces a corresponding notion of disparity, which we denote respectively as  $d_P^I(\mathbf{x})$ ,  $d_P^S(\mathbf{x})$ , and  $d_P^D(\mathbf{x})$ .

We are now ready to present the specific formulations of the *Equitable Vaccine Allocation (EVA)* problem.

**Equitable Vaccine Allocation Formulation 1 (EVA-1):** In our first formulation (see the box EVA-1 below), we maximize the proportion of infections averted minus a weighted disparity, i.e.,  $g_V(\mathbf{x}) - \alpha \cdot d_P(\mathbf{x})$ , subject to a vaccine budget. For the remainder of the paper, we use  $b_\alpha(\mathbf{x}) := g_V(\mathbf{x}) - \alpha \cdot d_P(\mathbf{x})$  for the disparity-aware “benefit” of the vaccine allocation vector  $\mathbf{x}$ . Clearly, this formulation favors allocations that save a lot of individuals *and* reduce disparity. When the weight  $\alpha = 0$ , this formulation ignores disparity. As  $\alpha$  increases, the formulation becomes more “disparity aware”.

This simple formulation is inspired by the classical and well-studied problem of maximizing a submodular “cover” function subject to a cardinality constraint (Nemhauser, Wolsey, and Fisher 1978). Submodularity refers to the “diminishing returns” property of set or lattice functions, which can lead to guarantees in the performance of greedy algorithms (Nemhauser, Wolsey, and Fisher 1978). However, the objective function  $b_\alpha(\mathbf{x})$  of EVA-1 is not submodular and this rich theory of submodular function optimization does not apply to EVA-1.

#### EVA-1

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{Z}_+^V} \quad & b_\alpha(\mathbf{x}) := g_V(\mathbf{x}) - \alpha \cdot d_P(\mathbf{x}) \\ \text{s.t.} \quad & \|\mathbf{x}\| \leq D \end{aligned}$$

**Equitable Vaccine Allocation Formulation 2 (EVA-2):** Formulation 2 (see the box EVA-2 below) is inspired by the *Submodular Cost Submodular Cover (SCSC)* problem (Iyer and Bilmes 2013; Wan et al. 2010; Crawford, Kuhnle, and Thai 2019), which minimizes a *cost* function while ensuring that a *cover* function has a high enough value.

For our formulation, the cost function is defined by a linear combination of the disparity (Equation (1)) and the size of the vaccine allocation, i.e.,  $\gamma \cdot d_P(\mathbf{x}) + (1 - \gamma) \cdot \|\mathbf{x}\|/|V|$ . For the remainder of the paper, we use  $c_\gamma(\mathbf{x})$  to denote this objective function. Notice that unlike in EVA-1, we don’t have a “hard” budget constraint. We treat the number of vaccines allocated as a type of cost, along with disparity. The

<sup>4</sup>Note that to be fully correct,  $f(\mathbf{x})$  should be denoted  $f(\mathbf{x} | \mathcal{A}, I_0)$  because the spread of infection is conditioned on disease model  $\mathcal{A}$  and the set  $I_0$  of initially infected individuals. Usually,  $\mathcal{A}$  and  $I_0$  will be understood from the context, so it is convenient to drop  $\mathcal{A}$  and  $I_0$  from the notation.

cover function is simply defined as  $g_V(\mathbf{x})$ , the proportion of infections averted, and this is required to be at least a minimum threshold value  $\tau$ . Thus the overall effectiveness of the vaccination is ensured by this coverage constraint. It is worth emphasizing that even though our problem formulation is inspired by the SCSC problem, neither our cost function nor our cover function are submodular and therefore the performance guarantees in (Iyer and Bilmes 2013; Wan et al. 2010; Crawford, Kuhnle, and Thai 2019) do not apply to EVA-2.

### EVA-2

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{Z}_+^V} c_\gamma(\mathbf{x}) &:= \gamma \cdot d_P(\mathbf{x}) + (1 - \gamma) \cdot \frac{\|\mathbf{x}\|}{|V|} \\ \text{s.t. } g_V(\mathbf{x}) &\geq \tau \end{aligned}$$

Having alternate formulations with different parameterizations and hard constraints provides flexibility to policy makers and they can use one or the other depending on the real-world constraints they are working with. This is advantageous to the modeler and the algorithm-designer as well who can pick an appropriate formulation depending on availability of data and tractability of algorithms.

## Algorithms

We consider three algorithms to solve EVA; two variants of the greedy algorithm and the simulated annealing metaheuristic, all of which we describe in this section.

