Modeling Access Differences to Reduce Disparity in Resource Allocation

Kenya S. Andrews
kandre32@uic.edu
University of Illinois at Chicago
Chicago, IL, USA

Mesrob I. Ohannessian
mesrob@uic.edu
University of Illinois at Chicago
Chicago, IL, USA

Tanya Berger-Wolf
berger-wolf.1@osu.edu
Ohio State University
Columbus, OH, USA

ABSTRACT

Motivated by COVID-19 vaccine allocation, where vulnerable subpopulations are simultaneously more impacted in terms of health and more disadvantaged in terms of access to the vaccine, we formalize and study the problem of resource allocation when there are inherent access differences that correlate with advantage and disadvantage. We identify reducing resource disparity as a key goal in this context and show its role as a proxy to more nuanced downstream impacts. We develop a concrete access model that helps quantify how a given allocation translates to resource flow for the advantaged vs. the disadvantaged, based on the access gap between them. We then provide a methodology for access-aware allocation. Intuitively, the resulting allocation leverages more vaccines in locations with higher vulnerable populations to mitigate the access gap and reduce overall disparity. Surprisingly, knowledge of the access gap is often not needed to perform access-aware allocation. To support this formalism, we provide empirical evidence for our access model and show that access-aware allocation can significantly reduce resource disparity and thus improve downstream outcomes. We demonstrate this at various scales, including at county, state, national, and global levels.

ACM Reference Format:
Kenya S. Andrews, Mesrob I. Ohannessian, and Tanya Berger-Wolf. 2022. Modeling Access Differences to Reduce Disparity in Resource Allocation. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '22), October 6–9, 2022, Arlington, VA, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3551624.3555302

1 INTRODUCTION

All common resources must be allocated in some way to members of a society. During such an allocation, the goal should be to optimize the benefits of the resources for that society. To be transformational, the way we allocate needs to be not only superficially fair but also deeply just and willfully acting to mitigate unfairness. We learned a great deal about what happens when we are not intentional in our allocation efforts through the COVID-19 virus health crisis. Yet, amidst the rise in monkeypox cases, we find ourselves in familiar territory with shortages and allocation issues for the monkeypox virus vaccines [14, 21]. For most of the COVID-19 pandemic, government officials and healthcare workers have been rapidly seeking ways to minimize loss caused by the COVID-19 virus. In this search, they found the virus affects members of the vulnerable population more severely; namely people with underlying conditions, members of racial minority groups, those who lack personal transportation, those of lower economic status, and older people [13]. Vulnerable people experience worse short and long term health side effects and higher mortality rates [6]. Additionally, the impact on cities and states when members of the vulnerable populations get sick from COVID-19 leads to higher hospitalization rates, lack of healthcare staff availability, and less hospital space to care for other conditions. Though these realities were recognized, many of the plans that were developed to minimize loss, including vaccination, did not fully take this into consideration. Instead, when vaccines began to be distributed, they were distributed according to age (oldest) or career (most essential workers) [18]. Additionally, points of distribution for vaccines were receiving proportional amounts of vaccines [8]. It is seemingly ideal to have such an allocation, but there are several examples where this logic breaks. For example, age is not the only good measure because we see Black Americans have a shorter life expectancy than their White counterparts due to health inequities [20]. Thus, only focusing on age exacerbates these issues for Black Americans.

Motivated by COVID-19 vaccine allocation, we formalize the problem of fair resource allocation when there are inherent access differences that correlate with advantage and disadvantage. We make and support three key premises:

(A) The population is stratified into disadvantaged and advantaged subpopulations.
(B) The disadvantaged subpopulation is more vulnerable to the lack of resources.
(C) The disadvantaged subpopulation has obstacles in acquiring allocated resources.

In Sec. 3, we describe the formal problem by closely following premises (A) and (B). Our first contribution there is to highlight reducing **resource disparity** between advantaged and disadvantaged subpopulations as a key goal of resource allocation. Then, in Sec. 4, our second contribution is a concrete **access model**, which we develop through premise (C). This allows us to determine how resources would flow to the disadvantaged vs. advantaged subpopulations, upon performing a given allocation. In Sec. 5, we show that resource disparity can be effectively thought of as a proxy to more complex impact disparities downstream. In Sec. 6, we combine the formal problem with the access model to obtain our third contribution, an **access-aware allocation** algorithm that reduces resource disparity while maintaining approximate geographic proportionality. A key finding here is that access-aware allocation can often be performed without explicit knowledge of the access gap

EAAMO '22, October 6–9, 2022, Arlington, VA, USA

ACM ISBN 978-1-4503-9477-2/22/10.
https://doi.org/10.1145/3551624.3555302

between the subpopulations. We acknowledge that there are many contributors to access disparities other than disadvantage. That said, in Sec. 7.1 we give empirical evidence in support of our access model. Then, in Sec. 7.2, we revisit the COVID-19 vaccine allocation problem as a use case. Our fourth contribution is, we show **significant reduction** in resource disparity and thus downstream impact can be achieved with minimal deviation from proportionality. In Sec. 8, we conclude the paper with discussion points. We start with a survey of related work in Sec. 2.

2 RELATED WORK

The importance of considering equity and justice when allocating resources has not been lost on the scientific community, especially for the COVID-19 vaccine distribution problem, the primary case study here. The role of access differences and its interplay with vulnerability and disadvantage is often alluded to, but not directly formalized as we do.

For example, in [18], the authors note that those 65 and older account for 80% of deaths from COVID-19; however, they warn that "if this is done without explicit attention to promoting health equity, it will, once again, exacerbate major health disparities." They attribute this to the fact that: making a vaccination appointment requires substantial time, reliable internet access and technology, ability to travel to a vaccination site, and trust in the safety of the vaccine. Note that inability to do any of these and the existence of pre-exisiting health conditions correlates with a person being more vulnerable. Several solutions are offered in [18], including prioritizing hardest hit and lowest economic ZIP Codes, which is very much in line with the access-aware allocation solutions that come out of our explicit model. Other suggestions are: local vaccine education, decreasing transportation barriers, and simplifying registration procedures. Though we do not discuss these, we do acknowledge that they could aid in grassroots efforts of vaccine administration.

In [33], the authors develop a tool to identify people who will benefit most from receiving the vaccine at a specific time period. They argue that those at most risk of contracting should be prioritized. However, access differences are not made explicit, and this focus does not necessarily address minimizing the harshest effects of contracting COVID-19.

In [12], data from electronic hospital records (EHRs) is used to train a model which can determine the likelihood of a person to die from contracting COVID-19. While we do not depend on health records to determine whether a person would die from COVID-19, we do use the CDC's full definition of what it means to be vulnerable. The findings of [12] (see also [17] for early variant findings) clearly show affect the level of hospitalization and death risk. It is also worth noting that health records for racial minorities, lower-income people, and those who are in more rural areas tend to not be complete profiles due to access barriers [26].

Many data-driven methods to guide vaccine allocation have also been proposed. In [9], the authors cluster states with similarly reported effective reproduction numbers[1] to look for an optimal distribution. Their focus does not include equity and is more on preventing spread, which we do not address. By concentrating on resource equity between vulnerable and non-vulnerable subpopulations, we focus on who would benefit most from receiving enough vaccines.

In [7], allocation under limited resources is addressed but without considerations of equity, disparity, or access. There, the importance of dynamic allocation is emphasized, which we do not address. In [29], the authors consider overall COVID-19 impacts (death, job loss, and changes in the economy). They build a linear optimization method to predict testing allocation, based on predicting the influx of cases, over different time periods to maximize detection of the virus amongst different population groups (symptomatic, asymptomatic, and contact tracing). They incorporate a basic geographic notion of equity, by allocating a minimum number of resources at each location, but do not model access disparity. In a different relevant line of work, resources are allocated to increase direct (spatial) access. For example, in [34], the placement of hospital beds is optimized using a model of spatial accessibility. This does not involve an explicit relationship between vulnerability and resource acquisition but is relevant nonetheless in that it centers on issues of access by the disadvantaged and findings include resulting improvements in equity also.

