On the Resolution of Ties in Fair Convex Allocation Problems

Bismark Singh[a]

[a] *School of Mathematical Sciences, University of Southampton, UK*

Abstract

We study the emergence of indistinguishable, but structurally distinct, allocation outcomes in convex resource allocation models. Such outcomes occur when different users receive proportionally identical allocations despite differences in initial conditions, eligibility sets, or priority weights. We formalize this behavior and analyze the structural conditions under which it arises, with a focus on fairness-oriented objectives. While the remedy — an infinitesimal perturbation — is classical, we show it can be interpreted as a principled secondary criterion for breaking ties that preserves the fairness structure encoded in the original problem. We apply this framework to a previously studied vaccine distribution model, demonstrating that such outcomes persist across multiple parameter regimes and can be systematically resolved without distorting the fairness scheme.

Keywords: convex optimization, fairness, tie-breaking, distinct solutions, perturbation

---

Email address: `b.singh@southampton.ac.uk` (Bismark Singh)

1. Background

Optimization models for resource allocation are pervasive in operations research and mathematical economics, see, e.g., [10, 5, 18]. Many such models employ convex objective functions combined with eligibility constraints [3, 8], and their practical motivation is often grounded in notions of fairness across differently weighted users [2, 9].

Singh [19] recently formalized a fairness criterion for one such class of models. The framework considers multiple users $i \in I$, each with an initial coverage level $f_i \in [0,1]$ and a final coverage level $y_i \in [0,1]$ after allocation. Substitutable resources $k \in K$ are available in limited quantities b_k, and users are eligible to receive only a subset of these. Each user has a weight $w_i > 0$ encoding their priority and a population size n_i. Similar to a social welfare framework (see, e.g., [6, 12]), the goal is to minimize a weighted total loss:

$$L(y) = \sum_{i \in I} w_i n_i F(y_i),$$

where F is a convex, decreasing function capturing the loss from a shortage (i.e., under coverage). This yields the following convex optimization model:

$$z_F^* = \min_y \sum_{i \in I} w_i n_i F(y_i) \tag{1a}$$

$$\text{s.t.} \quad f_i + \frac{1}{n_i} \sum_{k \in K_i} x_{ik} = y_i, \quad \forall i \in I, \tag{1b}$$

$$y_i \in [0,1], \quad \forall i \in I, \tag{1c}$$

$$\sum_{i \in I_k} x_{ik} \le b_k, \quad \forall k \in K, \tag{1d}$$

$$x_{ik} \ge 0. \tag{1e}$$

Under mild assumptions — F strictly convex, decreasing, and differentiable — Singh [19] prove a fairness guarantee: in an optimal solution, a higher-weighted user cannot achieve a lower coverage than a lower-weighted one unless the latter already began with a substantially larger f_i. Examples of such functions include $F(z) = (1-z)^m$ for $m > 1$, and $F(z) = -\ln(z+\varepsilon)$ for small $\varepsilon > 0$ [19].

However, even when prioritization is respected, the model may produce *strict unclear priority* (SUP) outcomes: situations where two or more users receive identical final coverage. We formally define this notion in Section 2. These ties can arise despite differences in weights, initial coverages, or eligibility sets. Although mathematically optimal and consistent with the fairness framework, these solutions obscure the intended prioritization and may complicate public justification of the allocation [1].

Example 1. *Consider an instance of model (1) with two users ($|I| = 2$) and one resource ($|K| = 1$), with both users eligible. Let $w_1 = 2$, $w_2 = 1$, $n_1 = n_2 = 100$, $f_1 = 0.1$, and $f_2 = 0.4$.*

(a) If $b_1 = 30$, the optimal solution allocates all units to user 1, giving $y_1 = y_2 = 0.4$.

(b) If $b_1 = 150$, both users reach full coverage $y_1 = y_2 = 1$.

In both cases the solution respects the priority orderings, however the final coverages are non-distinct. □

Ties such as those in Example 1 are not mere edge cases; they follow naturally from the convex structure and overlapping eligibilities of resources. While fairness (in the sense of [19]) is preserved, such solutions obscure the intended prioritization among users. This issue is not merely academic. In practice — whether in public health, education, or disaster relief — decision-makers must justify allocation outcomes to stakeholders, often under public scrutiny. Consider the following illustrative scenarios:

(a) Public health: During vaccine allocation, groups such as frontline workers, the elderly, and high-risk individuals may be assigned different weights. If two groups with clearly different risk profiles receive identical coverage, the allocation may appear opaque or unjustified—— even if mathematically fair; see also Yi and Marathe [21].

