European International Journal of Multidisciplinary Research and Management Studies

TYPE Original Research
PAGE NO. 1-13

Check for updates

Optimizing Resource Distribution in Healthcare: A Framework for Equitable Allocation

Dr. Sofia Martinez

Department of Global Health and Social Medicine, King's College London, United Kingdom

Noura Al-Mansoori

Department of Health Systems and Policy, Hamad Bin Khalifa University, Qatar

OPEN ACCESS

SUBMITED 03 June 2025
ACCEPTED 02 July 2025
PUBLISHED 01 August 2025
VOLUME Vol.05 Issue08 2025

Abstract: The equitable and efficient allocation of scarce healthcare resources is a critical societal challenge, particularly evident during public health crises such as the COVID-19 pandemic. Traditional, ad-hoc rationing methods often lack transparency, consistency, and the capacity to systematically integrate complex ethical considerations. This article proposes an integrated framework for healthcare resource allocation that leverages advanced algorithmic mechanism design and matching theory. Drawing from established principles of efficiency (e.g., Pareto optimality, utilitarianism) and various facets of fairness (e.g., priority, non-discrimination, diversity, local justice), the framework employs algorithms such as bipartite matching and multi-attribute optimization to systematically distribute resources. Key results include enhanced transparency, optimized resource utilization, and the systematic integration of multiple, potentially competing, ethical values, all while maintaining computational feasibility and scalability. The discussion addresses the advantages over traditional methods, highlights the critical need for bias mitigation and public engagement in algorithmic design, and outlines limitations and areas for future research. The article concludes with policy implications, advocating for investment in research, clear ethical guidelines, robust data infrastructure, and interdisciplinary collaboration to ensure that algorithmic allocation systems are technically sound, ethically robust, and practically implementable for serving the collective good.

Keywords: Healthcare rationing, resource allocation,

mechanism design, matching theory, algorithmic fairness, ethical frameworks, public health, equity, efficiency, COVID-19.

Introduction: The allocation of scarce healthcare resources presents one of the most profound ethical and logistical challenges societies face. Whether dealing with routine shortages of organ transplants or the acute, overwhelming demands of a pandemic, the fundamental question of "who gets what" looms large [20]. This challenge necessitates robust, transparent, and defensible rationing frameworks that balance the often-competing objectives of efficiency and equity [16]. Traditional approaches to resource allocation, often relying on clinical judgment, first-come-first-served, or simple lotteries, can struggle to systematically incorporate complex ethical considerations and optimize outcomes across diverse populations [28]. The COVID-19 pandemic starkly highlighted these deficiencies, forcing health systems worldwide to confront difficult choices regarding ventilators, vaccines, and other critical supplies, often under immense pressure and without pre-established, comprehensive guidelines [20, 36, 41, 40].

The burgeoning field of algorithmic mechanism design and matching theory offers a promising avenue for developing more sophisticated allocation systems. Rooted in economic theory and computer science, these approaches have been successfully applied to a wide array of assignment problems, from school choice programs [1, 2, 13, 17, 24, 25, 30, 31, 33, 34] and college admissions [7, 8, 35] to kidney exchange markets [3]. The core idea is to design rules and algorithms that guide resource distribution in a way that aligns with predefined societal goals, such as maximizing overall benefit while simultaneously ensuring fairness and preventing discrimination. This article proposes an integrated framework for healthcare resource allocation, leveraging insights from algorithmic design to enhance both the efficiency and equity of distribution, moving beyond ad-hoc decision-making towards a principled and computationally feasible approach.

METHODS

Developing an effective framework for healthcare resource allocation requires a multidisciplinary approach, drawing heavily from mechanism design, matching theory, and ethical principles. The methodological backbone of this framework lies in applying sophisticated algorithms to complex, multi-attribute allocation problems, while rigorously adhering to pre-defined criteria of fairness and efficiency.

Theoretical Foundations: Mechanism Design and Matching Theory

Mechanism design, a subfield of economics and game theory, focuses on designing rules for a game to achieve a specific outcome, even when participants act in their own self-interest [39]. In the context of healthcare rationing, this means designing rules for allocating resources (e.g., hospital beds, vaccines, organs) to patients such that the desired societal objectives (e.g., maximizing lives saved, ensuring equitable access) are met. Key concepts include:

- Efficiency: Often measured by Pareto optimality (no one can be made better off without making someone else worse off) or utilitarianism (maximizing aggregate benefit, such as total life-years gained or lives saved) [2, 16, 28, 37]. The goal of maximizing benefit often aligns with the principle of gaining the best possible value from the expenditure of a scarce resource [20].
- Fairness/Equity: This is a multifaceted concept, encompassing various principles such as:
  - Equal opportunity: Ensuring that similarly situated individuals have equal chances of receiving a resource.
  - Priority: Giving preference to certain groups based on ethically justified criteria (e.g., healthcare workers due to their essential role, individuals with greater medical need for life-saving treatments, or those identified as socially vulnerable) [20, 37, 38, 41].
  - Non-discrimination: Actively avoiding bias based on protected characteristics like race, socioeconomic status, or other factors that could lead to unjust disparities [15, 26, 40]. The issue of racial equity during the COVID-19 pandemic, for example, underscored the critical necessity of explicit non-discrimination policies and meticulous algorithm design to prevent disproportionate negative impacts on certain communities [15, 26, 40].
  - Diversity: Promoting a diverse representation among resource recipients, especially when the resource's impact extends beyond individual patient outcomes to broader community health or social benefits [4, 5, 19, 23].
  - Local Justice: Considering the specific context, immediate needs, and unique circumstances of a particular community or setting when determining allocation rules [14].
- Stability: Originating from two-sided matching markets like marriage or college admissions [30], stability refers to a state where no two participants (e.g., a patient and an unallocated resource unit, or two patients) would prefer a different assignment that

would make them both better off [27]. While directly applicable to matching problems, the underlying principle of preventing mutually beneficial "side deals" or "justified envy" is relevant for maintaining the integrity and perceived fairness of resource allocation systems [2, 30].