In the theoretical CS and OR communities, models of equity and allocation have also been proposed. Recently, [10] studied resource allocation across groups with varying demand distributions and showed tradeoffs between utilization and fairness across groups. The varying demand distributions could be interpreted as access differences, but it is important to note that upon instantiation of the demand, there is no inherent difference in the ease with which the resource is acquired. More importantly, in that model, individual allocations can be made for each group, which may not ideally describe shared public resource allocation, save for assistance models that occur through an application process. Though we focus on access differences in physical resource acquisition, we acknowledge the importance of differences in access to information and their impact, e.g., as is done in [16] for standardized testing. Of course, there are many dimensions beyond access, such as dynamics and correlations, that further complicate resource acquisition. For example, this was explored in depth by [22] for rationing medical supplies in the COVID-19 backdrop.

A key theme in the current paper is that the disadvantaged lack ease of access to resources and, simultaneously, are disproportionally affected by the burden of lack of resource [4, 27]. Sometimes they can be targeted directly (e.g., mobile clinics, government-supplied PPE, etc). Though these are helpful solutions, these efforts are expensive and resources often run out before having had sufficient impact [30]. Policies that place more of these resources where they are needed, such as the access-aware allocation methods proposed here, can boost resource acquisition indirectly. Reaching vulnerable populations and mitigating the obstacles they face are critical in making optimization and data-driven methods achieve their true clinical potential [23, 35]. We believe explicitly modeling access differences is an essential tool for achieving this potential.

[1] an epidemiological property that describes the expected number of new cases due directly to one case

3 PROBLEM DESCRIPTION

Allocation Task and Notation. We are tasked with allocating an amount N_Σ of resources, assumed to be arbitrarily divisible, across k locations indexed by $j \in [k] := \{1, \cdots, k\}$. The decision variables are the amounts to be allocated to each location j, denoted by N_j, or equivalently the fraction of the whole that these represent, denoted by $n_j := N_j/N_\Sigma$. It is desirable to entirely allocate the resources, so that $\sum_j N_j = N_\Sigma$ and $\sum_j n_j = 1$. We write $\mathbf{n} := (n_j)_{j\in[k]}$. The implicit target for this allocation are the population sizes P_j at each location. Using parallel notation to N_Σ, let $P_\Sigma = \sum_j P_j$ be the total population size and $p_j := P_j/P_\Sigma$ be the fraction of the population at location j, so that $\sum_j p_j = 1$. We write $\mathbf{p} := (p_j)_{j\in[k]}$. Resources are assumed to be limited, $N_\Sigma < P_\Sigma$, and we parametrize this explicitly as $\alpha := N_\Sigma/P_\Sigma \in (0, 1)$.

Disadvantage vs. Advantage. A central premise of this work, Premise (A), is that the population is stratified by disadvantage and that this stratification varies across locations. The *disadvantaged* should be regarded as particularly vulnerable in the absence of a resource. We adopt the simplest model that encapsulates this, by using a parameter $\beta_j \in [0, 1]$, which describes the fraction at location j who is a particularly disadvantaged subpopulation. While the remainder $1 - \beta_j$ of the population may also have some vulnerability, we will describe them as *advantaged*. This stratification influences the allocation task through a difference in how allocated resources are acquired by each subpopulation. At any given location j, given N_j resources, we assume there exists an acquisition function ρ, that determines the fraction of N_j acquired by the disadvantaged, while the remainder $1 - \rho$ is acquired by the advantaged.

How can we characterize ρ? We can only do this through a clear model of access that describes the process by which allocated resources are acquired by each subpopulation and the difference between them. This is the motivation of Sec. 4, which elaborates our proposed access model theoretically — this is later supported empirically in Sec. 7.1. The resulting ρ is found to depend on the local quantities N_j, P_j, and β_j, as well as a global parameter η that describes the *access gap* between disadvantaged and advantaged subpopulations. In what follows, for conciseness, we refer to the value that ρ takes at location j as ρ_j. For reference, please refer to Table 1 for a glossary of the main notations.

Resource Rate Disparity. Upon an allocation $\mathbf{n}$, each subpopulation acquires part of the N_Σ resources. The disadvantaged acquire $\sum_j \rho_j N_j = N_\Sigma \sum_j \rho_j n_j$, while the advantaged acquire the rest $\sum_j (1 - \rho_j)N_j = N_\Sigma \sum_j (1 - \rho_j)n_j$. As a result, resource acquisition rate may vary across the two subpopulations, which we quantify through *rate disparity*:

$$\begin{aligned} \mathsf{RD} &= \frac{\sum_j (1-\rho_j)N_j}{\sum_j (1-\beta_j)P_j} - \frac{\sum_j \rho_j N_j}{\sum_j \beta_j P_j} \\ &= \alpha \frac{\sum_j (1-\rho_j)n_j}{\sum_j (1-\beta_j)p_j} - \alpha \frac{\sum_j \rho_j n_j}{\sum_j \beta_j p_j} \end{aligned} \tag{1}$$

If the notion of disadvantage is true to its name, then RD will typically be non-negative, as the resource acquisition rate will be higher within the advantaged group. As we see in Sec. 7, this is indeed the case for the vaccine allocation problem. Moreover, it is desirable not only to have RD **close to zero**, but also to make it **negative**, as downstream outcomes may improve with higher allocation rates among the disadvantaged. We expand on this in Sec. 5. This codifies Premise (B) and is the reason why the definition in Eq. (1) does not enforce positivity. That said, |RD| may also be handled with minimal modifications within the present methodology.

Proportional Allocation and Geographic Disparity. The baseline allocation is proportional to the population, that is $n_j = p_j$, for all j. This is fundamentally a notion of geographic equity that is an expected and desirable property of a potential allocation in many policies, including some recent resource allocation work, e.g. in COVID-19 testing allocation [29]. To allow ourselves to remain close to this property, we relax the strict equality of $n_j = p_j$ for all j to a distance on the k-simplex between $\mathbf{n}$ and $\mathbf{p}$. One can think of this as a form of headroom for social intervention, up to a limit that could be potentially determined by law or by agreement. We primarily consider the absolute ℓ_1 or relative ℓ_∞ distances:

$$\begin{aligned} \mathsf{d}_1(\mathbf{n}, \mathbf{p}) &= \sum_j |n_j - p_j| \\ &\text{or} \\ \mathsf{d}_\infty(\mathbf{n}, \mathbf{p}) &= \max_j \left| \frac{n_j}{p_j} - 1 \right| \end{aligned} \tag{2}$$

Note that d_∞ is the more stringent of the two distances as $\mathsf{d}_1 = \sum_j p_j \left| \frac{n_j}{p_j} - 1 \right| \le \mathsf{d}_\infty$. (The methodology could also handle absolute ℓ_∞ and relative ℓ_1 distances, given by $\max_j |n_j - p_j|$ and $\sum_j \left| \frac{n_j}{p_j} - 1 \right|$ respectively, and weighted variants.)

Allocation Objective. It is worth remarking that if $\rho_j = \beta_j$, that is if disadvantaged and advantaged subpopulations get their proportional shares of the allocated resources at each location, then proportional allocation yields zero rate disparity, RD = 0. Intuitively, this is a scenario where there is no access gap between the disadvantaged and advantaged. When there is a gap, we expect proportional allocation to not yield the best rate disparity. How can we allocate resources to improve disparity across subpopulations while maintaining a reasonable geographic parity? This trade-off leads to the central objective of this paper, which we can express at a high level as:

$$\begin{aligned} \mathbf{n}^\star = \quad & \arg\min \quad && \mathsf{RD}(\mathbf{n}, \mathbf{p}, N_\Sigma, P_\Sigma, \beta, \eta) \\ & \text{subject to} && \mathsf{d}_\cdot(\mathbf{n}, \mathbf{p}) \le \varepsilon \\ & && \alpha\mathbf{n} \le \mathbf{p}, \end{aligned} \tag{3}$$

where the last inequality is a *no-waste constraint*, which prohibits the allocation of more resources than people at an location j: $N_j \le P_j$. The $\cdot$ can be either 1 or ∞, indicating the distance used.

4 MODELING ACCESS

In this section, our scope is a single location j. As such, for clarity, we drop all j subscripts.