UNITGREEDY is the first and most straightforward greedy algorithm we present. It assumes the input of a single objective function to be maximized subject to a budget, and in each iteration, allocates one shipment at a time to the subpopulation that provides the most marginal gain in the objective function. The algorithm continues in this way until the vaccine budget is exhausted. Informally, the algorithm starts with an “empty allocation” to evaluate where sending a single shipment of vaccines would do the most good (defined by  $b_\alpha(\mathbf{x})$  described above). This algorithm then chooses where to send that single shipment, and repeats this process of selecting the subpopulation that adds the most value one at a time until it has used all available vaccines. With each decision made, the algorithm uses disease model to measure the benefit. The running time for UNITGREEDY is  $O(|V| \cdot D \cdot \theta)$ , where  $\theta$  is the running time of a single disease model evaluation given a vaccine allocation (Keithley et al. 2024). We run the maximization step on line 3 in parallel, so in practical terms, the running time is  $O(\max(1, \lfloor |V|/T \rfloor) \cdot D \cdot \theta)$ , where  $T$  is the number of computing threads available.

The next greedy algorithm we consider is RATIOGREEDY, which finds an approximate solution to Formulation 2 by iteratively selecting the subpopulation to send a shipment of vaccines which maximizes the ratio of the marginal gain of infections averted to the cost. If the cost function was monotone (i.e., increases with each addition to the input), we could simply stop the algorithm when the requirement  $g_V(\mathbf{x}) \geq \tau$  is fulfilled. Since the cost function is not monotone, we must run the greedy algorithm until

---

### Algorithm 1: UNITGREEDY( $D$ )

---

```

1:  $\mathbf{x} \leftarrow \mathbf{0}$ 
2: while  $\|\mathbf{x}\|_1 < D$  do
3:    $k^* \leftarrow \operatorname{argmax}_{k \in V} b_\alpha(\mathbf{x} + \mathbf{e}_k) - b_\alpha(\mathbf{x})$ 
4:    $\mathbf{x} \leftarrow \mathbf{x} + \mathbf{e}_{k^*}$ 
5: end while
6:
7: return  $\mathbf{x}$ 

```

---



---

### Algorithm 2: RATIOGREEDY( $\tau$ )

---

```

1:  $\mathbf{x} \leftarrow \mathbf{0}$ 
2:  $\mathcal{F} \leftarrow \{\mathbf{x}\}$ 
3: while  $\|\mathbf{x}\|_1 < |V|$  do
4:    $k^* \leftarrow \operatorname{argmax}_{k \in V} \frac{g_V(\mathbf{x} + \mathbf{e}_k) - g_V(\mathbf{x})}{c(\mathbf{x} + \mathbf{e}_k)}$ 
5:    $\mathbf{x}_{k^*} \leftarrow \mathbf{x}_{k^*} + 1$ 
6:    $\mathcal{F} \leftarrow \mathcal{F} \cup \{\mathbf{x}\}$ 
7: end while
8:  $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{F}} \{c(\mathbf{x}) \mid g_V(\mathbf{x}) \geq \tau\}$ 
9:
10: return  $\mathbf{x}^*$ 

```

---

the entire population is vaccinated, then select the allocation with the lowest cost which fulfills the threshold constraint. Because of the necessity to check all allocation sizes for EVA-2, RATIOGREEDY has running time  $O(|V|^2 \cdot \theta)$ , with a parallelized running time of  $O(\max(1, \lfloor |V|/T \rfloor) \cdot |V| \cdot \theta)$ .

We also consider SIMULATEDANNEALING, which is a randomized algorithm designed to imitate metallurgical annealing (Kirkpatrick, Gelatt, and Vecchi 1983) and explore the solution spaces of problems that have many local minima (due to space being limited, the pseudocode is presented in the full paper (Keithley, Bonner, and Pemmaraju 2025)). It begins with a random initial solution (randomly allocate a vaccine budget to subpopulations) and iteratively considers similar solutions generated by random perturbations, which are performed by moving a small number of vaccine shipments from one randomly selected subpopulation to another. If the algorithm finds a solution *better* than what it has so far, it accepts that new solution. In order to escape local minima, SIMULATEDANNEALING can also accept solutions that are *worse* (if no better solution is found) with a probability that gets smaller as the algorithm progresses. This probability is a function of the “temperature”, which is lowered with each iteration, so it is less likely to accept worse solutions as the algorithm progresses.