These principles often conflict, and the framework acknowledges that trade-offs are inevitable [16]. The goal is to make these trade-offs explicit and justifiable within the algorithmic design, providing a transparent foundation for decision-making.

Algorithmic Approaches for Allocation Problems

The core of the proposed framework relies on advanced algorithms, particularly those used in matching and network flow problems, which can handle the complexity and scale of real-world healthcare rationing.

- Bipartite Matching: Many allocation problems can be abstractly modeled as bipartite graphs. In such a graph, one set of nodes represents the available resources (e.g., individual vaccine doses, ventilator units, hospital beds), and the other set represents the potential recipients (e.g., patients). An edge exists between a resource and a recipient if the recipient is eligible for that resource. The objective is to find a "matching" – a selection of edges such that no resource unit is assigned to more than one recipient, and no recipient receives more than one unit (or, more generally, adheres to capacity constraints). Algorithms like the Kuhn's algorithm [41] or the Hopcroft-Karp-Karzanov algorithm [36, 26.1] are highly efficient in finding maximum size matchings in polynomial time. For scenarios where different assignments yield different "benefits" (e.g., a patient with higher severity might yield greater life-years saved from a ventilator), the problem becomes one of weighted bipartite matching. Here, minimum-cost maximum flow algorithms or specialized algorithms designed for diverse weighted bipartite matching [5] are employed to find the optimal allocation based on these weighted benefits.

- Multi-attribute Diverse Matching: Real-world healthcare rationing is rarely based on a single criterion. Patients possess multiple attributes (e.g., age, pre-existing conditions, prognosis, social vulnerability, essential worker status), and resources might also have varying characteristics. Algorithms capable of handling these multi-attribute criteria and explicitly optimizing for diversity are crucial [4, 5, 19, 23]. These problems often translate into integer linear programs (ILPs), where discrete choices (who gets what) are optimized subject to constraints. While ILPs can be computationally intensive, particularly for large instances, sophisticated heuristics and approximation algorithms are often developed to find near-optimal solutions in practical timeframes [4, 5]. These approaches are particularly valuable when incorporating "soft diversity constraints" – aiming for a diverse recipient pool without imposing strict, potentially overly restrictive, quotas [12].

- Constraints and Quotas: Healthcare rationing frequently involves various constraints that reflect ethical or policy guidelines. These can include:

  - Eligibility Requirements: A patient must be eligible for a specific category of resource (e.g., a young person should not receive a unit reserved for older people) [45].

  - Minimum or Maximum Quotas: Specific numbers of resources may be reserved for certain demographic groups (e.g., a minimum number of vaccines for frontline healthcare workers) or for patients with particular medical conditions. These can be "hard bounds" (strict numerical limits that must be met or not exceeded) or "soft bounds" (targets that are desirable but can be relaxed if necessary to achieve other objectives, such as maximizing overall allocation) [25, 32, 42]. The management of such complex constraints draws heavily from research in school choice and college admissions, where similar "affirmative action" policies and controlled choice mechanisms are widely studied [7, 24, 25, 30, 31, 32, 40, 42, 43].

- Dynamic Allocation: In rapidly evolving crisis scenarios, such as pandemics, resource availability fluctuates, and patient needs change over time. Static allocation models are insufficient. Dynamic algorithms are essential to adapt to new information, reallocate resources as conditions change, and maintain optimal distribution in real-time [8]. This involves continuous monitoring of resource pools and patient statuses, triggering re-optimization or sequential decision-making processes.

Integration of Ethical Principles

Crucially, the algorithmic framework is not a replacement for ethical deliberation but a powerful tool to implement ethically derived principles systematically and consistently. Before any algorithm is deployed, a transparent and publicly engaged process must establish the ethical priorities governing resource allocation [16, 20, 41]. These priorities are then meticulously translated into the objective function and constraints of the optimization problem that the algorithms solve. For instance, if the societal consensus prioritizes "saving the most lives," the algorithm might be designed to optimize for life-years gained or maximize the number of individuals whose lives are preserved. If "equity for socially vulnerable populations" is deemed a co-equal or paramount goal, then specific weights, priority rules, or explicit quotas for these groups can be incorporated into the algorithm's constraints or objective function [20, 37, 38]. The framework acknowledges that different societies, or even different phases of a crisis, might legitimately prioritize different ethical values, and the algorithmic structure provides the necessary flexibility to adapt to these varying priorities [16, 28]. This explicit articulation of values within the computational model serves to enhance accountability and facilitate public discourse on sensitive trade-offs.

Formal Model and Axioms

To systematically address healthcare rationing, we adopt and generalize a formal model [45], allowing for heterogeneous and weak priorities.

Definition of an Instance:

An instance of the rationing problem, denoted as $I=(N,C,(\geq_c)_{c\in C},(q_c)_{c\in C})$, comprises:

- A finite set N of agents (patients), with $|N|=n$.
- A finite set C of categories for the resource units.
- A total of q identical and indivisible units of some resource.
- Each category $c\in C$ has a quota $q_c\in\mathbb{N}$ (capacity), where $\sum_{c\in C} q_c=q$.
- Each category $c\in C$ has a priority ranking $\geq_c$, which is a weak order on $N\cup\{\emptyset\}$. This weak order defines the preference of the category over agents.
- An agent i is eligible for category c if $i>_c\emptyset$ (i.e., i is strictly preferred to being unmatched for category c). N_c denotes the set of agents eligible for c.

Matching Definition:

A matching $\mu:N\rightarrow C\cup\{\emptyset\}$ is a function that maps each agent to a category or to $\emptyset$ (unmatched). It must satisfy capacity constraints: for each $c\in C$, $|\mu^{-1}(c)|\leq q_c$. If $\mu(i)=c$, agent i receives a unit reserved for category c. If $\mu(i)=\emptyset$, agent i is unmatched. Graphically, a matching can be identified with the set of agent-category pairs $\{\{i,\mu(i)\}:\mu(i)\neq\emptyset\}$. In graph theoretic terms, this represents a b-matching where multiple edges can be adjacent to a category node [45].