4.1 Resource Acquisition

To model access differences across disadvantaged and advantaged populations, we need to model how resources are acquired upon allocation. For this, we adopt a simple yet effective model: *Poisson acquisition until saturation or exhaustion.* Each subpopulation

P_Σ	total population size	$\mathbf{n}$	vector of all n_j
N_Σ	total resources available	β_j	percentage of disadvantaged at location j
α	resource availability, N_Σ/P_Σ	ρ_j	percentage of N_j acquired by the disadvantaged
k	number of locations	η	access gap between disadvantaged and advantaged
P_j	population size at location $j \in \{1, \cdots, k\}$	d_1	ℓ_1 distance
N_j	resources allocated to location j	d_∞	relative ℓ_∞ distance
p_j	fraction of total population at location j, P_j/P_Σ	ε	tolerance in distance from proportionality
n_j	fraction of total resources allocated to j, N_j/N_Σ	$\mathbf{n}^\star$	optimal allocation when distance d. is used
$\mathbf{p}$	vector of all p_j	RD	resource disparity, Eq. (1)

Table 1: Glossary of main notations

acquires a unit of resource according to an independent Poisson arrival process, which stops when one of two things happen: (1) an amount of resource equal to the subpopulation size is acquired by the subpopulation (saturation) or (2) a total amount of resource equal to the allocated size is acquired (exhaustion). This model allows us to account for access differences via the relative rate of each Poisson process. In what follows, we focus on a specific location j with population P where β fraction of the population is disadvantaged and N resources have been already allocated. Our goals are to specify the acquisition model, describe access gaps, and derive the shape of the acquisition function ρ. We assume that resources and subpopulation sizes are integers.

Base Model. Let λ^- and λ^+ be the rates of acquisition of the disadvantaged and advantaged respectively. When there are no access gaps, these rates reflect only the relative sizes of the subpopulations. We set the expected time for a single resource acquisition per person in this case as our unit of time, and treat the two populations as a split. This gives us the base rates of $\lambda^- = \beta P$ and $\lambda^+ = (1-\beta)P$.

Access Gap. Premise (C) states that the disadvantaged face obstacles in acquiring resources. We can model this as a *slowdown* relative to the base rate. Let η be the (relative) access gap between the disadvantaged and advantaged populations, where $\eta = 1$ signifies no gap and $\eta = 0$ is the limit of infinite gap. Under this model, the modified rate for the disadvantaged becomes $\lambda^- = \eta\beta P$ while the rate of the advantaged remains unchanged $\lambda^+ = (1-\beta)P$.

4.2 Naïve Acquisition Function

With this model, we can now concretely define the acquisition function ρ to be the expected fraction of the total allocated N resources that will be acquired by the disadvantaged. To develop intuition, let us start with a naïve characterization that ignores saturation. This can be thought of as the limit of very scarce resources ($\alpha \to 0$), so much so that no subpopulation in any location can be saturated. In this case, subpopulations acquire resources until exhaustion. If $\ddot{U}_t$ represents the disadvantaged process and $\ddot{V}_t$ represents the advantaged process, then the exhaustion time can be defined as $\ddot{S} = \min\{t \ : \ \ddot{U}_t + \ddot{V}_t \geq N\}$. With this, the acquisition function can be defined as $\ddot{\rho} := \mathbf{E}\left[\ddot{U}_{\ddot{S}}\right]/N$. The $\ddot{}$ notation is used to evoke *"naïve"*, and not second derivative. We have (the proof is given in the supplementary material):

Proposition 4.1. *If saturation is ignored, the naïve acquisition function $\ddot{\rho}$ depends only on β and η, and is of the form:*

$$\ddot{\rho} = \frac{\lambda^-}{\lambda^- + \lambda^+} = \frac{\eta\beta}{\eta\beta + 1 - \beta} \tag{4}$$

4.3 Exact Acquisition Function

The model above is unrealistic, because it has a positive probability for $\ddot{U}_{\ddot{S}} > \beta P$ or $\ddot{V}_{\ddot{S}} > (1-\beta)P$. As the availability of the resource increases, i.e., for larger values of α, this probability is non-negligible as subpopulations may saturate. We thus need an exact characterization of ρ that reflects this. Let the disadvantaged and advantaged processes in this case be U_t and V_t respectively. These are now *stopped* versions of the Poisson processes, $\ddot{U}_t$ and $\ddot{V}_t$, of the naïve case. More precisely, if we define saturation times of these processes as $T_U = \min\{t \ : \ \ddot{U}_t \geq \beta P\}$ and $T_V = \min\{t \ : \ \ddot{V}_t \geq (1-\beta)P\}$ respectively, i.e., the times when they hit their respective population sizes, then

$$U_t = \begin{cases} \ddot{U}_t & ; \quad \text{if } t \leq T_U \\ \beta P & ; \quad \text{if } t > T_U \end{cases} \qquad V_t = \begin{cases} \ddot{V}_t & ; \quad \text{if } t \leq T_V \\ (1-\beta)P & ; \quad \text{if } t > T_V \end{cases}$$

The exhaustion time is now modified to be: $S = \min\{t \ : \ U_t + V_t \geq N\}$, which is well-defined assuming that no wasteful allocations are allowed, so that $N \leq P$ always. The exact characterization of the acquisition function is now $\rho := \mathbf{E}\left[U_S\right]/N$. We have (the proof is given in the supplementary material):

Proposition 4.2. *When both saturation and exhaustion occur, the exact acquisition function ρ depends not only on β and η, but also on N and P. It can be written in terms of the naïve acquisition function $\ddot{\rho}$ of Eq.* (4) *as follows:*

$$\begin{aligned} \rho = \ddot{\rho} \quad & - \quad \ddot{\rho}\overline{F}\left(\beta P - 1; N-1, \ddot{\rho}\right) + \frac{\beta P}{N}\overline{F}\left(\beta P; N, \ddot{\rho}\right) \\ & + \quad (1-\ddot{\rho})\overline{F}\left((1-\beta)P - 1; N-1, 1-\ddot{\rho}\right) \\ & - \quad \frac{(1-\beta)P}{N}\overline{F}\left((1-\beta)P, 1-\ddot{\rho}\right) \end{aligned} \tag{5}$$

4.3.1 Tractable Approximation. The acquisition function directly affects the allocation objective expressed in Eq. 3 through the rate disparity, as seen in Eq. 1. While the characterization of ρ in Prop. 4.2 is exact, it is not directly amenable to reduction to a convex optimization problem. Fortunately, we can develop a tractable approximation to ρ that also helps develop intuition about its behavior. One can think of the acquisition process in two stages: (1) both subpopulations acquire simultaneously, (2) one subpopulation saturates, while the other continues to acquire.

To quantify these stages, first note that with unlimited resources, the advantaged saturation time is $t = 1$, on average. This is because the expected m^{th} arrival time of $\ddot{V}$ is m/λ^+, with $m = \lambda^+ = (1-\beta)P$. Similarly, the disadvantaged saturation time is $t = 1/\eta$, on average. As a result, it is highly unlikely for the disadvantaged to saturate prior to the advantaged. This sets the stage for the approximation, by taking these times to be equal to their expectations. The strength

of such an approximation stems from the fact that for large parameters, a Poisson random variable concentrates very sharply around its mean. The approximate model is:

- Before $t = 1$, $\mathbf{E}[U_t]$ and $\mathbf{E}[V_t]$ grow linearly at the rate of $\eta\beta P$ and $(1-\beta)P$ respectively. This stage thus reaches a maximal acquisition of $(\eta\beta + 1 - \beta)P$.
- After $t = 1$, $\mathbf{E}[V_t]$ saturates at $(1-\beta)P$, while $\mathbf{E}[U_t]$ continues to grow linearly at the rate of $\eta\beta P$.
- Exhaustion happens when $\mathbf{E}[U_t] + \mathbf{E}[V_t] = N$.