(b) Educational grants: A funding agency may prioritize applicants from underrepresented backgrounds by assigning higher weights. If the resulting grant amounts are equal across groups, the process may be perceived as arbitrary or inconsistent with stated goals; see also Corbett-Davies et al. [4].

(c) Disaster relief: Ties in aid distribution between regions with differing levels of damage and vulnerability may undermine the perceived responsiveness or credibility of relief efforts; see also Mazepus and van Leeuwen [11].

In each of these examples, decision-makers may be challenged to explain why distinct priorities did not lead to visibly distinct outcomes. This motivates our focus on *distinctness*: ensuring that users with different weights receive different final coverages. To achieve this, we propose a classical idea in optimization: adding a small, separable penalty via a proximal term to discourage equal coverage levels. This perturbation resolves ties by seeking directions that cause the smallest possible increase in the original loss function. Consequently, users whose marginal cost is smallest are the ones whose coverage is most likely to be adjusted; we formalize this later. In this sense, tie-breaking respects the priority structure already embedded in the model, without imposing any additional ranking beyond that implied by the original loss function itself.

The following are the main contributions of this work:

(i) We define and classify strict unclear priority outcomes in convex resource-allocation models with attention to how eligibility structures and smooth objectives interact.

(ii) We show that a vanishing perturbation resolves non-distinct outcomes while preserving the original fairness guarantees.

(iii) We validate the approach using a case study derived from a real-world vaccine allocation problem during the 2009 H1N1 pandemic, showing that SUP outcomes arise both under equal and unequal weights, and across a range of resource regimes.

The structure of the rest of the work is as follows. Section 2 formalizes the problem formally and presents a taxonomy of SUP cases. Section 3 introduces the perturbed model and studies its key properties. Section 4 presents our numerical study, while Section 5 concludes with a summary and implications for both theory and practice.

2. Problem Setup

We study the distinctness of optimal allocations under the following assumptions:

Assumption 1. *Throughout this work, we assume:*

(a) Distinct weights*: Users are relabeled such that $w_1 > w_2 > \cdots > w_{|I|}$; i.e., $w_j > w_i$ for all pairs $j < i$.*

(b) Unmet initial demand: *Every user has* $f_i < 1$.

(c) Scarce resources: *For each* $k \in K$, $\sum_{i \in I_k} x_{ik} = b_k < \sum_{i \in I_k} n_i(1 - f_i)$, *implying every resource is fully utilized [19, Theorem 3].*

(d) Loss function: *The function* $F\colon [0,1] \to \mathbb{R}$ *is continuously differentiable, strictly convex, and strictly decreasing; i.e.,* $F'(z) < 0$ *for all* $z \in [0,1]$.

Assumption 1 enforces a strict priority ordering and rules out trivial saturation of users. Further, we restrict attention to user pairs that actually compete for at least one common resource. Suppose a pair (j, i) ties; i.e., $y_j = y_i$, and $K_j \cap K_i \neq \varnothing$. If $x_{ik} = x_{jk} = 0$ for all shared $k \in K_j \cap K_i$, two cases arise:

(i) Neither user receives any allocation. This can occur only if initial coverage is already so high that no resource is allocated to them.

(ii) Each user receives allocations exclusively from non-shared resources.

Case (ii) implies no direct competition and is therefore excluded; case (i) yields no priority comparison. Hence, we analyze only pairs receiving some shared resource.

Definition 1. *A pair* (j, i) *with* $j < i$ *is* relevant *if at least one of them receives a positive allocation of a shared resource; i.e., for some* $k \in K_j \cap K_i$, *we have* $\max\{x_{ik}, x_{jk}\} > 0$.