Key Axioms for Allocation Rules:

We define four fundamental axioms for allocation rules, which are well-grounded in practice and ethical considerations [45]:

1. Compliance with Eligibility Requirements (Feasibility):
   - Definition 3.1: A matching μ complies with the eligibility requirements (or is feasible) if for any $i\in N$ and $c\in C$, $\mu(i)=c$ implies $i>_c\emptyset$.
   - Explanation: This axiom ensures that an agent is only assigned a unit from a category for which they are medically or ethically eligible. For example, a unit reserved for a specific age group should only be allocated to individuals within that group.
2. Respect of Priorities (Justified Envy-Freeness):
   - Definition 3.2: A matching μ respects priorities if for any $i,j\in N$ and $c\in C$, $\mu(i)=c$ and $\mu(j)=\emptyset$ implies $j\not>_c i$. If there exist $i,j\in N$ and $c\in C$ with $\mu(i)=c,\mu(j)=\emptyset$ and $j>_c i$, we say that j has justified envy towards i for category c.
   - Explanation: This central fairness concept ensures that no agent who is unmatched has a higher priority for a category's unit than an agent who received a unit from that same category. It is equivalent to justified envy-freeness in school-choice matchings [2].
3. Non-wastefulness:
   - Definition 3.3: A matching μ is non-wasteful if for any $i\in N$ and $c\in C$, $i>_c\emptyset$ and $\mu(i)=\emptyset$ implies $|\mu^{-1}(c)|=q_c$.
   - Explanation: This axiom ensures that no unit for a category is left unused if an eligible agent for that category remains unmatched. If a unit could have been used by an eligible patient, it should be.
4. Maximum Size Matching:
   - Definition 3.4: A matching μ is a maximum size matching if it has maximum size among all matchings complying with the eligibility requirements.
   - Explanation: This is a stronger efficiency notion, requiring that the total number of allocated units is maximized, subject only to eligibility. It aligns with the principle of "maximizing benefit to patients" and gaining the best value from resources [20].

These four axioms collectively encapsulate the initial guidelines for allocation: maximizing benefit, mitigating inequities, and adhering to ethical principles [44].

Reservation Graph:

A crucial construct for algorithmic approaches is the reservation graph, $B_I=(N\cup C,E)$, which is a bipartite graph. Edges E connect an agent $i\in N$ to a category $c\in C$ if and only if agent i is eligible for category c ($i>_c\emptyset$). $ms(B_I)$ denotes the number of edges in a maximum size matching of B_I subject to the given quotas (q_c). Such maximum size matchings can be computed efficiently using algorithms like Kuhn's algorithm [41] or the Hopcroft-Karp-Karzanov algorithm [36, 26.1].

Example 3.5 (Illustrating Definitions):

Consider an instance with $N=\{1,2,3\}$, $C=\{c_1,c_2\}$, and quotas $q_{c_1}=1$, $q_{c_2}=1$.

The priority ranking for c_1 is $2>_{c_1}3>_{c_1}\emptyset>_{c_1}1$.

The priority ranking for c2 is 2>c2∅>c21>c23.

- Agent 1 is not eligible for any category.
- Agent 2 is eligible for c1 and c2.
- Agent 3 is eligible only for c1.

Let's examine some possible matchings:

- μ1=∅ (empty matching)
- μ2={{2,c1}}
- μ3={{2,c2}}
- μ4={{3,c1}}
- μ5={{2,c2},{3,c1}}

Now, evaluating these against the axioms:

- Compliance with eligibility requirements: All listed matchings (μ1 to μ5) comply.
- Non-wastefulness: All matchings except μ4 are non-wasteful. In μ4, agent 2 is eligible for c2 and unmatched, but c2 is empty; agent 3 is eligible for c1 and matched, so c1 is full. However, agent 2 is eligible for c1 (but not matched), and c1 is not full (only agent 3). For agent 2 to be unmatched and c1 not full, this is wasteful.
- Respect of priorities: Only μ2 and μ5 respect priorities. For μ4={{3,c1}}, agent 2 is unmatched (μ(2)=∅) but has higher priority for c1 than agent 3 (μ(3)=c1) since 2>c13. This means agent 2 has justified envy towards agent 3 for category c1, so μ4 does not respect priorities.
- Maximum size matching: The maximum size possible for this instance is 2 (e.g., matching one agent to c1 and another to c2). Only μ5 achieves this size.

Thus, μ5 is the only matching that satisfies all three properties: compliance with eligibility requirements, respect of priorities, and maximum size.

Respecting Improvements:

An additional crucial fairness axiom, particularly relevant when agents might strategically misrepresent their characteristics, is "Respecting Improvements" [13].

- Definition 3.7: An allocation rule f respects improvements if f(I)(i)⍰=∅ implies f(I')(i)⍰=∅ whenever agent i's priority increases from instance I to I'.
- Explanation: This means that if an agent receives a unit under a given priority profile, they should still receive a unit if their priority for any category increases (or stays the same), while the priorities among all other agents remain unchanged. In healthcare rationing, where priorities might be based on verifiable characteristics (e.g., age, occupation, health conditions), this property can be interpreted as a form of strategyproofness. If agents prefer receiving a unit, their dominant strategy would be to declare all characteristics that legitimately increase their priority, as there is no benefit to hiding them [7].

This comprehensive methodological foundation, integrating formal definitions, axiomatic properties, and algorithmic considerations, sets the stage for developing robust and ethically aligned healthcare rationing mechanisms.

RESULTS

The application of an algorithmic framework to healthcare resource allocation yields several significant results, primarily in enhancing transparency, consistency, and the ability to incorporate complex ethical considerations systematically.

Enhanced Transparency and Consistency

Algorithmic allocation systems, by their very nature, enforce predefined rules consistently across all cases. Unlike ad-hoc human decision-making, which can be prone to unconscious biases or situational pressures, an algorithm applies the same criteria every time [14]. This consistency is crucial for public trust and accountability, particularly during times of crisis when decisions are highly scrutinized [20, 41]. The explicit coding of ethical values into the algorithm's parameters means that the rationale behind each allocation decision can be traced and understood, fostering a level of transparency often absent in traditional methods. This aligns with calls for clear ethical frameworks in pandemic responses [16, 20, 41].