In this model, if $N < (\eta\beta + 1 - \beta)P$ we have $S \approx (N/P)/(\eta\beta + 1 - \beta)$ and thus $\rho = \mathbf{E}[U_S]/N \approx (\eta\beta PS)/N = \ddot{\rho}$. In other words, if exhaustion occurs during stage (1), the naïve model's acquisition function applies. Otherwise, the acquisition function is fully determined by the size of the saturated subpopulation, since all that is left for the other is the remainder of the resources. More precisely, in that case we have $\rho \approx [N - (1-\beta)P]/N$. We can write:

$$\rho \approx \begin{cases} \frac{\eta\beta}{\eta\beta+1-\beta} & ; \quad \text{if } N < \frac{1}{\alpha}P(\eta\beta + 1 - \beta) \\ 1 - \frac{(1-\beta)P}{N} & ; \quad \text{otherwise.} \end{cases}$$

Further, using the fact that $N/P = \alpha n/p$ and observing that $\frac{\eta\beta}{\eta\beta+1-\beta} > 1 - \frac{(1-\beta)p}{\alpha n}$ is equivalent to $n < \frac{1}{\alpha}p(\eta\beta + 1 - \beta)$, we can write the approximate acquisition function as:

$$\tilde{\rho} := \max\left\{\ddot{\rho}, 1 - \frac{(1-\beta)p}{\alpha n}\right\} \tag{6}$$

Eq. (6) clearly demarcates the two possible behaviors of $\tilde{\rho}$: the first dominated by exhaustion ($\ddot{\rho}$) and the second dominated by saturation. Note that the absolute sizes of N and P do not affect this approximation, only their relative size N/P matters. In fact, we can think of this approximation as the limit of ρ for large N and P. Fig. 1a illustrates this visually.

5 RESOURCE DISPARITY AND DOWNSTREAM IMPACT

Reducing resource disparity may be thought of as a goal in its own right, as a form of virtue or deontological ethics. However, in the present context one may argue that this goes hand in hand with a more consequentialist notion. Assume that, in the absence of the resource, the advantaged are likely to incur an adverse outcome with probability x and that the disadvantage are more likely to incur such an outcome, say with probability $(1+\delta)x$. These outcomes represent an abstraction of potential downstream impact. In the presence of the resource, say the adverse outcome's likelihood drops by a factor $q \in (0,1)$ for the advantaged and by a factor of q' for the disadvantaged. Thus, with the resource, the advantaged and disadvantaged have a probability qx and $q'(1+\delta)x$ of adverse outcomes respectively. We can then make the following observation.

Proposition 5.1. *In the adverse outcome scenario described above, the expected number of adverse outcomes is a linear function of* RD. *Furthermore, the slope of the relationship is positive whenever* $\delta > \frac{q'-q}{1-q'}$.

One can see this by noting that RD in Eq. (1) is a linear function of the number of vaccines that go to the disadvantaged and that the same is also true for the expected number of adverse outcomes. The condition of Prop. 5.1 then follows from a direct calculation. We omit the proof, as it is straightforward. Thus, if the disadvantage is large enough at first (large enough δ) and if the resource allows a decent improvement for the disadvantaged (q' not much larger than q) then reducing RD also reduces the expected number of adverse outcomes. The first facet is intuitive. To see why the second is also necessary, imagine that the resource did nothing to the disadvantaged, i.e., $q' = 1$. In that case, it doesn't help to channel more resources to them, and it is in fact more advantageous (for downstream impact) to channel more to the advantaged. However, these are arguably satisfied in practice, especially if the resource is even more effective for the disadvantaged ($q' < q$).

Example. Using the CDC estimates [5] in the period 10/2021 to 2/2022, the average rate of COVID19 hospitalization per 100,000 people in the unvaccinated (advantaged) 18-64 age group is 99 and the (disadvantage) 65+ group is 415. For the fully vaccinated, the respective rates are 18 and 78. This gives $q = 18/99 \approx 0.182$ and $q' = 78/415 \approx 0.188$, corresponding to a minuscule threshold of 0.007 that $\delta = 415/99 - 1 \approx 3.19$ readily exceeds. If one considers the 50+ group as advantaged instead, then one gets $q' < q$ and the condition holds trivially. □

6 ACCESS-AWARE ALLOCATION

We are now in position to instantiate concrete version of Eq. (3). To develop intuition, we first do this for the naïve acquisition model and show that it reduces to a linear program. We then show how the resulting solution can be used iteratively, to heuristically solve allocation under the approximate access model. As for the exact access model, one could in principle use nonlinear optimization with Eq. 5 directly, but we do not address this currently. Before proceeding, let us rewrite resource disparity in Eq. (1) as follows:

$$\begin{aligned} \text{RD} &= \alpha\frac{\sum_j(1-\rho_j)n_j}{\sum_j(1-\beta_j)p_j} - \alpha\frac{\sum_j \rho_j n_j}{\sum_j \beta_j p_j} \\ &= \sum_j \underbrace{\alpha\left(\frac{1-\rho_j}{\sum_{j'}(1-\beta_{j'})p_{j'}} - \frac{\rho_j}{\sum_{j'}\beta_{j'}p_{j'}}\right)}_{c_j} n_j \end{aligned} \tag{7}$$

6.1 Naïve Allocation

In the naïve model, resource disparity becomes a linear function of $\mathbf{n}$, because plugging in $\ddot{\rho}$ instead of ρ in the expression for c_j in Eq. (7) gives:

$$c_j = \frac{\alpha}{\eta\beta_j + 1 - \beta_j}\left(\frac{1-\beta_j}{\sum_{j'}(1-\beta_{j'})p_{j'}} - \frac{\eta\beta}{\sum_{j'}\beta_{j'}p_{j'}}\right),$$

which does not depend on $\mathbf{n}$. Furthermore, the no-waste constraint is already linear and the constraints with both $\mathrm{d}_1(\mathbf{n},\mathbf{p})$ and $\mathrm{d}_\infty(\mathbf{n},\mathbf{p})$ of Eq. (2) can be expanded into a set of linear constraints, via slack variables. As such, in this case the allocation problem (3) can be written as a linear program. Here are the instances that result for each distance:

(a) Acquisition Behavior

(b) Robustness to η

Figure 1: Visualizations — (a) Acquisition behavior versus time (with $\eta = 0.5$, $\beta = 0.5$). Blue and red plots are resources acquired by all and by the advantaged, respectively for $P = 10, 100, 1000, 10000$ (faintest to darkest), for the maximal expected time. For a specific N, resources are exhausted sooner (the orange example). For each case, the average resource is the solid line. The error bars represent the 5 to 95 percentile range. The limits (dotted red and blue) correspond to the approximation model in Section 4.3.1. These are indistinguishable from the average $P = 1000$ case and from the entire confidence range for $P = 10000$. The acquisition function is always the expected fraction of resources acquired by the disadvantaged upon exhaustion. This can have two distinct behaviors, based on whether only exhaustion occurs or both saturation and exhaustion occur. (b) Constraint set of the example in Sec. 6.3 (interior of the blue region). Direction of the vector c_j of RD in Eq. (7), as η varies (in pink). The narrow fan implies that there exists a single discrepancy-optimizing allocation that is optimal for all η, and thus robust to lack of specification of η.

$$
\begin{array}{lcll}
\mathbf{n}_1^\star & = & \arg\min & \sum_j c_j n_j \\
 & & \text{subject to} & \forall j \quad n_j, s_j \geq 0 \\
 & & & n_j - s_j \leq p_j \\
 & & & -n_j - s_j \leq -p_j \\
 & & & \alpha n_j \leq p_j \\
 & & & \textstyle\sum_j s_j \leq \varepsilon \\
 & & & \textstyle\sum_j n_j = 1 \\
\text{or} & & & \\
\mathbf{n}_\infty^\star & = & \arg\min & \sum_j c_j n_j \\
 & & \text{subject to} & \forall j \quad n_j, s_j \geq 0 \\
 & & & n_j - s_j \leq p_j \\
 & & & -n_j - s_j \leq -p_j \\
 & & & \alpha n_j \leq p_j \\
 & & & \forall j \quad s_j \leq \varepsilon p_j \\
 & & & \textstyle\sum_j n_j = 1.
\end{array}
\tag{8}
$$

6.2 Better Allocation

Eq. 8 demonstrates concretely how the access model influences the allocation decision. Access determines resource disparity, which in turn can be traded off with geographic disparity. While conveniently in linear program form, these do not capture the phenomenon of saturation that will occur whenever sufficient resources are available at a given location. The exact acquisition function in Eq. (5) does address this, but is highly nonlinear. Let us consider instead the approximate acquisition function in Eq. (6). The resulting RD is still nonlinear, because $\tilde{\rho}_j$ depends on n_j and p_j, and thus so will c_j. Moreover, at each location, Eq. (7) multiplies Eq. (6) by $-n$, yielding the minimum of two linear functions, which is thus concave. Since RD is the sum of these functions, the result is a concave minimization problem, which can generally be non-tractable. That said, this informs us that the solution remains on the boundary of the constraint set. In addition, the two distinct behaviors of $\tilde{\rho}$ suggest that one could start with a naïve model and then "correct" the model and repeat. This gives the following iterative heuristic:

(1) Solve the naïve allocation problem 8, assuming that exhaustion does not occur.
(2) Use the current allocation to detect exhaustion locations j and correct $\tilde{\rho}$ according to (6) at these locations.
(3) Re-solve the naïve allocation problem 8, but now with the c_j weights in (7) using the updated $\tilde{\rho}_j$.
(4) Repeat Steps (2) and (3) until convergence.