## Results

In this section, we experimentally evaluate our model, problem formulations, and algorithms for disparity-aware vaccine allocation in Johnson County, which is located in Eastern Iowa. Johnson County (population 152,854 according to the 2020 US Census<sup>5</sup>; 24 census tracts; 72 census block

<sup>5</sup>The contact network has  $\sim 7K$  fewer individuals represented because its construction uses the 2019 American Community Sur-

groups) includes Iowa City, which is the home of the University of Iowa. Because of the university, the population skews young, but a substantial fraction of the county is rural. Almost 11% of the population lives in census tracts with SVI above 0.8 and despite a comprehensive University-based medical center at the University of Iowa, almost 16% of the population lives in census block groups with TAI above 0.8. See Table 1 for more details. We define three protected classes for our experiments:

- **SVI-based:** Individuals residing in census tracts with Social Vulnerability Index (SVI) above 0.8.
- **TAI-based:** Individuals residing in census block groups with Treatment Accessibility Index (TAI) above 0.8. The TAI definition presented here uses a hospital catchment radius of 60 minutes.
- **Essential workers:** Individuals employed in retail, transportation, education, healthcare, or food service.

The correlation between high SVI and poor health outcomes in general (Higginbotham et al. 2025) and higher prevalence of COVID-19 (Biggs et al. 2021; Karaye and Horney 2020) motivates the definition of the SVI-based protected class. The definition of a TAI-based protected class is an attempt to explicitly identify groups of individuals with relatively poor healthcare access, taking travel time and capacity of healthcare facilities into account. The definition of the Essential workers protected class is motivated by the poor health outcomes of essential workers during COVID-19 (Gaitens et al. 2021). We first present what we refer to as our *basic experiment*: using UNITGREEDY to solve EVA-1 with an SVI-based protected class definition. We use  $\alpha = 0, 0.1, 0.3$  and  $0.5$  for the basic experiment and all experiments involving EVA-1. We call algorithms solving EVA-1 with  $\alpha = 0$  *disparity-oblivious*. Algorithms solving EVA-1 with  $\alpha > 0$  are called *disparity-aware*. We then present 5 additional experiments that extend the basic experiment as follows: **i)** extending the protected class definition to TAI-based and essential workers, **ii)** evaluating the importance of incorporating disparity into the disease model, **iii)** comparing the results of UNITGREEDY for EVA-1 to that of RATIOGREEDY for EVA-2, **iv)** evaluating the performance of SIMULATEDANNEALING compared to UNITGREEDY for EVA-1, and **v)** assessing the impact of adding a subregion-level policy which prioritizes older individuals and essential workers at vaccination sites. Results for a wider variety of settings may be found in the full paper (Keithley, Bonner, and Pemmaraju 2025).

### Improved Outcomes for SVI-based Protected Class

Figure 4 shows the results of the basic experiment, with subfigure A showing the infection prevalence  $f^I(\mathbf{x})$  and subfigure B showing the disparity  $d_P^I(\mathbf{x})$ . These experiments show that, over all vaccine allocation budgets considered (0% to 50% of the population) the infection prevalence of the disparity-aware algorithms is at most 3% higher than the infection prevalence of the disparity-oblivious algorithm. On the other hand, the disparity-aware algorithms for higher  $\alpha$

vey, which uses different sampling methods than the Census.




Figure 4: Results from using UNITGREEDY on EVA-1, with an SVI-based protected class, for budgets of up to 50% of the population vaccinated. (A) shows the overall infection prevalence  $f^I(\mathbf{x})$  and (B) shows the disparity  $d_P^I(\mathbf{x})$ .

( $\alpha = 0.3, 0.5$ , shown in red with triangle and diamond markers) reduce disparity very quickly; disparity stays below 1 after 5% of the population has been vaccinated. Disparity falls for the disparity-oblivious algorithm also, but more slowly; about 15% of the population needs to be vaccinated before disparity falls below 1 for this algorithm. At 5% vaccination, all disparity-aware algorithms have achieved disparity 1, whereas the disparity-oblivious algorithm has a disparity of 1.4. In general, higher values of  $\alpha$  achieve lower values of disparity than lower values of  $\alpha$ , without raising the infection prevalence for the overall population substantially.

### Improved Outcomes for TAI-based Protected Class and Essential Workers

Figure 5 and Figure 6 show outcomes using the same setup as the previous subsection, but using TAI-based and essential worker based protected classes, respectively. Here, prevalence and disparity ( $f^S(\mathbf{x})$  and  $d^S(\mathbf{x})$ ) for the TAI-based protected class are in terms of serious infections, while prevalence and disparity ( $f^I(\mathbf{x})$  and  $d^I(\mathbf{x})$ ) for the essential worker protected class are in terms of overall infections.