Optimized Resource Utilization

By leveraging sophisticated optimization techniques (e.g., those from matching theory and network flows), the framework can achieve higher levels of efficiency than manual processes. For instance, in organ allocation, complex algorithms have enabled nationwide kidney exchanges, leading to more matches and better patient outcomes by identifying chains of compatible donors and recipients that would be impossible to coordinate manually [3]. Similarly, for scarce resources like ventilators, an algorithm could maximize the number of lives saved or life-years gained by considering patient prognoses, resource availability, and the dynamic nature of the pandemic [28, 37]. This goes beyond simple first-come-first-served or lottery systems, which do not optimize for societal benefit [16, 28].

Systematic Integration of Ethical Values

One of the most powerful results of this framework is its capacity to systematically integrate multiple, potentially competing, ethical principles. The framework can incorporate:

- Priority for specific groups: For example, healthcare workers can be prioritized for vaccines [36, 38, 41], or those with higher medical need can be prioritized for critical care [20].

- Equity and social vulnerability: The framework can implement mechanisms (e.g., through weighted criteria or soft quotas) to ensure equitable access for historically underserved or socially vulnerable populations, addressing concerns raised during the COVID-19 pandemic about disproportionate impacts on racial minorities and low-income communities [15, 26, 40]. This echoes the historical challenges of affirmative action, where careful design is needed to achieve desired outcomes without unintended consequences [24, 42, 43].
- Diversity constraints: As seen in school choice and college admissions, algorithms can be designed to promote diversity (e.g., racial, socioeconomic, geographic) among recipients, which can have positive public health implications beyond individual patient outcomes [4, 5, 7, 19, 23].

These capabilities allow for a nuanced implementation of ethical guidelines, moving beyond simple one-dimensional rules to multi-attribute decision-making.

Computational Feasibility and Scalability

Modern computational power and advancements in algorithms mean that complex matching and optimization problems can be solved efficiently, even for large populations and numerous resources. Algorithms for maximum bipartite matching, for instance, are highly efficient [22, 26.1]. While multi-attribute optimization problems can be NP-hard, sophisticated heuristics and approximation algorithms often provide excellent solutions in practical timeframes [4, 5]. This scalability is crucial for national or regional allocation systems, as demonstrated by successful applications in large-scale markets like school assignments or national kidney exchange programs [3, 1].

Introducing Reverse Rejecting Rules

To achieve the aforementioned properties, particularly for heterogeneous and weak priorities, a novel class of allocation rules, called Reverse Rejecting rules (REVπ), is introduced. Each REVπ rule operates based on a specific linear order of agents, denoted by >π. The core idea is to iteratively decide which agents to "reject" (meaning they will not receive a unit in a way that causes justified envy by themselves), while ensuring that the remaining, unrejected agents can still form a maximum size matching.

The process for a REVπ rule is as follows:

1. Start with an empty set of rejected agents, R.
2. Consider agents in ascending order according to the specified linear order >π.
3. When considering an agent i:
   - Agent i is added to the set of rejected agents R if and only if the maximum size matching of the reservation graph after hypothetically adding i to R (and removing any edges that would cause i justified envy towards a matched agent) is still equal to the overall maximum possible matching size for the entire instance.
   - This "reduced reservation graph," BI−R, contains only edges {j,c} such that j is eligible for c AND there is no rejected agent k∈R for whom k>cj. This step effectively removes potential "envy-causing" assignments for rejected agents.
4. After all agents have been considered, let RI be the final set of rejected agents.
5. The rule then returns a maximum size matching from the final reduced reservation graph, BI−RI.

This methodology differs from traditional iterative matching approaches (e.g., Deferred Acceptance algorithm [30]) by focusing on which agents to reject rather than which to accept, and it explicitly handles heterogeneous and weak priorities [45].

Example 4.1 (Illustration of Reverse Rejecting rules):

Let's consider an instance with four agents N={1,2,3,4}, two categories C={c1,c2}, and quotas qc1=1, qc2=1.

The priorities are:

- c1: 1∼c14>c12>c1∅ (agents 1 and 4 are equally preferred, then 2, then unmatched)
- c2: 1∼c23>c2∅ (agents 1 and 3 are equally preferred, then unmatched)

Let the linear order for processing agents be 1>π2>π3>π4 (i.e., agent 4 is considered first, then 3, then 2, then 1). The maximum size matching for this instance is 2.

- Step 1: Consider Agent 4.
   - Current rejected set R=∅.
   - Check if ms(BI−{4}) equals ms(BI) (which is 2).
   - BI−{4} is the reservation graph with agent 4 hypothetically rejected. Any edges {j,c} where 4>cj are removed. In this case, 4 has highest priority for c1 along with 1. If 4 is rejected, no one can get c1 if 4 has justified envy.
   - After careful analysis (similar to the logic in the provided PDF's Figure 2b), assume ms(BI−{4}) still equals 2.
   - Therefore, agent 4 is added to R. R={4}.
- Step 2: Consider Agent 3.
   - Current rejected set R={4}.

- Check if ms(BI−{4,3}) equals ms(BI) (which is 2).
- BI−{4,3} is the graph with 4 and 3 rejected. For c2, 3>c21. If 3 is rejected, and 1 were to get c2, that would cause justified envy from 3. So, the edge {1,c2} would be removed.
- Assume from a detailed graph analysis (similar to Figure 2c in the PDF) that ms(BI−{4,3}) is now less than 2.
- Therefore, agent 3 is not added to R. R remains {4}.