Definition 2. *For a relevant pair* (j, i), *exactly one of the following holds* *at optimality:*

$$\begin{aligned} &\textit{Strict Appropriate Priority (SAP):} && y_j > y_i, \\ &\textit{Strict Inappropriate Priority (SIP):} && y_j < y_i, \\ &\textit{Strict Unclear Priority (SUP):} && y_j = y_i. \end{aligned}$$

Definition 3. *An optimal allocation* (x, y) *of model* (1) *is* distinct *if* $y_j \neq y_i$ *for every relevant pair* $j < i$*. Equivalently, no SUP cases occur.*

Under Assumption 1, model (1) is strictly convex with linear constraints; hence the Karush Kuhn Tucker(KKT) conditions are necessary and sufficient for optimality. Let $\alpha_i \in \mathbb{R}$, $\nu_k \geq 0$ and $\mu_{ik} \geq 0$ be the dual multipliers for constraints (1b), (1d) and (1e), respectively, and λ_i be that for $y_i \leq 1$. Then, the KKT system is:

$$\text{stationarity:} \quad \mu_{ik} = \nu_k + w_i F'(y_i) + \tfrac{\lambda_i}{n_i}, \quad \forall i \in I,\ k \in K_i \tag{2a}$$

$$\text{complementary slackness:} \quad \mu_{ik} \cdot x_{ik} = 0, \qquad \forall i \in I,\ k \in K_i \tag{2b}$$

$$\nu_k \cdot \left(b_k - \sum_{i \in I_k} x_{ik} \right) = 0, \qquad \forall k \in K \tag{2c}$$

$$\lambda_i \cdot (1 - y_i) = 0, \qquad \forall i \in I \tag{2d}$$

$$\text{primal feasibility:} \quad x_{ik} \geq 0, \quad \sum_{i \in I_k} x_{ik} \leq b_k, \qquad \forall i \in I,\ k \in K \tag{2e}$$

$$f_i + \tfrac{1}{n_i} \sum_{k \in K_i} x_{ik} = y_i \leq 1, \qquad \forall i \in I \tag{2f}$$

For a relevant pair (j, i) sharing a resource $k \in K_j \cap K_i$, three allocation patterns are possible:

$$
\begin{aligned}
&\text{(P1)} \quad x_{jk} > 0,\ x_{ik} > 0,\\
&\text{(P2)} \quad x_{jk} = 0,\ x_{ik} > 0,\\
&\text{(P3)} \quad x_{jk} > 0,\ x_{ik} = 0,
\end{aligned}
$$

Proposition 1. *Under Assumption 1, let (j, i) be a relevant pair with $y_j = y_i = y^\star$ in an optimal solution. Then exactly one of the following holds:*

A. $y^\star = 1$: any pattern (P1)–(P3) can occur.

B. $y^\star < 1$: only pattern (P3) can occur.

Proof. Choose $k \in K_j \cap K_i$ with $\max\{x_{ik}, x_{jk}\} > 0$. From (2a),

$$(w_j - w_i)F'(y^\star) = \mu_{ik} - \mu_{jk} + \frac{\lambda_i}{n_i} - \frac{\lambda_j}{n_j}.$$

A. If $y^\star = 1$, λ_i and λ_j may be non-zero, and any sign on the right-hand side is possible; patterns (P1)—(P3) are feasible.

B. If $y^\star < 1$, complementary slackness gives $\lambda_i = \lambda_j = 0$. Because $F'(y^\star) < 0$ and $w_i > w_j$, the left-hand side is positive, implying $\mu_{ik} > \mu_{jk}$. Hence $\mu_{ik} > 0$ and, by (2b), $x_{ik} = 0$. Only pattern (P3) remains consistent.

□

Proposition 1 enumerates all SUP configurations. In the next section, we present a perturbation approach that converts any optimal SUP solution into a distinct one while seeking to preserving fairness.

3. Perturbations of the Objective Function

The analysis so far has classifies when and how ties (i.e., SUP outcomes) may arise in optimal solutions of model (1) using the concept of relevant priority pairs. We now propose a refinement that enforces $y_i \neq y_j$ not only for such relevant pairs, but for *all* pairs of users satisfying Assumption 1. This perturbation eliminates all SUP outcomes by construction irrespective of the resource overlap structure.

The refinement proceeds by augmenting the objective function of model (1) with a deterrent term that penalizes equal coverage levels. Such regularization has precedent not only in continuous optimization through proximal terms for strict convexity or faster convergence [17, 13] but also in combinatorial optimization, for instance, to break symmetry in mixed-integer programs [15, 20].

For a small $\varepsilon > 0$, define a penalty term

$$\Psi_\varepsilon(y) = \varepsilon \sum_{\substack{i \in I,\, j \in I \\ j < i}} h(y_i - y_j),$$

where the function $h : \mathbb{R} \setminus \{0\} \to \mathbb{R}$ satisfies:

(i) $h(d) \to \infty$ as $d \to 0$;

(ii) $h(d) > 0$ for all $d \neq 0$.

Representative examples include $h(d) = d^{-2p}$ for $p > 0$ and $h(d) = -\log |d|$.