This is a heuristic both because $\tilde{\rho}$ is an approximation, even if an accurate one, and also because a priori the iterations could either not converge or converge to a suboptimal stable point. Nevertheless, in practice the iterations converge extremely quickly and to good solutions, i.e., solutions that significantly reduce resource disparity, as we illustrate experimentally in Sec. 7. In the supplementary material, we empirically verify the quality of these using exhaustive search.

6.3 Robustness to the Access Gap

To consider the allocation scheme of Sec. 6.2 as a viable alternative to the baseline proportional allocation, one should be able to readily implement it. Apart from design choices, such as the type of allowable deviation from proportionality, all variables are known but one: the access gap η. How should this be addressed? One approach is to fit the model to historical data, e.g., using a similar approach to the one we use to provide empirical evidence for the access model in Sec. 7.1. While viable, this requires careful elimination of confounding variables. Surprisingly, we experimentally observe that $n_1^\star$ and $n_\infty^\star$ are *robust* to the choice of η, in that the solution remains stable for a large range of values and in some cases, for *all* η. We give here an intuitive understanding of this robustness, through a simple example. In the supplementary material, we suggest more direct means to address not knowing η.

Example. Consider a case with only three locations, $j = 1, 2, 3$. Let the population be evenly distributed, $\mathbf{p} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Say locations have varying proportions of disadvantaged: low $\beta_1 = 0.2$, medium $\beta_2 = 0.5$, and high $\beta_3 = 0.8$. Assume resources are limited to 70% of the population, $\alpha = 0.7$. We are willing to deviate our allocation as much as $\varepsilon = 0.4$ away from $\mathbf{n}$ in ℓ_1, to reduce the resource disparity RD between the disadvantaged and advantaged. Fig. 1b displays the constraint set centered around the interior point $\mathbf{p}$. The fan shape represents the direction of the objective function vector c_j in RD of Eq. (7) at the optimal allocation, as η varies. Note that η does not affect the constraint set. Even if RD depends nonlinearly on $\mathbf{n}$, this informs us that the optimal allocation is indeed at the upper left corner $n_1^\star = \left(\frac{2}{15}, \frac{41}{105}, \frac{10}{21}\right)$ of the constraint polytope, since otherwise an infinitesimal gradient descent would move away. This implies that the solution is completely unaffected by η! Interestingly, this phenomenon also occurs with most of the real data cases that we cover in Sec. 7.2. Pragmatically, this means that we have an alternative to proportional allocation that *does not require knowledge of the access gaps* and yet results in a much more equitable distribution of resources. Note that the resulting allocation is intuitive: it boosts the highly disadvantaged location's allocation, reduces the low's allocation, and keeps the medium's roughly the same. □

7 EMPIRICAL VALIDATION

7.1 Evidence for the Access Model

To determine whether the proposed model of access has merit in practice, we look at how COVID-19 vaccine acquisition varies across locations. We cannot directly observe the number of vaccines acquired by the advantaged. We can, however, observe the overall vaccination rate. According to the the naïve model, i.e., assuming no saturation for simplicity, at any time t the overall vaccination rate is $\mathbf{E}\left[\ddot{U}_t + \ddot{V}_t\right]/(Pt) = (\eta\beta + 1 - \beta) = -(1-\eta)\beta + 1$. In what follows, we empirically show that we indeed observe this linear dependence on β, despite the potential presence of confounding variables.

Data. For the rates of COVID-19 vaccines, we use publicly available timeseries data from the CDC at the county level [3]. This data is provided to the CDC by state or territory health departments by healthcare providers. A few counties have no reported vaccinations, we discard these. The features we captured were the dates (`Date`), county name (`Recip_County`), state name (`Recip_State`), number of people who received 1 dose of a COVID-19 vaccine (`Administered_Dose1_Recip`), and their percentage at the county level (`Administered_Dose1_Pop_Pct`). This data reflects recipients over the age of 5. To characterize each county's overall vulnerability, we use publicly available data from Surgo Ventures' Precision for Community Vulnerability to COVID-19: Explore the U.S. Data Tool [31]. They look at many of the factors the CDC uses to judge vulnerability for COVID-19 such as socioeconomic status, household composition, racial status, existence of pre-exiting conditions, and accessibility to healthcare and transportation. They use these factors to calculate overall vulnerability at the county-level. From Surgo, we captured county name, state name, and vulnerability at (very high, high, moderate, low, very low) levels. To explore the applicability of our observations at every scale, we also look at the whole world, using Gallup numbers for vulnerability [11] and vaccination numbers from One World in Data [24]. We include only a subset of 142 countries that have clean representation across data sets.

Experiments. For each county in California, Illinois, Ohio, and Pennsylvania, we look at the percentage of the population that has acquired at least the first-dose of the COVID-19 vaccine on December 30, 2021. We do the same at the global level, per country. Let's call these numbers y_j. At the state level, we consider the proportion of the vulnerable at each county to be $\beta_j = \text{moderate}_j + \text{high}_j + 0.5 \times \text{very high}$, where the very high vulnerability range is slightly weighted down because it is often noisy and incomplete. For each of the 4 states, Fig. 2a shows a scatter plot of y (the bubbles have areas proportional to population size), along with its soft nearest-neighbor interpolation (with $\lambda = 20$): $\hat{y}(\beta) = \frac{\sum_j y_j \exp(-\lambda|\beta - \beta_j|)}{\sum_j \exp(-\lambda|\beta - \beta_j|)}$. In the global case, we set $\beta_j = \text{moderate}_j + \text{high}_j$, the 2 categories of Gallup, and illustrate the parallel results in Fig. 2b.

Results. We observe that, despite noise and confounding variables (e.g., differences between urban/suburban/rural regions, the deliberately ignored saturation which affects the low-β range, noise and outliers, etc.), the general trend of y_j versus beta is not only monotonically decreasing, but is indeed roughly linear. In particular, we do see clear segments of linear behavior for primary clusters of counties. We conclude that, though highly idealized, the presented access model has the potential to capture a realistic correlation between the disadvantaged's lack of resource and access obstacles. Note also the common behavior across states, highlighted with the red dotted line in Fig. 2a. We conjecture that this may be formalized as a methodology to learn η from the data. More strikingly, the same behavior extends all the way to the global scale, showing the ubiquity of the phenomenon anticipated by our model.

7.2 Behavior and Outcomes of Access-Aware Allocation

We now explore, numerically, the behavior and outcomes to expect when employing the proposed access-aware allocation, in comparison to proportional allocation. The latter has a simple characterization: we always have $\frac{n_j}{p_j} = \alpha$. To assess how access-aware allocation deviates from proportional allocation, we visualize the behavior of $\frac{n_j}{p_j}$ in this case, for locations with varying β_j. In order to demonstrate the possible gains achieved versus proportional allocation, we simultaneously visualize the latter's discrepancy RD along with the discrepancy achieved by access-aware allocation.