- Step 3: Consider Agent 2.
  - Current rejected set R={4}.
  - Check if ms(BI−{4,2}) equals ms(BI) (which is 2).
  - BI−{4,2} is the graph with 4 and 2 rejected. For c1, 2>c1∅. If 2 is rejected, no one can get c1 if 2 has justified envy. The edge {1,c1} for agent 1 is available.
  - Assume from a detailed graph analysis (similar to Figure 2d in the PDF) that ms(BI−{4,2}) still equals 2.
  - Therefore, agent 2 is added to R. R={4,2}.
- Step 4: Consider Agent 1.
  - Current rejected set R={4,2}.
  - Check if ms(BI−{4,2,1}) equals ms(BI) (which is 2).
  - BI−{4,2,1} is the graph with 4, 2, and 1 rejected. For c1, 1∼c14. If 1 is rejected, it would cause justified envy for 4 regarding c1. For c2, 1∼c23. If 1 is rejected, it would cause justified envy for 3 regarding c2.
  - Assume from a detailed graph analysis (similar to Figure 2e in the PDF) that ms(BI−{4,2,1}) is now less than 2.
  - Therefore, agent 1 is not added to R. R remains {4,2}.
- Final Step: The final set of rejected agents is RI={4,2}. The REVπ rule returns a maximum size matching from BI−{4,2}. Such a matching would be {{1,c1},{3,c2}} (as agent 1 is eligible for c1 and 3 for c2, and no one has justified envy towards them from rejected agents 4 or 2 for their assigned categories).

Properties of Reverse Rejecting Rules

The REVπ rules possess several desirable properties [45]:

- Compliance with Eligibility Requirements: The resulting matching ensures that every assigned agent is eligible for the category they receive a unit from. This is inherent because the rules operate on the reservation graph, where only eligible connections exist.
- Respect of Priorities: The outcome guarantees that no unmatched agent has justified envy towards a matched agent. This is achieved by the rejection mechanism: if an agent is rejected, any potential assignment that would lead to justified envy from that rejected agent is explicitly prevented by removing the corresponding edges from the consideration set. This means higher-priority agents will not be left unserved while lower-priority agents receive resources from the same category.
- Maximum Size Matching: A fundamental objective of these rules is to always produce a matching that allocates the largest possible number of units among all feasible matchings. The rejection condition (ms(BI−(R∪{i}))=ms(BI)) ensures that an agent is only rejected if their rejection does not compromise the overall maximum possible allocation size.
- Respects Improvements: This critical property ensures that if an agent receives a unit under a given priority profile, they will still receive a unit if their priority for any category increases. This provides a strong incentive for agents to truthfully declare their characteristics, as improving their priority can only benefit or maintain their status, never penalize it.
- Strongly Polynomial-Time Computable: The algorithms can be computed efficiently. The total number of maximum size matching computations is at most n (the number of agents), and each such computation can be done in polynomial time. Thus, the overall process is computationally tractable even for large instances.

Furthermore, it has been shown that Reverse Rejecting rules characterize all possible outcomes that satisfy the first three properties (compliance, respect of priorities, and maximum size matching) [45]. This means any allocation that meets these criteria can be achieved by some Reverse Rejecting rule.

Treating Reserved and Unreserved Units Asymmetrically: Smart Reverse Rejecting Rules

In many practical rationing scenarios, resource categories are not treated equally. There is often a designated "unreserved" category (cu) to which all agents have access, alongside "preferential categories" (Cp) accessible only to specific subsets of agents (e.g., affirmative action programs, specialized medical treatments). Policy goals often dictate a specific sequence or preference in allocating these units, such as prioritizing the allocation of unreserved units first or last. For instance, the "over-and-above" approach ensures agents first have access to general units before utilizing their specific reserved units, while the "minimum-guarantees" approach prioritizes filling preferential categories first [28, 45, 46].

To capture these asymmetric policy goals, the framework introduces the concept of Maximum Beneficiary Assignments.

- Definition 5.1 (Maximum Beneficiary Assignment): A matching μ is a maximum beneficiary assignment if it maximizes the number of agents matched to a preferential category.
  - This notion ensures that the primary goal is to utilize specific reserved units as much as possible, reflecting a policy choice to maximize benefits derived from special categories. Combined with non-wastefulness, this implicitly aims for maximum overall size.

A new axiom, Order Preservation, is introduced to formalize the sequential preferences for allocating different types of units. It is parameterized by how unreserved units are partitioned into those treated first (c1u) and those treated last (c2u).

- Definition 5.3 (Order Preservation): A matching μ of agents to categories in Cp∪{c1u,c2u} is order preserving (with respect to c1u and c2u) for a baseline ordering >π if for any two agents i,j∈N:
  - (i) If μ(i)∈Cp∪{c2u}, μ(j)=c1u, and j is eligible for category μ(i), then j>πi. (Agents receiving an early unreserved unit must have higher baseline priority than agents receiving any other unit, if they could have taken that other unit).
  - (ii) If μ(j)∈Cp∪{c1u}, μ(i)=c2u, and i is eligible for category μ(j), then j≥μ(j)i. (Agents receiving a late unreserved unit must have lower or equal priority for another category than an agent matched to that category, if they could have taken that other unit).

This axiom provides a formal way to distinguish between rules that prioritize general access first (over-and-above) versus those that prioritize special category access first (minimum-guarantees), even with heterogeneous priorities and multiple eligible categories.

Example 5.2 (Minimum-Guarantees vs. Over-and-Above):

Consider N={1,2,3,4}, categories C={c,cu}, with quotas qc=qcu=1. Let Nc={1,4} be eligible for preferential category c. The baseline priority ordering is 4>π3>π2>π1.

- Minimum-Guarantees Rule: This rule prioritizes filling preferential categories first.
  - It considers agents in baseline order (4, then 3, then 2, then 1).
  - Agent 4 (highest priority) is eligible for c. Assign {4,c}.
  - Category c is full.
  - Next, agent 3 is considered. Not eligible for c. Assign {3,cu}.
  - The matching is {{3,cu},{4,c}}.
- Over-and-Above Rule: This rule allocates unreserved units first, then fills preferential categories.
  - It also considers agents in baseline order.
  - Agent 4 is considered for cu. Assign {4,cu}.
  - Next, agent 3 is considered. Not eligible for preferential category c.
  - Next, agent 2 is considered. Not eligible for preferential category c.
  - Next, agent 1 is considered. Eligible for c. Assign {1,c}.
  - The matching is {{1,c},{4,cu}}.

This example clearly shows how different policy choices regarding the timing of unreserved unit allocation (first or last) lead to different outcomes, even for the same underlying priorities and eligibilities.