Figure 5 shows that using UNITGREEDY to solve EVA-1 reduces disparity quickly for the disparity-aware algorithm for nonzero values of  $\alpha$ , while disparity decreases slightly more slowly for the disparity-oblivious algorithm but does not consistently stay below 1. The disparity-aware algorithm with  $\alpha = 0.5$  results in at most a 1% higher level of severe infections as compared to the disparity-oblivious algorithm, but we note that the disparity-aware algorithm using  $\alpha = 0.1$  and  $\alpha = 0.3$  achieves a severe infection prevalence almost as low as that of the disparity-oblivious algorithm, while keeping disparity below 1.

As shown in Figure 6, the disparity-aware algorithm with  $\alpha = 0.5$  achieves a *lower* infection prevalence than the other algorithms, while also achieving the lowest disparity. Every algorithm is eventually able to achieve equity, but the disparity-aware algorithm with  $\alpha = 0.5$  and  $\alpha = 0.3$  achieve it with about 10% fewer vaccines than the disparity-oblivious or disparity-aware algorithm with  $\alpha = 0.1$ . This indicates that in certain cases, the disparity-aware algorithm can reach a vaccine allocation that is both more equitable and averts more infections than the disparity-oblivious algorithm.




Figure 5: Results from using UNITGREEDY on EVA-1, with a TAI-based protected class, for budgets of up to 50% of the population vaccinated. (A) shows the overall infection prevalence  $f^S(x)$  and (B) shows the disparity  $d_P^S(x)$ .




Figure 6: Results from using UNITGREEDY on EVA-1, with an essential worker based protected class, for budgets of up to 50% of the population vaccinated. (A) shows the overall infection prevalence  $f^I(x)$  and (B) shows the disparity  $d_P^I(x)$ .

### Disparity-aware Modeling is Crucial

The goal of this subsection is to show the importance of incorporating disparity into the disease model used to make vaccine allocation decisions. Here, we run disparity-oblivious and disparity-aware algorithms using UNITGREEDY conditioned on a model that ignores all factors of disparity (which we shall refer to as the *disparity-unadjusted* disease model) discussed in the Modeling section. We then evaluate the resulting vaccine allocations using the model that does account for disparity, which we call the *disparity-adjusted* disease model. In general, we call this comparison *cross-evaluation*, and we use it to demonstrate outcomes that result from making vaccination allocation decisions without the knowledge of how those decisions affect disparity.

As shown in Figure 7, UNITGREEDY achieves an infection prevalence similar to that of the basic experiment with both disparity-oblivious and disparity-aware algorithms, but all perform far worse with respect to disparity. The protected class experiences infection rates up to about 3 times higher than the rest of the population, and the disparity increases with as the vaccine budget grows.

Notably, both disparity-oblivious and disparity-aware algorithms exhibit similar infection prevalence rates as in previous subsections across all vaccine budgets, demonstrating that a dual goal to reduce infections and minimize disparity does not compromise effectiveness. This highlights the importance of incorporating disparity into our disease model, as we can achieve substantially lower disparity while maintaining a similar level of infections averted. These results also show how allocation strategies that appear effective




Figure 7: Results from using UNITGREEDY (cross evaluated) on EVA-1, with an SVI-based protected class, for budgets of up to 50% of the population vaccinated. (A) shows the overall infection prevalence  $f^I(x)$  and (B) shows the disparity  $d_P^I(x)$ .




Figure 8: Results from using UNITGREEDY on EVA-1, with an SVI-based protected class, for budgets of up to 50% of the population vaccinated and compared with results from the same setup using RATIOGREEDY. (A) shows the overall infection prevalence  $f^I(x)$  and (B) shows the disparity  $d_P^I(x)$ .

from a high-level perspective may inadvertently aggravate significant health disparities when underlying disparities are not accounted for in the modeling process.

### UNITGREEDY Trades Efficiency for Disparity Compared to RATIOGREEDY

So far, we have focused solely on evaluating UNITGREEDY applied to EVA-1. Here, we compare those results to that of RATIOGREEDY applied to EVA-2. In particular, for each fixed value of  $\alpha$  (we show results for  $\alpha = 0.5$  here), we find a corresponding value of  $\gamma$  in EVA-2 that results in a similar infection prevalence for each vaccine budget. Note that we do not consider a disparity-oblivious algorithm here - we are only comparing the disparity-aware algorithm with  $\alpha = 0.5$  to results from RATIOGREEDY applied to EVA-2.