Now, consider the modified convex optimization model:

$$\min_{x,y} \quad \sum_{i \in I} w_i n_i F(y_i) + \varepsilon \sum_{\substack{i \in I,\, j \in I \\ j < i}} h(y_i - y_j) \tag{3a}$$

$$\text{s.t.} \quad f_i + \frac{1}{n_i} \sum_{k \in K_i} x_{ik} = y_i, \qquad \forall i \in I \tag{3b}$$

$$y_i \in [0, 1], \qquad \forall i \in I \tag{3c}$$

$$x \in \mathbb{R}^+, \quad \sum_{i \in I_k} x_{ik} \leq b_k, \qquad \forall k \in K. \tag{3d}$$

Since the augmented objective in (3a) penalizes identical coverage levels, any solution with $y_i = y_j$ for $i \neq j$ becomes suboptimal. Thus, model (3) admits only distinct solutions, in a strictly stronger sense than Definition 3, without requiring further assumptions on the eligibility or resource structure. Further, the deterrent term $\Psi_\varepsilon(y)$ grows rapidly as any pair $y_i \to y_j$, producing a *sharp minimum* in the sense of Polyak [16, Chapter 5]. This sharpness implies not only uniqueness of the optimal coverage vector y in model (3), but also stability under small convex perturbations of the objective [16, Theorem 6, Chapter 5]; this feature is known as the superstability property. In the next section, we show via numerical experiments that the perturbation serves as a consistent and interpretable tie-breaking mechanism that aligns with the underlying fairness structure.

4. Case Study

As a numerical case study, we analyze the vaccine distribution model presented in [7]. This model considers the 254 counties in Texas with five priority groups — (i) 0–3 year olds, (ii) 4–24 year olds, (iii) 25–64 year olds at high risk, (iv) pregnant women, and (v) infant caregivers — competing for four vaccine types — (i) pre-filled syringe for infants (PFS baby), (ii) standard pre-filled syringe (PFS), (iii) multi-dose vial (MDV), and (iv) live attenuated influenza vaccine (LAIV). In our notation, the five priority groups correspond to five users $i = 1, \ldots, 5$, while the four vaccine types correspond to four resource types $k = 1, \ldots, 4$. Each user is eligible to receive only a subset of the available vaccines. The suitability matrix, $K_i \subseteq K$, defining these eligibility sets is given in Table 1.

Table 1: Eligibility matrix, K_i, from [7]. An entry (i, k) of 1 denotes that user i is eligible for resource k.

User/Resource		$k = 1$	$k = 2$	$k = 3$	$k = 4$
		PFS baby	PFS	MDV	LAIV
$i = 1$	0–3 years	1	0	0	0
$i = 2$	4–24 years	0	1	1	1
$i = 3$	25–64 (high-risk)	0	1	1	0
$i = 4$	Pregnant women	0	1	1	0
$i = 5$	Infant caregivers	0	1	1	1

[7] consider individual county level data, however for the purposes of this work we aggregate the data to the state-level to define the input parameters of model (1). Table 2 presents the data we use.

Table 2: Aggregated user and resource data for the baseline instance

(a) User data

User i	n_i	f_i
1	1,568,427	0.1549
2	7,632,499	0.7158
3	3,276,939	0.5219
4	342,432	0.5378
5	681,930	0.7266

(b) Resource data

Resource k	b_k
1	17,711
2	100,365
3	354,228
4	118,076

We first consider an instance with uniform weights: $w_i = 1$ for all $i \in I$, as assumed in the original setting of [7]. Solving model (1) yields the final coverage levels

$$y^{F,\text{eq}} = \Big[0.166,\ 0.731,\ 0.649,\ 0.649,\ 0.731\Big].$$

Ties occur between users $i = 2$ and $i = 5$, and between $i = 3$ and $i = 4$. Solving the perturbed model (3) with penalty $h(d) = 1/d^4$ eliminates all ties, yielding

$$y^{P,\text{eq}} = \Big[0.166,\ 0.729,\ 0.647,\ 0.667,\ 0.747\Big].$$

Notably, such ties may persist even when weights are unequal. Consider, for example, the following weight vector:

$$\bar{w} = \Big[0.767,\ 0.702,\ 0.806,\ 1.079,\ 0.677\Big],$$

constructed via a simple randomized search. Solving model (1) with these weights and the same inputs yields