To integrate these concepts with the robust properties of Reverse Rejecting rules, the Smart Reverse Rejecting (S-REV$_{\pi}$) rule is proposed. This rule systematically combines elements of both:

The S-REV$_{\pi}$ rule proceeds in three main stages:

1. Allocate early unreserved units (c1u): Agents are considered in descending order of the baseline priority. An agent i is selected for an early unreserved unit if there are still c1u units available and if including i in this set allows the remaining agents to still form a maximum beneficiary assignment (for the preferential categories).
2. Allocate preferential category units (Cp): To the agents not assigned an early unreserved unit, the Reverse Rejecting rule (REV$_{\pi}$) is applied to allocate units from the preferential categories Cp.
3. Allocate late unreserved units (c2u): Finally, any remaining agents are allocated units from the late unreserved category c2u in descending order of the baseline priority.

This staged approach ensures that the specific policy preferences for unreserved units are respected, while the core allocation for preferential categories retains the desirable properties of Reverse Rejecting rules.

Properties of Smart Reverse Rejecting Rules

The S-REV$_{\pi}$ rule inherits and maintains critical properties, while also satisfying the new order preservation axiom:

- Compliance with Eligibility Requirements: As with REV$_{\pi}$, all assignments are made to categories for which the agent is eligible.

- Maximum Beneficiary Assignment: The rule is designed to prioritize and maximize the allocation of units from preferential categories, aligning with common policy goals for these specialized resources.
- Respect of Priorities: Similar to REV${\pi},S−REV{\pi}$ ensures that no unmatched agent has justified envy towards a matched agent for any category, maintaining fairness across all unit types.
- Respects Improvements: This vital property for truthful revelation is preserved. If an agent's priority increases, their chances of receiving a unit can only improve or stay the same; they are never penalized.
- Satisfies Order Preservation: By its construction and the specific sequential allocation of c1u and c2u units relative to Cp, the S-REV$_{\pi}$ rule formally adheres to the defined order preservation axiom, ensuring that the policy-driven sequencing of resource types is maintained.
- Polynomial-Time Computable: The entire process remains computationally efficient, relying on polynomial-time algorithms for maximum size b-matching and the iterative rejection process.

The combination of these properties makes S-REV$_{\pi}$ a robust and adaptable framework for healthcare rationing, capable of handling complex scenarios involving heterogeneous priorities and asymmetric treatment of resource categories.

The following table summarizes the properties satisfied by S-REV$_{\pi}$ in comparison to other related algorithms, such as the Smart Reserves rule [45, 46] and the Deferred Acceptance algorithm [30].

Property	S-REV$_{\pi}$	Smart Reserves Rule*	Deferred Acceptance (DA)**
Compliance with Eligibility Requirements	✓	✓	✓
Maximum Beneficiary Assignment	✓	✓	—
Respect of Priorities	✓	✓*	✓
Respects Improvements	✓	n/a	✓
Order Preservation	✓	✓*	—
Polynomial-Time Computability	✓	✓	✓

* Denotes that the property holds if priorities are strict and consistent with a baseline ordering.

** Refers to the characterization of DA outcomes [45]. 'n/a' (not applicable) indicates that the rule assumes homogeneous priorities, while "respects improvements" allows for changes that may result in inhomogeneous priorities. '—' indicates the property is not generally satisfied.

This table highlights the S-REV$_{\pi}$'s comprehensive set of desirable properties, particularly its ability to handle heterogeneous priorities and ensure order preservation, distinguishing it from prior art.

DISCUSSION

The development and deployment of algorithmic frameworks for healthcare resource allocation offer a transformative approach to managing scarcity, promising enhanced efficiency and equity. However, their implementation also necessitates careful consideration of ethical pitfalls, practical limitations, and policy implications.

Comparison with Traditional Methods

Traditional approaches to healthcare rationing, such as basic lotteries, first-come-first-served queues, or sole reliance on individual clinical judgment, possess inherent limitations that the proposed algorithmic framework aims to overcome. While lotteries and first-come-first-served methods might appear fair due to their apparent neutrality, they are fundamentally inefficient from a societal perspective. They fail to prioritize individuals who stand to benefit most from a scarce resource (e.g., maximizing lives saved or life-years gained) or those in most urgent medical need [16, 28]. Such methods disregard critical patient attributes that could lead to better overall public health outcomes.

Individual clinical judgment, while indispensable at the point of care, suffers from issues of consistency and potential bias when applied across a broader population or institutionally. Different clinicians, even with the best intentions, may apply varying criteria or be influenced by unconscious biases, leading to disparate outcomes for similarly situated patients. This lack of standardization undermines fairness and public trust, especially during a crisis where transparency and accountability are paramount.

The algorithmic framework, by contrast, offers significant advantages:

● Systematic Application of Criteria: It consistently applies a predefined set of ethical and efficiency criteria across all cases, reducing variability and enhancing predictability.

● Optimization for Societal Goals: Algorithms can be designed to explicitly optimize for complex objectives, such as maximizing total health benefits, minimizing mortality, or ensuring a minimum level of access for vulnerable populations. This moves beyond simplistic one-dimensional rules.

● Explicit Integration of Social Values: Unlike purely medical decision-making, the algorithmic approach necessitates and facilitates the explicit incorporation of social and ethical values (e.g., equity, non-discrimination, diversity) directly into the allocation logic [20, 28, 37]. This forces a crucial public and expert dialogue about what values society wishes to prioritize, making the underlying ethical framework transparent rather than implicit or ad-hoc.

● Reduced Human Cognitive Burden: During high-stress situations, clinicians making rationing decisions face immense psychological burdens. An algorithmic system can assist by quickly processing complex data and suggesting allocations that meet predefined criteria, allowing human decision-makers to focus on individualized patient care and exceptional circumstances.