Figure 8 demonstrates that UNITGREEDY and RATIOGREEDY achieve close performance in terms of infection prevalence. The highest difference is at a 30% vaccine budget, where UNITGREEDY results in about 5% higher infection prevalence than that of RATIOGREEDY. This comes with the advantage of UNITGREEDY consistently achieving a lower disparity than RATIOGREEDY, although the disparity for RATIOGREEDY is never above about 1.1 for vaccine budgets of over 10%. Overall, this result suggests that UNITGREEDY applied to EVA-1 is better suited for balancing efficiency and equity than RATIOGREEDY but may sacrifice efficiency in some cases.




Figure 9: Results from using UNITGREEDY on EVA-1, with an SVI-based protected class, for budgets of up to 50% of the population vaccinated and compared with results from the same setup using SIMULATEDANNEALING. (A) shows the overall infection prevalence  $f^I(\mathbf{x})$  and (B) shows the disparity  $d_P^I(\mathbf{x})$ .

### SIMULATEDANNEALING Shows Similar Trends

In this subsection, we evaluate the performance of SIMULATEDANNEALING compared to that of the UNITGREEDY, both disparity-aware with  $\alpha = 0.5$  applied to EVA-1. We use the SVI-based protected class and run the algorithms for vaccine budgets between 0% and 50%. Figure 9 shows UNITGREEDY and SIMULATEDANNEALING result in close infection prevalence for vaccine budgets of over 30%, but UNITGREEDY achieves an infection prevalence of about 7% lower than that of SIMULATEDANNEALING for a 20% vaccine budget. This result does not come in the form of a trade-off between effectiveness and equity, since UNITGREEDY lowers disparity below 1 with fewer vaccines than SIMULATEDANNEALING, although both methods consistently have a low disparity with a vaccine budget of greater than 10%. These results reinforce the general effectiveness of using a disparity-aware UNITGREEDY algorithm for equitable vaccine allocation.

### A Local Allocation Policy Targeting Essential Workers and Older Individuals

So far, we have only considered settings in which the vaccines allocated to a subregion are distributed to the members of its population uniformly at random (i.e., every member of the subregion’s population has an equal chance of receiving one of the vaccines distributed to that subregion). In general, we refer to the method that determines which individuals receive one of a subregion’s assigned vaccines as a *local allocation policy*. In this experiment, we explore a non-uniform local allocation policy, where vaccines are distributed with preference to individuals older than 65 and essential workers, where those groups are 5 times more likely to receive a vaccine from a shipment allocated to their home subregion. We consider this local allocation policy where vaccines are allocated to each subregion by UNITGREEDY using the essential worker based protected class.

Figure 10 shows that both disparity-aware and disparity-oblivious algorithms using UnitGreedy achieve similar results regardless of  $\alpha$ , with all values of  $\alpha$  producing nearly identical infection prevalence and disparity outcomes across all vaccine budgets. This suggests that the local allocation policy used here protects the essential worker protected class




Figure 10: Results from using UNITGREEDY (cross evaluated) on EVA-1, with an essential worker based protected class, for budgets of up to 50% of the population vaccinated. (A) shows the overall infection prevalence  $f^I(\mathbf{x})$  and (B) shows the disparity  $d_P^I(\mathbf{x})$ .

effectively, making the regional disparity consideration governed by  $\alpha$  less influential on the results, indicating that targeted local allocation can achieve equity goals independently of the regional optimization approach.

## Discussion

Our primary focus of this paper is to translate existing ethical vaccine allocation frameworks (NASEM 2020; Toner et al. 2020; World Health Organization 2020) into a computational framework using modern mathematical modeling and optimization methods. Through this translation, we enable public health decision-makers to leverage computational resources for working towards equity goals within established policy frameworks. Our framework prioritizes transparency, flexibility, and extendability. Our results demonstrate how explicitly incorporating disparity into disease modeling and allocation methods can reduce disparity without worsening health outcomes in the general population. While we focus on a specific setting (allocating vaccines along geographical lines to Johnson County, Iowa) in our experiments, our framework is general and can be adapted or extended along multiple dimensions.

Several limitations merit consideration. We make several simplifying assumptions to keep our paper focused. We assume perfect vaccine acceptance and efficacy; the former is reasonable when demand exceeds supply, but not in general. Our framework could readily incorporate empirical data on vaccine hesitancy, demographic-specific efficacy rates, or logistical constraints that limit vaccine delivery to certain subpopulations. Our current approach focuses on a specific definition of health disparity, but alternative definitions (e.g., Gini index, Slope Inequality Index) merit exploration. Future work could also include the protection of multiple groups at once. Our focus on preemptive allocation decisions, while practical for vaccine allocation at the early stages of a pandemic, could be extended to non-preemptive settings, where allocations need to be made over time. Future extensions could incorporate real-time manufacturing constraints, distribution bottlenecks, or adaptive allocation strategies that respond to changing epidemiological conditions. Finally, future work could also explore other, more fine-grained models for vaccine allocation.