$$y^{F,\bar{w}} = \Big[0.166,\ 0.731,\ 0.640,\ 0.731,\ 0.727\Big],$$

with a tie between users $i = 2$ and $i = 4$ despite strict differences in weights, eligibility sets, and initial coverage levels. This is a SUP outcome of type B in the taxonomy of Proposition 1. Specifically, user 2 receives no doses of resource types 2 or 3 (i.e., $x_{22} = x_{23} = 0$), while user 4 receives positive allocations of both ($x_{42}, x_{43} > 0$). Solving the perturbed model (3) with the same data yields

$$y^{P,\bar{w}} = \Big[0.166,\ 0.729,\ 0.647,\ 0.667,\ 0.747\Big],$$

in which all coverage levels are distinct.

The perturbed model first preserves the original loss function $L(y) = \sum_i w_i n_i (1 - y_i)^2$. Then, as a secondary step, it breaks ties by penalizing equality. To minimize the impact on L, the optimizer perturbs users with the smallest marginal cost $\frac{\partial L}{\partial y_i} = -2 w_i n_i (1 - y_i^*)$. In this case, $|\partial L / \partial y_4|$ is significantly smaller than $|\partial L / \partial y_2|$, so the optimizer lowers y_4 while keeping y_2 essentially unchanged. The penalty function $h(\cdot)$ determines only how far the two values separate; the direction of movement is governed entirely by L. Indeed, with $h(d) = d^4$, the gap is modest ($y_2 - y_4 = 0.062$), while with $h(d) = d^8$, the separation becomes more pronounced ($y_2 - y_4 = 0.178$).

As discussed in Section 2, SUP outcomes become more pronounced when resources are less scarce. To illustrate, we solve a version of the baseline instance in which each resource supply is scaled by a factor of five. The unperturbed model then yields

$$y^{F,\text{scaled}} = \Big[0.211,\ 0.898,\ 0.898,\ 0.898,\ 0.898\Big],$$

assigning identical coverage to four distinct users. The first user is only eligible for a non-shared resource and receives a unique coverage. Solving the perturbed model (3), with $h(d) = \frac{1}{d^4}$ yields

$$y^{P,\text{scaled}} = \Big[0.211,\ 0.905,\ 0.880,\ 0.924,\ 0.892\Big],$$

restoring distinctness without significantly altering the objective value. This demonstrates that SUP outcomes can arise even when resources are not so abundant as to fully satisfy all users. In this instance, the total supply is insufficient to assign every user full coverage ($y = 1$), yet all users who share resources receive exactly the same final coverage under the unperturbed model. The perturbation mechanism corrects this by restoring distinctness without significantly altering the allocation.

5. Conclusions

Systematic tie-breaking mechanisms in optimization-based allocation are relevant in high-stakes settings where transparency, interpretability, and adherence to prioritization principles are essential. In this work, we studied the occurrence of non-distinct outcomes in convex resource allocation problems, using a vaccine distribution model as a motivating case study. We demonstrated that such outcomes — characterized by ties in final coverages across users — can naturally arise even in realistic, policy-relevant instances under both uniform and non-uniform weighting schemes.

Further, our theoretical analysis showed that strict unclear priority (SUP) outcomes can persist not only in severely resource-constrained settings, but also in moderately resourced regimes where the total supply is still insufficient to satisfy all users fully. In such cases, ties emerge not due to a relaxation of the constraints, but rather from overlapping eligibility sets and the structure of the convex loss function.

To resolve these ambiguities, we revisited a classical idea: small perturbations to the objective function. We showed that a separable deterrent term, designed to penalize equality in coverage levels, ensures that the resulting solution is distinct in a strong sense. The direction of this perturbation aligns with the gradient of the original objective, thereby respecting the prioritization already embedded in the model. While such perturbations have a long history in the optimization literature (see also [14]), we are unaware of their application to fairness-constrained convex allocation models.

Finally, our numerical experiments demonstrated that this approach, despite its simplicity, consistently yields distinct and interpretable solutions with negligible impact on the original objective. The method is thus suitable as a lightweight, principled enhancement to existing allocation models when interpretability and distinctness are critical.

Acknowledgements

The research was funded by the Deutsche Forschungsgemeinschaft (DFG) project number 490948220.

References

[1] Atila Abdulkadiroğlu and Aram Grigoryan. Priority-based assignment with reserves and quotas. Working Paper 28689, April 2021. URL http://www.nber.org/papers/w28689.