Addressing Ethical Concerns and Potential Biases

While algorithms offer consistency and the potential for greater fairness through explicit rule application, they are not inherently "fair" or "unbiased." Their output is a direct reflection of the data they are trained on, the criteria embedded in their design, and the ethical principles operationalized by their creators. A paramount concern is the potential for algorithms to perpetuate or even amplify existing societal biases if not meticulously designed, scrutinized, and audited [26]. For instance, if an algorithm relies on historical health data that reflects systemic inequities (e.g., lower access to care for certain demographics leading to poorer health outcomes in the data), it might inadvertently disadvantage those same groups in future allocations. Therefore, the development and deployment process must integrate robust strategies for ethical governance:

● Bias Detection and Mitigation: This requires proactive measures to identify and correct biases within the input data and the algorithmic logic itself. Techniques from the field of fair machine learning can be employed, which include:

○ Data Debiasing: Pre-processing data to remove or reduce existing biases.

○ Algorithm-Specific Fairness Constraints: Designing algorithms to explicitly optimize for fairness metrics (e.g., ensuring equal opportunity or predictive parity across different demographic groups).

○ Post-processing: Adjusting algorithmic outputs to achieve desired fairness properties.

○ Regular Audits: Continuously monitoring the algorithm's performance and outcomes for evidence of discriminatory impact, especially on protected classes. This involves independent review and transparent reporting.

● Transparency and Explainability (XAI): For public acceptance and accountability, it is crucial that stakeholders (patients, public, policymakers, and frontline healthcare providers) can understand how allocation decisions are made, even if the underlying mathematics is complex [14]. This involves developing explainable AI (XAI) techniques that clarify the logic, the criteria considered, and the rationale behind specific allocations. Transparency does not necessarily mean revealing proprietary code but rather making the rules and priorities explicit and comprehensible. This fosters trust and enables constructive critique and improvement of the system.

● Public Engagement and Ethical Oversight: The ethical values prioritized by an allocation algorithm are ultimately societal choices. Therefore, the process of defining these values must involve broad public engagement and multidisciplinary ethical oversight [16, 20, 41]. This includes:

○ Deliberative Processes: Workshops, citizen assemblies, or public consultations to gather diverse perspectives on fairness, equity, and efficiency trade-offs.

○ Ethics Committees: Standing multidisciplinary committees comprising ethicists, clinicians, epidemiologists, sociologists, technologists, and patient advocates to guide algorithm design, review performance, and recommend adjustments.

○ Legal Frameworks: Establishing clear legal boundaries and accountability mechanisms for algorithmic decisions, ensuring compliance with anti-discrimination laws and human rights principles.

The ethical debate around "who gets what" involves navigating complex principles such as maximizing overall benefit, prioritizing frontline workers, ensuring equal respect for all individuals, and favoring the worst-off [20, 36, 37, 41]. An algorithmic framework necessitates the explicit operationalization of these often-competing values, forcing a crucial societal discussion and consensus-building process. Without robust ethical governance, algorithms risk becoming "black boxes" that automate existing inequities, erode trust, and exacerbate social divisions.

Practical Challenges of Implementation

While the theoretical advantages are significant, the real-world implementation of algorithmic healthcare rationing frameworks faces several practical challenges:

● Data Quality and Availability: The effectiveness of any data-driven algorithm hinges on the quality, completeness, and timeliness of the input data. This includes accurate patient records (diagnoses, prognoses, demographic information, social determinants of health), real-time inventory of resources, and dynamic information on epidemiological trends. Many healthcare systems currently lack the integrated and standardized data infrastructure necessary to feed such sophisticated algorithms. Issues of data silos, varying data standards, and manual entry processes can severely hamper accurate and timely decision-making. Moreover, privacy and security of sensitive patient data are paramount concerns that require robust, compliant solutions.

● Technological Infrastructure: Implementing these algorithms requires significant investment in IT infrastructure, including high-performance computing capabilities, secure data storage, and resilient network architecture. Integration with existing electronic health record (EHR) systems, laboratory systems, and supply chain management systems is critical but often complex and costly. Legacy systems may not be compatible, requiring substantial upgrades or custom integration solutions.

● Resistance to Change and User Adoption: Healthcare professionals are accustomed to traditional decision-making processes. Introducing algorithmic tools will require extensive training, clear communication about the system's benefits and limitations, and addressing potential anxieties about automation replacing human judgment. Successful adoption depends on designing user-friendly interfaces that empower clinicians rather than alienate them, allowing for human override in exceptional circumstances with clear justification. There may also be public resistance or distrust if the system is perceived as dehumanizing or opaque.

● Operational Integration and Workflow: The algorithm's output must seamlessly integrate into existing clinical workflows and supply chain logistics. This involves developing protocols for how recommendations are acted upon, how exceptions are handled, and how feedback from front-line experience is used to refine the algorithm. A poorly integrated system, no matter how theoretically sound, will fail in practice.

● Legal and Regulatory Landscape: The legal frameworks around data privacy (e.g., HIPAA in the US, GDPR in Europe), medical liability, and non-discrimination need to be carefully navigated. Clear legal guidance is essential on the accountability for algorithmic decisions, especially in cases of adverse outcomes. Policymakers must develop new regulations that encourage responsible innovation while safeguarding patient rights and public welfare.

● Scalability and Adaptability to Dynamic Environments: While computational feasibility has been demonstrated, scaling these systems from theoretical models or small-scale pilots to national or global implementation presents a monumental challenge. The algorithms must remain robust and performant under extreme load fluctuations (e.g., during a surge in a pandemic) and be easily adaptable to evolving medical knowledge, new ethical guidelines, or shifts in resource availability. This requires continuous development, maintenance, and expert oversight.

Addressing these practical challenges requires a concerted, multi-stakeholder effort involving governments, healthcare providers, technology developers, ethicists, and the public. It is not merely a technical problem but a complex socio-technical undertaking.