## Ethical Impact

Our research relies on population and epidemiological data, raising several ethical considerations. First, individual-level data must be properly anonymized to protect privacy, following established practices such as those used by the American Community Survey. While aggregate distributions should reflect real-world patterns, no personally identifiable information should be included in allocation algorithms.

Inaccurate or biased population data can lead to inequitable vaccine allocation decisions that may systematically disadvantage certain communities. Similarly, errors in disease surveillance data or epidemiological modeling can misrepresent disease spread patterns, potentially directing vaccines away from areas of greatest need. We emphasize the critical importance of using high-quality, representative data sources and regularly validating model assumptions against observed outcomes.

## Positionality Statement

The authors' expertise lies in the area of computer science and optimization applied to problems in infectious disease modeling and intervention. While well-versed in techniques for studying disease *computationally*, the authors acknowledge gaps in their expertise with regards to real-world implementation. To date, the authors have sought feedback from experts in fields such as infectious disease, biostatistics, and epidemiology, but feel additional feedback from and constructive discussion with practitioners in public health would be very helpful.

## Broader Impact

We discuss potential adverse impacts of the research outlined in this paper in two directions: **i)** unintentional adverse impacts and **ii)** intentional adverse impacts.

- i) Our framework is intended for use in regional public health departments alongside expert decision-makers who can use it as a decision-support tool. Rigid adherence to recommendations made by any framework (including ours) may reduce responsiveness to local conditions, community needs, or emerging equity concerns during implementation. Since modeling a complex system such as disease spread has inherent uncertainty and cannot fully capture all details, our framework should serve as a starting point for informed decisions rather than a substitute for expert judgment and community engagement.
- ii) While intended as a tool to decrease disparity in disease outcomes, our framework may be used to justify discriminatory practices under the guise of scientific objectivity. To prevent intentional misuse, we emphasize that all model parameters, assumptions, and trade-offs should be made transparent to stakeholders and communities. Our framework should never be used as a “black box” to avoid accountability for allocation decisions or to circumvent community engagement in policy-making.

## Acknowledgments

The authors acknowledge feedback from members of the Computational Epidemiology research group at the University of Iowa and the CDC MInD-Healthcare group. Funding for this research was provided in part by the CDC MInD Healthcare group under cooperative agreement U01CK000594 and associated Covid19 supplemental funding for authors JK and SVP. Funding for author MB was provided by NSF REU site grant 2050526 titled “Computing for Health and Well-Being”.

## References

- Biggs, E. N.; Maloney, P. M.; Rung, A. L.; Peters, E. S.; and Robinson, W. T. 2021. The Relationship Between Social Vulnerability and COVID-19 Incidence Among Louisiana Census Tracts. *Frontiers in Public Health*, 8: 617976.
- Bissett, K. R.; Cadena, J.; Khan, M.; and Kuhlman, C. J. 2021. Agent-Based Computational Epidemiological Modeling. *Journal of the Indian Institute of Science*, 101(3): 303–327.
- CDC. 2021a. Categories of Essential Workers: COVID-19 Vaccination.
- CDC. 2021b. Social Vulnerability Index 2020 Database.
- CDPH. 2020. California COVID-19 Vaccination Plan.
- Cohen, J. A.; Mistry, D.; Kerr, C. C.; and Klein, D. J. 2020. Schools are not islands: Balancing COVID-19 risk and educational benefits using structural and temporal countermeasures. *medRxiv*.
- Crawford, V. G.; Kuhnle, A.; and Thai, M. T. 2019. Submodular Cost Submodular Cover with an Approximate Oracle. arXiv:1908.00653.
- Emanuel, E. J.; and Wertheimer, A. 2006. Who Should Get Influenza Vaccine When Not All Can? *Science*, 312(5775): 854–855.
- Gaffney, A.; Himmelstein, D. U.; McCormick, D.; and Woolhandler, S. 2023. COVID-19 Risk by Workers' Occupation and Industry in the United States, 2020–2021. *American Journal of Public Health*, 113(6): 647–656.
- Gaitens, J.; Condon, M.; Fernandes, E.; and McDiarmid, M. 2021. COVID-19 and essential workers: a narrative review of health outcomes and moral injury. *International Journal of Environmental Research and Public Health*, 18(4): 1446.
- Germann, T. C.; Kadau, K.; Longini, I. M.; and Macken, C. A. 2006. Mitigation strategies for pandemic influenza in the United States. *Proceedings of the National Academy of Sciences*, 103(15): 5935–5940.
- Gu, T.; Mack, J. A.; Salvatore, M.; Prabhu Sankar, S.; Valley, T. S.; Singh, K.; Nallamotheu, B. K.; Kheterpal, S.; Lisabeth, L.; Fritsche, L. G.; and Mukherjee, B. 2020. Characteristics Associated With Racial/Ethnic Disparities in COVID-19 Outcomes in an Academic Health Care System. *JAMA Network Open*, 3(10): e2025197.
- Higginbotham, J. K.; Segovia, L. M.; Rohm, K. L.; Anderson, C. M.; and Breitenstein, S. M. 2025. Social Vulnerability Index and Health Outcomes in the United States: A