[2] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, February 2011. ISSN 1526-5463. doi: 10.1287/opre.1100.0865.

[3] Gabriel R. Bitran and Arnoldo C. Hax. Disaggregation and resource allocation using convex knapsack problems with bounded variables. *Management Science*, 27(4):431–441, April 1981. ISSN 1526-5501. doi: 10.1287/mnsc.27.4.431.

[4] Sam Corbett-Davies, Johann D. Gaebler, Hamed Nilforoshan, Ravi Shroff, and Sharad Goel. The measure and mismeasure of fairness. *J. Mach. Learn. Res.*, 24(1), January 2023. ISSN 1532-4435.

[5] Mehmet C. Demirci, Andrew J. Schaefer, H. Edwin Romeijn, and Mark S. Roberts. An exact method for balancing efficiency and equity in the liver allocation hierarchy. *INFORMS Journal on Computing*, 24 (2):260–275, May 2012. ISSN 1526-5528. doi: 10.1287/ijoc.1110.0445.

[6] Larry G. Epstein and Uzi Segal. Quadratic social welfare functions. *Journal of Political Economy*, 100(4):691–712, August 1992. ISSN 1537-534X. doi: 10.1086/261836.

[7] Hsin-Chan Huang, Bismark Singh, David P. Morton, Gregory P. Johnson, Bruce Clements, and Lauren Ancel Meyers. Equalizing access to pandemic influenza vaccines through optimal allocation to public health distribution points. *PLOS ONE*, 12(8):e0182720, August 2017. doi: 10.1371/journal.pone.0182720.

[8] Anastasiya Ivanova, Pavel Dvurechensky, Alexander Gasnikov, and Dmitry Kamzolov. Composite optimization for the resource allocation problem. *Optimization Methods and Software*, 36(4):720–754, February 2020. ISSN 1029-4937. doi: 10.1080/10556788.2020.1712599.

[9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, March 1998. ISSN 1476-9360. doi: 10.1057/palgrave.jors.2600523.

[10] Hanan Luss. On equitable resource allocation problems: A lexico-

graphic minimax approach. *Operations Research*, 47(3):361–378, June 1999. ISSN 1526-5463. doi: 10.1287/opre.47.3.361.

[11] Honorata Mazepus and Florian van Leeuwen. Fairness matters when responding to disasters: An experimental study of government legitimacy. *Governance*, 33(3):621–637, August 2019. ISSN 1468-0491. doi: 10.1111/gove.12440.

[12] Nhan-Tam Nguyen, Trung Thanh Nguyen, Magnus Roos, and Jörg Rothe. Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multi-Agent Systems*, 28(2):256–289, April 2013. ISSN 1573-7454. doi: 10.1007/s10458-013-9224-2.

[13] Neal Parikh and Stephen Boyd. *Proximal algorithms*, volume 1. Now Publishers Inc, 2014. ISBN 9781601987174. doi: 10.1561/9781601987174.

[14] Alexander Plakhov. Local properties of the surface measure of convex bodies. *Journal of Convex Analysis*, 26(4):1373–1402, 2019.

[15] Nicholas G. Polson, James G. Scott, and Brandon T. Willard. Proximal algorithms in statistics and machine learning. *Statistical Science*, 30(4), November 2015. ISSN 0883-4237. doi: 10.1214/15-sts530.

[16] Boris T. Polyak. *Introduction to optimization.* Translations series in mathematics and Engineering. Optimization Software, New York, 1987. ISBN 0911575146.

[17] R. Tyrrell Rockafellar. Monotone operators and the proximal point

algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, August 1976. ISSN 1095-7138. doi: 10.1137/0314056.

[18] Bismark Singh. *Optimal spatiotemporal resource allocation in public health and renewable energy*. phdthesis, The University of Texas at Austin, December 2016. URL https://repositories.lib.utexas.edu/handle/2152/44589.

[19] Bismark Singh. Fairness criteria for allocating scarce resources. *Optimization Letters*, 14(6):1533–1541, March 2020. doi: 10.1007/s11590-020-01568-1.

[20] Jean-Paul Watson and David L. Woodruff. Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. *Computational Management Science*, 8(4):355–370, July 2010. ISSN 1619-6988. doi: 10.1007/s10287-010-0125-4.

[21] Ming Yi and Achla Marathe. Fairness versus efficiency of vaccine allocation strategies. *Value in Health*, 18(2):278–283, March 2015. ISSN 1098-3015. doi: 10.1016/j.jval.2014.11.009.