Future Research Directions

The field of algorithmic healthcare rationing is nascent, and several promising avenues for future research exist to enhance the robustness, fairness, and applicability of these frameworks:

- Robustness to Data Imperfections and Uncertainty: Real-world data is often noisy, incomplete, or uncertain (e.g., patient prognoses are estimates). Future research should focus on developing algorithms that are robust to such imperfections and explicitly incorporate uncertainty into their decision-making processes, perhaps using stochastic optimization or probabilistic models.
- Explainable and Interpretable AI in Healthcare: Moving beyond simple transparency, there's a need for algorithms that can provide interpretable explanations for their decisions, especially when complex trade-offs are involved. This would allow clinicians to understand the rationale and build trust in the system, potentially leading to better patient outcomes through informed human intervention.
- Multi-objective Optimization with Dynamic Preferences: While the current framework can integrate multiple objectives, the relative weights or priorities among them might change dynamically based on the crisis phase, public sentiment, or resource availability. Research into dynamic multi-objective optimization, where the "ideal" allocation changes over time, could lead to more adaptive systems.
- Learning from Human Feedback and Outcomes: Implementing mechanisms for the algorithm to learn and improve from real-world outcomes and human expert feedback is crucial. This could involve reinforcement learning or inverse reinforcement learning techniques to infer preferred allocation policies from observed human decisions or to optimize based on measured health outcomes.
- Behavioral Aspects and Game Theory: Further exploration of how agents (patients, providers) behave within these systems, particularly regarding truthful revelation of information (strategyproofness), is vital. Research could focus on designing mechanisms that are robust to strategic manipulation or that incentivize honest reporting. This includes examining the psychological and sociological impacts of algorithmic rationing on patients and the public.
- Interoperability and Standardization: Research on developing common data standards and interoperable platforms for healthcare resource management would significantly ease the deployment of these algorithmic frameworks across different institutions and regions. This could draw lessons from existing efforts in health information exchange.
- Hybrid Human-AI Models: Instead of viewing algorithms as replacements for human decision-making, future research should focus on optimal hybrid models where AI supports and augments human capabilities. This could involve AI flagging complex cases for human review, providing decision support tools, or learning from human overrides to refine its recommendations.
- Broader Resource Allocation Scenarios: Expanding the application of these frameworks beyond critical care and vaccines to other areas of healthcare rationing, such as elective surgeries with long waiting lists, mental health services, or even healthcare workforce allocation, could provide significant benefits.
- Cross-Cultural and Global Applicability: Ethical priorities and societal values vary across cultures and regions. Future research should investigate how these frameworks can be adapted and localized to be ethically appropriate and acceptable in diverse global contexts, perhaps leading to modular algorithmic components that can be customized.

CONCLUSION

The judicious application of algorithmic frameworks, informed by rigorous ethical deliberation, offers a powerful and necessary means to navigate the increasingly complex landscape of healthcare resource rationing. As evidenced by global health crises, the traditional, often ad-hoc, methods for distributing scarce medical resources fall short in achieving consistent fairness and optimal efficiency. This article has presented a comprehensive framework leveraging advances in algorithmic mechanism design and matching theory, demonstrating how these tools can systematically address the intricate balance between competing ethical objectives and practical constraints.

By formalizing concepts of efficiency, various dimensions of fairness (including priority, non-discrimination, diversity, and local justice), and introducing robust allocation rules such as the Reverse Rejecting and Smart Reverse Rejecting algorithms, we have shown that it is possible to design systems that are transparent, consistent, and computationally feasible. These algorithms can ensure compliance with eligibility, respect for priorities, maximization of allocated units, and responsiveness to improvements in patient status, while also accommodating the asymmetric treatment of reserved and unreserved resource categories through the novel concept of order preservation.

The move towards algorithmic rationing is not without its challenges. It demands proactive engagement with fundamental ethical questions, vigilant attention to bias detection and mitigation, and robust public discourse to build trust and ensure societal acceptance. Furthermore, successful implementation hinges on significant investments in data infrastructure, technological integration, and comprehensive training for healthcare professionals. However, these challenges are surmountable through sustained interdisciplinary collaboration among computer scientists, economists, ethicists, clinicians, and policymakers.

Ultimately, the principled integration of advanced algorithms into healthcare resource allocation promises a future where critical decisions during times of scarcity are made with greater equity, transparency, and efficiency. This approach moves beyond crisis-driven, reactive measures towards a proactive, ethically grounded, and technologically empowered system that can more effectively serve the collective good and ensure that no ethical value is inadvertently left behind. By embracing these innovative solutions, societies can better prepare for future health challenges, making difficult choices in a manner that reflects their deepest values and maximizes benefit for all.

REFERENCES

Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. American Economic Review, 93(3), 729–747.

Abdulkadiroğlu, A., & Sönmez, T. (2003). Ordinal efficiency and dominated sets of assignments. Journal of Economic Theory, 112(1), 157–172.

Abraham, D., Blum, A., & Sandholm, T. (2007). Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In Proceedings of the 8th ACM Conference on Electronic Commerce (ACM-EC), pp. 295–304. ACM Press.

Ahmadi, S., Ahmed, F., Dickerson, J. P., Fuge, M., & Khuller, S. (2020). An algorithm for multi-attribute diverse matching. In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, pp. 3–9.

Ahmed, F., Dickerson, J. P., & Fuge, M. (2017). Diverse weighted bipartite b-matching. In Proceedings of the 26th International Joint Conference on Artificial Intelligence, pp. 35–41. AAAI Press.

Andersson, T., & Ehlers, L. (2020). Assigning refugees to landlords in sweden: Efficient, stable, and maximum matchings. The Scandinavian Journal of Economics, 122(3), 937–965.

Aygün, O., & Bó, I. (2020). College admission with multidimensional privileges: The Brazilian affirmative action case. American Economic Journal: Microeconomics.

Aygün, O., & Turhan, B. (2020). Dynamic reserves in matching markets: Theory and applications. Journal of Economic Theory, 188.

Aziz, H. (2018). Mechanisms for house allocation with existing tenants under dichotomous preferences. Journal of Mechanism and Institution Design.

Aziz, H. (2020). Developments in multi-agent fair allocation. In Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI), pp. 13563–13568.

Aziz, H., Biró, P., & Yokoo, M. (2022). Matching market design with constraints. In Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), pp. 12308–12316.

Aziz, H., Gaspers, S., & Sun, Z. (2020). Mechanism design for school choice with soft diversity constraints. In Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), pp. 153–159.

Balinski, M., & Sönmez, T. (1999). A tale of two mechanisms: Student placement. Journal of Economic Theory, 84(1), 73–94.