Data. In addition to the county-level CDC data, we also used data privately communicated by the Ohio Health Department, including tract-level maps with information about the number of vulnerable people in each tract, for 22 counties in Ohio[28]. The criterion for vulnerability in this case differs from Surgo's, but is based on similar metrics. Using tract-level data from the Census Bureau [2], we could then determine the population size at each tract, and thus the percentage vulnerable (β) in each. This allows us to explore the applicability of the present methodology at such finer granularity. We also explore a much coarser granularity, with allocation of vaccines across all 50 states in the US. The data for this is an aggregation of that obtained from CDC and Surgo, as

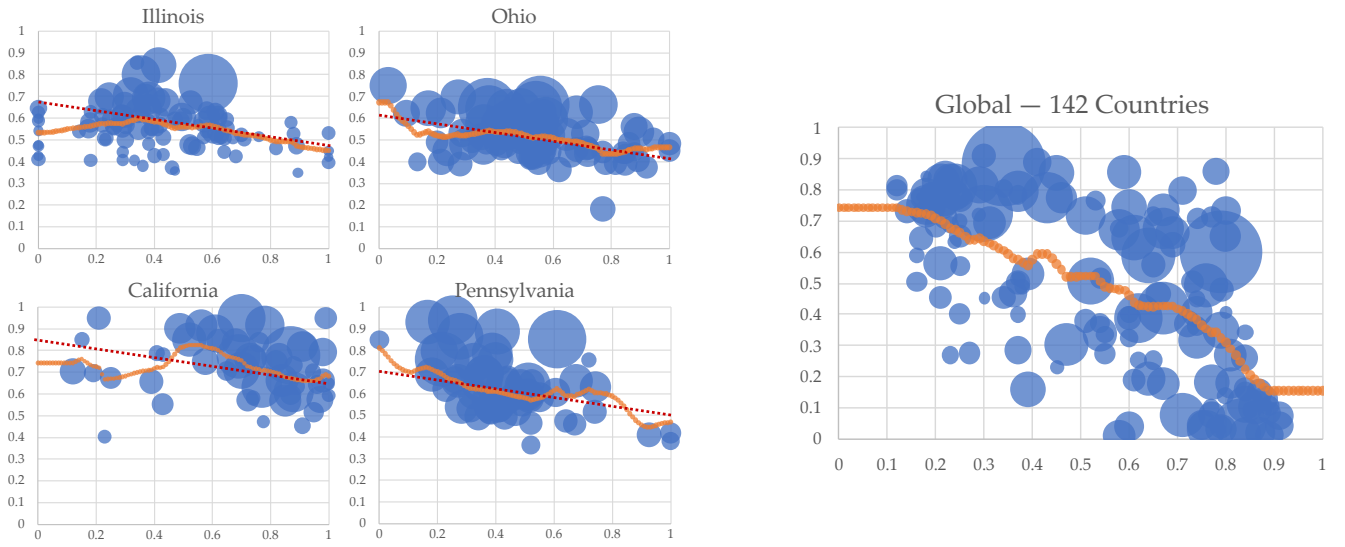

(a) U.S. States **(b) Global**

Figure 2: Vaccination rate vs. vulnerability — (a) Each bubble represents a county, with area proportional to its population. The x-axis is the estimated percentage of vulnerable in the county (β). The y-axis is the percentage of the overall population with their first COVID-19 vaccination dose as of 12/30/2021. The orange curve is a soft nearest-neighbor interpolation $\hat{y}(\beta)$ of the relationship between β and y. The dotted red is a visual guide illustrating the common dominant behavior across states. (b) Global manifestation of the same phenomenon as in Fig. 2a — A roughly linear inverse relationship between vulnerability numbers and vaccination rates across countries.

described in Sec. 7.1. The vulnerable percentage in this section is taken to be everyone categorized as very high, high, or moderate.

Experiments. We fix the value of ε to 0.1 throughout. The influence of ε on the allocation is apparent nonetheless, and we comment on it below. We look at the states of California, Illinois, Georgia, Ohio, and Pennsylvania (locations = counties), in addition to Franklin County, OH (locations = Census tracts), and the US (locations = states). We consider $\alpha = 0.1$ (low), 0.5 (medium), and 0.9 (high) levels of resource availability. We vary η from 0 to 1. In each scenario, we solve for $n_1^\star$ and $n_2^\star$ using the iterative approach in Sec 7.2, by solving Eq. 8 and alternating with updates of ρ.

Results — Behavior. Fig. 3a illustrates the behavior of the $n_1^\star$ allocation through the ratio $\frac{n_j}{p_j}$. The dotted line corresponds to α, the proportional allocation case. We observe that in all cases the access-aware allocation has roughly the same form: no allocation $\frac{n}{p} = 0$ below a lower threshold, full allocation $\frac{n}{p} = 1$ above an upper threshold, and proportional allocation $\frac{n}{p} = \alpha$ in between. The value of ε influences how close the upper and lower thresholds are for fixed α. This is a harsh policy and is due to the lax nature of d_1 that allows such deviation from proportionality. It is worth noting that this is not always the form of $n_1^\star$, as we have seen in the example of Sec. 6.3.

Fig. 3b illustrates the behavior of the $n_\infty^\star$ allocation, which is much simpler: slightly less allocation below a threshold $\frac{n}{p} < \alpha$ and slightly more allocation above a threshold $\frac{n}{p} > \alpha$. The value of ε determines the amount of change from α. It is also interesting to note that in this case, unlike the thresholds of $n_1^\star$, this threshold does not depend on α. This is a lax policy, giving a much gentler deviation from proportionality, due to the harsh nature of d_∞.

Lastly, we note that we do not display multiple policies for each η, because under each scenario the policy is unaffected by the choice of η, thus supporting the claims of robustness in Sec. 6.3.

Results — Outcomes. Fig. 4a and Fig. 4b illustrate the improvement in resource disparity that $n_1^\star$ and $n_\infty^\star$ afford. Each constraint results in specific characteristics. As we could expect from their harsh and lax respective behaviors, $n_1^\star$ can deliver a more marked improvement to resource disparity than $n_\infty^\star$ can.

More interestingly, under $n_1^\star$, access-awareness is generally (though not always) more impactful at higher access gaps (small η) and less so at lower gaps (large η). There isn't such a monotonic relationship for $n_\infty^\star$, with RD often peaking in η in that case. This seems to suggest that when too much deviation from proportionality isn't desirable (so one would best use d_∞), then allocation, to be effective, needs to be compounded with enough social effort to bring the disadvantaged above a certain level of access beyond which it is able to mitigate any remaining difference. However, if more bold redistribution is accepted (thus one could use d_1) then allocation on its own can go a long way to mitigate disparity. Thus this work reveals an interesting tradeoff. To borrow some political jargon, if one is "conservative" in resource allocation, they ought to be more "progressive" in providing social safety nets. Conversely, if one is "socialist" in resource allocation, they can afford to be less regulated and more "libertarian" when it comes to society.

Both instances do however share a common yet counterintuitive pattern, in that the disparity reduction is relatively higher when resources are more available ($\alpha = 0.9$), especially in moderate access gaps (mid to high η). This is due to saturation effects, as without access-awareness, the advantaged subpopulation can quickly acquire all the needed resources, leaving scraps to the disadvantaged.