- Systematic Review. *Family & Community Health*, 48(2): 81–96.
- HRSA. 2025. MUA Find: Find data on Medically Underserved Area and Medically Underserved Population designations.
- Iyer, R. K.; and Bilmes, J. A. 2013. Submodular optimization with submodular cover and submodular knapsack constraints. *Advances in neural information processing systems*, 26.
- Karaye, I. M.; and Horney, J. A. 2020. The Impact of Social Vulnerability on COVID-19 in the U.S.: An Analysis of Spatially Varying Relationships. *American Journal of Preventive Medicine*, 59(3): 317–325.
- Keithley, J.; Bonner, M.; and Pemmaraju, S. 2025. Models and Algorithms for Balancing Efficiency and Equity in Vaccine Allocation. *medRxiv*. In preparation.
- Keithley, J.; Choudhuri, A.; Adhikari, B.; and Pemmaraju, S. V. 2024. Analyzing Greedy Vaccine Allocation Algorithms for Metapopulation Disease Models. *medRxiv*.
- Kermack, W.; and McKendrick, A. 1927. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London*, 15(772): 700–721.
- Kerr, C. C.; Stuart, R. M.; Mistry, D.; Abey Suriya, R. G.; Hart, G. R.; Rosenfeld, K.; Nunez, R. C.; Cohen, J. A.; Selvaraj, P.; Hagedorn, B.; George, L.; Jastrzebski, M.; Izzo, A. S.; Fowler, G.; Palmer, A.; Delport, D.; Scott, N.; Kelly, S. L.; Bennette, C. S.; and Klein, D. J. 2021. Covasim: An agent-based model of COVID-19 dynamics and interventions. *PLOS Computational Biology*, 17(7): e1009149.
- Khan, N.; Mu, K.; Moharrami, M.; and Subramanian, V. 2023. Backward and Forward Inference in Interacting Independent-Cascade Processes: A Scalable and Convergent Message-Passing Approach. arXiv:2310.19138.
- Kirkpatrick, S.; Gelatt, C. D.; and Vecchi, M. P. 1983. Optimization by Simulated Annealing. *Science*, 220(4598): 671–680.
- Kiss, I. Z.; Miller, J. C.; and Simon, P. L. 2017. *Mathematics of Epidemics on Networks: From Exact to Approximate Models*, volume 46 of *Interdisciplinary Applied Mathematics*. Springer International Publishing.
- Liu, K.; and Lou, Y. 2022. Optimizing COVID-19 Vaccination Programs during Vaccine Shortages. *Infectious Disease Modelling*, 7(1): 286–98.
- Luxen, D.; and Vetter, C. 2011. Real-time routing with OpenStreetMap data. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, GIS '11, 513–516. New York, NY, USA: ACM. ISBN 978-1-4503-1031-4.
- Medlock, J.; and Galvani, A. P. 2009. Optimizing Influenza Vaccine Distribution. *Science*, 325(5948): 1705–1708.
- MichiganDHHS. 2020. Michigan COVID-19 Vaccination Plan.
- Mossong, J.; Hens, N.; Jit, M.; Beutels, P.; Auranen, K.; Mikolajczyk, R.; Massari, M.; Salmaso, S.; Tomba, G. S.; Wallinga, J.; Heijne, J.; Sadkowska-Todys, M.; Rosinska, M.; and Edmunds, W. J. 2008. Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases. *PLOS Medicine*, 5(3): e74.
- NASEM. 2020. *Framework for Equitable Allocation of COVID-19 Vaccines*. Washington, DC: The National Academies Press.
- Nemhauser, G.; Wolsey, L.; and Fisher, M. 1978. An Analysis of Approximations for Maximizing Submodular Set Functions—I. *Mathematical Programming*, 14(1): 265–94.
- NMDOH. 2020. New Mexico Vaccine Distribution Plan.
- OSHA. 2025. COVID-19 - Control and Prevention.
- OSRM. 2010. Open Source Routing Machine (OSRM) Route API.
- Panovska-Griffiths, J.; Kerr, C. C.; Stuart, R. M.; Mistry, D.; Klein, D. J.; and Viner, R. M. 2020. Determining the optimal strategy for reopening schools, the impact of test and trace interventions, and the risk of occurrence of a second COVID-19 epidemic wave in the UK: a modelling study. *The Lancet Child & Adolescent Health*, 4(11): 817–827.
- Pijls, B. G.; Jolani, S.; Atherley, A.; Derckx, R. T.; Dijkstra, J. I. R.; Franssen, G. H. L.; and Hendriks, S. 2021. Demographic Risk Factors for COVID-19 Infection, Severity, ICU Admission and Death: A Meta-Analysis of 59 Studies. *BMJ Open*, 11(1): e044640.
- Prem, K.; Cook, A. R.; and Jit, M. 2017. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLOS Computational Biology*, 13(9): e1005697.
- Prem, K.; van Zandvoort, K.; Klepac, P.; Eggo, R. M.; Davies, N. G.; Cook, A. R.; and Jit, M. 2021. Projecting contact matrices in 177 geographical regions: An update and comparison with empirical data for the COVID-19 era. *PLOS Computational Biology*, 17(7): e1009098.
- ProximityOne. 2024. Census Block Groups and Block Group Codes.
- Scott, N.; Palmer, A.; Delport, D.; Abey Suriya, R.; Stuart, R.; and Kerr, C. C. 2020. Modelling the impact of reducing control measures on the COVID-19 pandemic in a low transmission setting. *The Medical Journal of Australia*, 1.
- Srivastava, V.; and Priyadarshini, S. 2021. Vaccine Shortage Dents India's Coronavirus Adult Immunisation Drive. *Nature India*.
- TexasDSHS. 2020. Texas COVID-19 Vaccination Plan.
- Tipirneni, R.; Karmakar, M.; O'Malley, M.; Prescott, H. C.; and Chopra, V. 2022. Contribution of Individual- and Neighborhood-Level Social, Demographic, and Health Factors to COVID-19 Hospitalization Outcomes. *Annals of Internal Medicine*, M21–2615.
- Toner, E.; Barnhill, A.; Krubiner, C.; Bernstein, J.; Privordum, L.; Watson, M.; Martin, E.; Rutkow, L.; Cicero, A.; Inglesby, T.; et al. 2020. Interim Framework for COVID-19 Vaccine Allocation and Distribution in the United States. Technical report, Johns Hopkins Center for Health Security, Baltimore, MD.