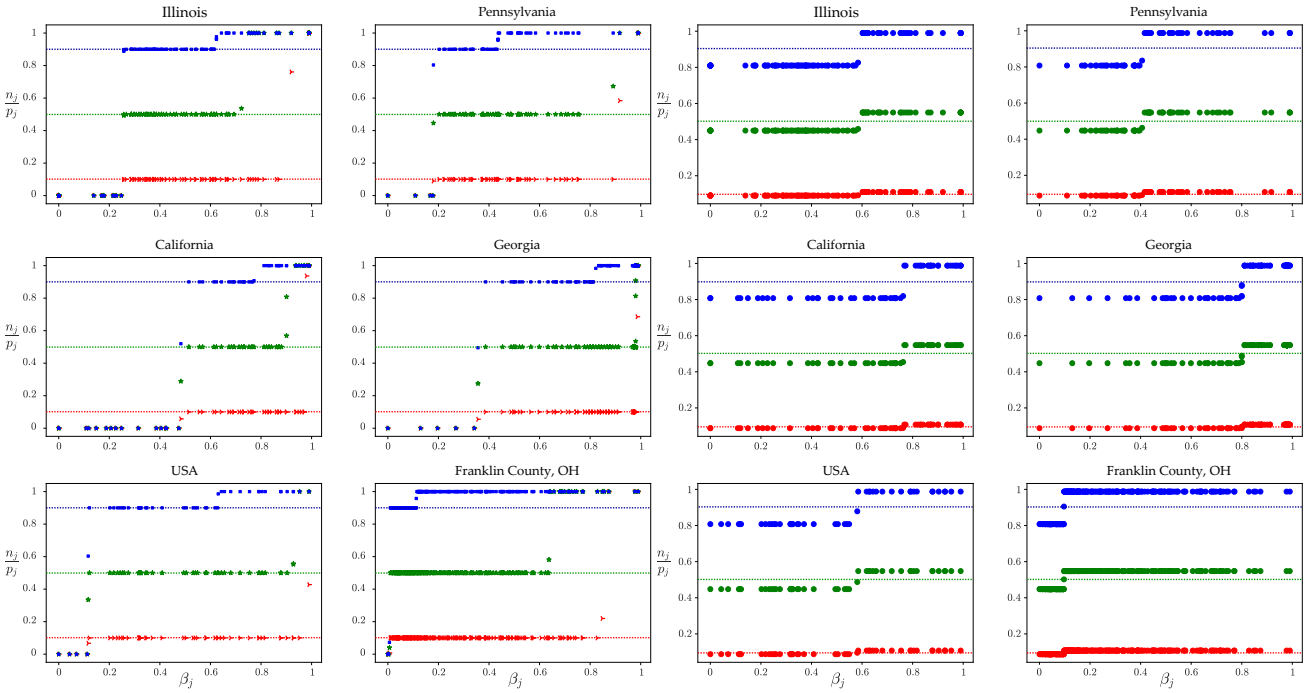

(a) Behavior of $n_1^\star$ **(b) Behavior of $n_\infty^\star$**

Figure 3: Behavior of Allocations — Each point is a location j. The x-axis is the percentage vulnerable β at location j and the y-axis is the ratio of allocation to its population $\frac{n}{p}$. The allocation for $\alpha = 0.1$, $\alpha = 0.5$, and $\alpha = 0.9$ are also indicated with dotted lines. Main observations — (a) $n_1^\star$ specifies two thresholds, does not allocate below the lower, fully allocates above the higher, and maintains α in between. (b) $n_\infty^\star$ fixes a threshold and alters the proportional allocation slightly lower below it and slightly higher above it.

What is practically most significant, however, is that at moderate availability ($\alpha = 0.5$) the improvement in RD is consistent in both instances, no matter what the value of η is. We see this visually as the RD plots having shifted down. This highlights the potential impact that access awareness could have on distributing resources more equitably, no matter what the gap is between the advantaged and disadvantaged.

A final phenomenon that is worth clarifying is the following. Imagine $\eta = 0$, meaning that no matter how resources are allocated, the disadvantaged acquire *nothing*. In that case, one cannot hope to improve on proportional allocation (or for that matter on any non-wasteful allocation), since all resources will then flow to the advantaged, and no reallocation can change it. One would then expect, near $\eta = 0$, for the RD of the access-aware allocations to approach that of the proportional allocation. Yet, we observe a clear separation between the two. How can we explain this? This is due to the fact that our access model operates with an infinite time horizon, which creates a discontinuity. While the disadvantaged acquire nothing at $\eta = 0$, even a modicum of access $\eta > 0$ no matter how small, the disadvantaged can acquire whatever is leftover after the advantaged saturate. They can do this, even if it takes them an arbitrarily long time. And this offers an opportunity for access-aware allocation to make a difference. This also opens up an avenue to explore finite time horizon effects by appropriately modifying the access model.

8 DISCUSSIONS

Throughout this work, we have considered technical means by which to address issues with access to limited resources for members of vulnerable groups. We acknowledge that this is an oversimplification of society and that there are many factors, both endogenous and exogenous, that can hamper a person's acquisition of a limited resource.

Other Factors of Slowed Access. Often, members of vulnerable groups have been marginalized and mistreated by systems offering aid and assistance to illness or societal issues (e.g. Tuskegee Experiments) [32]. This mistreatment has lead to mistrust and reluctance to accepting proposed solutions amongst some subpopulations. In addition, there is often a wave of misinformation distributed in mass on public concerns [25]. Without the correct information on a resource, people are robbed of the opportunity to decide whether obtaining a resource will truly be beneficial for them. Educating people on the true benefits and disadvantages of obtaining a resource could help in their decision-making. Even with a trustworthy system, true education on obtaining a resource, and the solution of proximity issues, more barriers exist, e.g., lack of internet, inflexible jobs, mobility issues, or no health insurance [18].

Calls for Justice. Our approach to mitigating slowed down access through increased proximity could be interpreted as a form of justice, in that those a part of the disadvantaged population have often been marginalized, creating lack of access and exacerbating

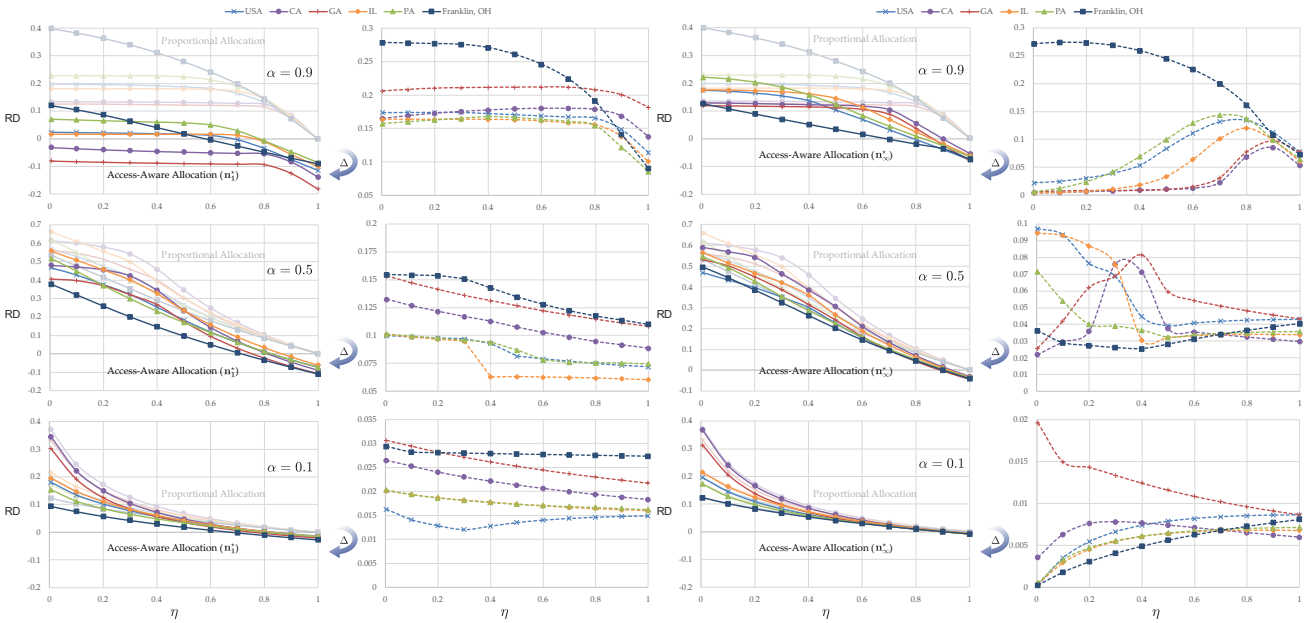

(a) Outcomes of $n_1^\star$ **(b) Outcomes of $n_\infty^\star$**

Figure 4: Outcomes of Allocations — (Left) Resource (vaccination) disparity (y-axis, RD) of both the access-aware allocation $n_\infty^\star$ offers and the proportional allocation (faint plots), for $\alpha = 0.1, 0.5, 0.9$. (Right) Difference between the two, which captures the improvement in disparity by using the access-aware allocation. The x-axis is the access gap, η. Main observations — (a) For $n_1^\star$, improvements are more marked when the access gap is larger, i.e. when η is smaller. Counterintuitively, more availability (larger α) offers relatively more opportunity to mitigate disparity, as explained in the text. (b) For $n_\infty^\star$, the general behavior is comparable to that of $n_1^\star$. However, in this case we are more constrained not to deviate from proportionality, and the gains are thus less.

the effects of this deficit. This marginalization has been caused by multiple factors and reared in various ways, some include: redlining which hurt access to jobs and healthcare [15], social exclusion which can lead to lack of education, and inability to gather those resources [1], etc. Such history has shifted the distribution of power [19]. Though these effects are deeply rooted and ring throughout the lives of the disadvantaged, the hope is that taking such steps will encourage appropriate visibility of these members of society and, by doing so, work toward reversing injustice.

ACKNOWLEDGMENTS

This paper is based upon work supported in part by the NSF LSAMP Bridge to the Doctorate Program under Award No. HRD-1906075 and the NSF Program on Fairness in AI in Collaboration with Amazon under Award No. IIS-1939743, titled FAI: Addressing the 3D Challenges for Data-Driven Fairness: Deficiency, Dynamics, and Disagreement. Any opinion, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation or Amazon.

REFERENCES

[1] Margaret M Alston and Jenny Kent. 2003. Educational access for Australia's rural young people: A case of social exclusion. *Australian Journal of Education* 47, 1 (2003), 5–17.

[2] US Census Bureau. 2020. *2020 Census Data.* https://www2.census.gov/geo/docs/reference/2020addresscountlist/39_Ohio_AddressBlockCountList_012020.txt

[3] CDC. 2021. *COVID-19 Vaccinations in the United States,County.*

[4] CDC. 2021. *Covid-19 vaccine administration, by Race and ethnicity - North Carolina, December 14, 2020–April 6, 2021.* https://www.cdc.gov/mmwr/volumes/70/wr/mm7028a2.htm

[5] CDC. 2022. *CDC Covid Data tracker.* https://covid.cdc.gov/covid-data-tracker/#covidnet-hospitalizations-vaccination

[6] CDC. 2022. *Risk for COVID-19 infection, hospitalization, and death by age group.* https://www.cdc.gov/coronavirus/2019-ncov/covid-data/investigations-discovery/hospitalization-death-by-age.html

[7] Xin Chen, Menglong Li, David Simchi-Levi, and Tiancheng Zhao. 2020. Allocation of COVID-19 vaccines under limited supply. *Available at SSRN 3678986* (2020).

[8] Terri Moon Cronk. 2020. *Officials Optimistic About COVID-19 Vaccines in Near Future.* Retrieved August 8, 2022 from https://www.defense.gov/News/News-Stories/Article/Article/2420190/officials-optimistic-about-covid-19-vaccines-in-near-future/

[9] Mohammad Reza Davahli, Waldemar Karwowski, and Krzysztof Fiok. 2021. Optimizing COVID-19 vaccine distribution across the United States using deterministic and stochastic recurrent neural networks. *PLOS One* 16, 7 (2021), e0253925.

[10] Kate Donahue and Jon Kleinberg. 2020. Fairness and utilization in allocating resources with uncertain demand. In *Proceedings of the 2020 conference on fairness, accountability, and transparency.* 658–668.

[11] Neli Esipova, Julie Ray, and Han Ying. 2020. *750 Million Struggling to Meet Basic Needs With No Safety Net.* Retrieved August 8, 2022 from https://news.gallup.com/poll/312401/750-million-struggling-meet-basic-needs-no-safety-net.aspx

[12] Hossein Estiri, Zachary H Strasser, Jeffy G Klann, Pourandokht Naseri, Kavishwar B Wagholikar, and Shawn N Murphy. 2021. Predicting COVID-19 mortality with electronic medical records. *NPJ digital medicine* 4, 1 (2021), 1–10.

[13] National Center for Immunization, Division of Viral Diseases Respiratory Diseases (NCIRD), and CDC. 2022. *Covid-19 information for specific groups of people.* https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/index.html

[14] Assistant Secretary for Preparedness and Response (ASPR). 2022. *HHS Expands Availability of Monkeypox Vaccine to More Than 1.1 Million Doses.* Retrieved August 8, 2022 from https://www.hhs.gov/about/news/2022/07/28/hhs-expands-availability-of-monkeypox-vaccine-to-more-than-1-million-doses.html1

[15] Judith Garber. 2021. *Racist redlining policies still have an impact on health.* https://lowninstitute.org/racist-redlining-policies-still-have-an-impact-on-health/

[16] Nikhil Garg, Hannah Li, and Faidra Monachou. 2020. Dropping Standardized Testing for Admissions Trades Off Information and Access. *arXiv preprint arXiv:2010.04396* (2020).

[17] Shikha Garg, Lindsay Kim, Michael Whitaker, Alissa O'Halloran, Charisse Cummings, Rachel Holstein, Mila Prill, Shua J Chai, Pam D Kirley, Nisha B Alden, et al. 2020. Hospitalization rates and characteristics of patients hospitalized with laboratory-confirmed coronavirus disease 2019—COVID-NET, 14 States, March 1–30, 2020. *Morbidity and mortality weekly report* 69, 15 (2020), 458.

[18] Muriel Jean-Jacques and Howard Bauchner. 2021. Vaccine distribution—equity left behind? *JAMA* 325, 9 (2021), 829–830.

[19] Maximilian Kasy and Rediet Abebe. 2021. Fairness, equality, and power in algorithmic decision-making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 576–586.

[20] James B Kirby and Toshiko Kaneda. 2010. Unhealthy and uninsured: exploring racial differences in health and health insurance coverage using a life table approach. *Demography* 47, 4 (2010), 1035–1051.

[21] Sharon LaFraniere, Noah Weiland, and Joseph Goldstein. 2020. *Officials Optimistic About COVID-19 Vaccines in Near Future.* Retrieved August 8, 2022 from https://www.nytimes.com/2022/08/03/us/politics/monkeypox-vaccine-doses-us.html?campaign_id=60&emc=edit_na_20220803&instance_id=0&nl=breaking-news&ref=cta&regi_id=170156702&segment_id=100369&user_id=c5dcb52b3d26acd5c75e18a3d0cf5e91

[22] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. 2021. Fair dynamic rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 694–695.

[23] Jayson S. Marwaha and Joseph C. Kvedar. 2022. Crossing the chasm from Model performance to clinical impact: The need to improve implementation and evaluation of ai. *npj Digital Medicine* 5, 1 (2022). https://doi.org/10.1038/s41746-022-00572-2

[24] Edouard Mathieu, Hannah Ritchie, Esteban Ortiz-Ospina, Max Roser, Joe Hasell, Cameron Appel, Charlie Giattino, and Lucas Rodés-Guirao. 2021. A global database of COVID-19 vaccinations. *Nature human behaviour* 5, 7 (2021), 947–953.

[25] Salman Bin Naeem, Rubina Bhatti, and Aqsa Khan. 2021. An exploration of how fake news is taking over social media and putting public health at risk. *Health Information & Libraries Journal* 38, 2 (2021), 143–149.

[26] Keith Norris and Allen R Nissenson. 2008. Race, gender, and socioeconomic disparities in CKD in the United States. *Journal of the American Society of Nephrology* 19, 7 (2008), 1261–1270.

[27] NC Department of Health and Human Services. 2021. NC Providers Increase Equitable Vaccine Distribution With NCDHHS Census Tract COVID-19 Map. *NCDHHS* (2021). https://www.ncdhhs.gov/news/press-releases/nc-providers-increase-equitable-vaccine-distribution-ncdhhs-census-tract-covid

[28] Ohio Department of Health. 2020. *Health360 Coronavirus Response Management Analytics.*

[29] Georgia Perakis, Divya Singhvi, Omar Skali Lami, and Leann Thayaparan. 2021. COVID-19: A Multipeak SIR Based Model for Learning Waves and Optimizing Testing. *Available at SSRN 3817680* (2021).

[30] Selena Simmons-Duffin. 2022. *Free Covid tests and treatments no longer free for uninsured, as funding runs out.* https://www.npr.org/sections/health-shots/2022/03/29/1089355997/free-covid-tests-and-treatments-no-longer-free-for-uninsured-as-funding-runs-out

[31] Surgo Ventures. 2021. *Precision for COVID- Community Vulnerability to COVID-19: Explore the U.S. Data Tool.*

[32] Jason Wasserman, MA Flannery, and JM Clair. 2007. Rasing the ivory tower: the production of knowledge and distrust of medicine among African Americans. *Journal of Medical ethics* 33, 3 (2007), 177–180.

[33] Leia Wedlund and Joseph Kvedar. 2021. New machine learning model predicts who may benefit most from COVID-19 vaccination. *NPJ Digital Medicine* 4, 1 (2021), 1–1.

[34] Yumeng Zhang, Huazhen Yang, and Jay Pan. 2021. Gaining from rational health planning: Spatial reallocation of top-tier General Hospital Beds in China. *Computers & Industrial Engineering* 157 (2021), 107344. https://doi.org/10.1016/j.cie.2021.107344

[35] Qian Zhou, Zhi-hang Chen, Yi-heng Cao, and Sui Peng. 2021. Clinical impact and quality of randomized controlled trials involving interventions evaluating artificial intelligence prediction tools: a systematic review. *NPJ digital medicine* 4, 1 (2021), 1–12.