- Tulchinsky, A. Y.; Haghpanah, F.; Hamilton, A.; Kipshidze, N.; and Klein, E. Y. 2024. Generating geographically and economically realistic large-scale synthetic contact networks: A general method using publicly available data. arXiv:2406.14698.
- USCensus. 2021. 2020 Census Demographic Data Map Viewer.
- USCensus. 2025a. North American Industry Classification System (NAICS).
- USCensus. 2025b. Public Use Microdata Sample (PUMS).
- Wan, P.-J.; Du, D.-Z.; Pardalos, P.; and Wu, W. 2010. Greedy Approximations for Minimum Submodular Cover with Submodular Cost. *Computational Optimization and Applications*, 45(2): 463–474.
- World Health Organization. 2020. WHO Concept for Fair Access and Equitable Allocation of COVID-19 Health Products. Final working version, World Health Organization.
- Ye, P.; Ye, Z.; Xia, J.; Zhong, L.; Zhang, M.; Lv, L.; Tu, W.; Yue, Y.; and Li, Q. 2024. National-scale 1-km maps of hospital travel time and hospital accessibility in China. *Scientific Data*, 11(1): 1130.
- Yeh, M.-H.; Metevier, B.; Hoag, A.; and Thomas, P. 2024. Analyzing the Relationship Between Difference and Ratio-Based Fairness Metrics. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 518–528. ACM.
- Yi, M.; and Marathe, A. 2015. Fairness versus Efficiency of Vaccine Allocation Strategies. *Value in Health: The Journal of the International Society for Pharmacoeconomics and Outcomes Research*, 18(2): 278–283.
- Zhang, Y.; Adiga, A.; Saha, S.; Vullikanti, A.; and Prakash, B. 2016. Near-Optimal Algorithms for Controlling Propagation at Group Scale on Networks. *IEEE Transactions on Knowledge and Data Engineering*, 28(12): 3339–52.