

Adaptive Epidemic Supply Allocation: Integrating Prediction and Fairness

Ziyu Chen*

Notre Dame Preparatory School of Sacred Heart University, Fairfield, CT 06825, United States of America

*Corresponding author's e-mail: Chenziyu568@gmail.com

Abstract

Public health emergencies can quickly create a mismatch between scarce supplies and surging demand. The main challenge is to allocate limited resources efficiently without undermining fairness. Existing methods often prioritize one objective and perform poorly when epidemic conditions change rapidly. This study develops a simulation system for epidemic supply allocation and proposes a predictive-equity adaptive algorithm. The algorithm combines a waiting-time threshold with strategy switching to adjust the balance between efficiency and fairness during distribution. It is tested under baseline, stress-test, and extreme scenarios with random outbreak spikes. Under high uncertainty, the proposed method achieves the highest overall score of 38.61 while maintaining acceptable fairness and stability. The findings indicate that dynamic fairness compensation can reduce persistent allocation disparities and support emergency resource distribution during public health crises.

Keywords

Epidemic Material Allocation; Fairness; Resource Scheduling; Adaptive Algorithm; Simulation Study.

1. Introduction

When COVID-19 hit, pretty much every country ran into medical supply shortages-masks, gowns, ventilators, you name it. That bottleneck seriously hampered containment efforts. The WHO even pointed out that how well a health system functions during a crisis depends a lot on supply capacity and how fast stuff gets where it's needed [1].

As the pandemic dragged on, researchers started paying more attention to allocation problems. The old-school approach? Give priority to places with more cases or higher risk. Makes sense for short-term efficiency, but it also tends to pile resources into a few hotspots while other hospitals go dry for weeks. Dai et al. warned that obsessing over efficiency can widen regional gaps, while going all-in on fairness might drag down overall treatment outcomes [2].

Huang et al. put it bluntly: emergency distribution is basically a tug-of-war between efficiency and fairness [3]. And when resources are tight, fixed rules just don't cut it-they can't keep up with the way an epidemic evolves. That's why more studies are now leaning toward dynamic decision-making, hoping to give systems some much-needed flexibility.

Still, there are gaps. For one, most methods rely on static rules and can't react in real time. For another, some optimize efficiency but ignore long-term fairness. And extreme situations-like sudden random outbreaks-haven't gotten enough attention.

So here's our shot: This study built a simulation platform for epidemic supply allocation, designed a prediction-fairness adaptive algorithm, and pitted it against a static greedy algorithm and a hybrid strategy. Through multi-scenario tests, this study zoom in on

performance under tough, unpredictable conditions, hoping to offer something fresh for public health resource management.

2. Related Research

2.1. Emergency Resource Allocation

This is a big deal in operations research and emergency management. Over the years, people have cooked up all sorts of optimization models-for fires, disaster relief, disease outbreaks, you name it. Zhang et al., for example, designed a city-level fire resource scheduler that did a solid job boosting response times [4].

But pandemics add a whole new layer of mess. Case numbers explode, demand shifts overnight, and static methods just can't keep pace. So researchers are increasingly turning to dynamic and intelligent optimization to make allocation more agile.

2.2. The Efficiency-Fairness Tradeoff

Fairness has become a hot topic lately. Dai et al. pointed out that during major outbreaks, purely demand-based allocation can leave some regions in the lurch for too long [2]. Huang et al. doubled down, saying you have to tackle both efficiency and fairness head-on [3].

Most agree these two goals are inherently at odds. Demand-first tends to ignore some areas; equal-split wastes resources. So how do you strike a balance? That's the million-dollar question right now.

2.3. Supply Chain Resilience

Wright et al. argue that a resilient medical supply chain is the backbone of any effective public health response [5]. In a pandemic, supply chains need to be fast and robust. Ghahremani-Nahr et al., coming from a humanitarian logistics angle, proposed a robust optimization model to keep systems stable under uncertainty [6]. These studies lay important groundwork for emergency scheduling.

Overall, existing work touches on efficiency, fairness, and resilience-but dynamic allocation under random outbreak scenarios is still underexplored. That's where this study step in, trying to design a strategy that adapts and balances both goals.

2.4. Resource Scheduling Algorithms

Scheduling algorithms have been heavily studied in operations research, CS, and intelligent decision-making. The goal is usually to make the most of limited resources while meeting system constraints [7]. Greedy algorithms, adaptive strategies, and fairness-aware methods are among the most popular.

Greedy algorithms are fast and simple-they pick the highest-priority node at each step-so they're common in emergency logistics, healthcare, and cloud computing. But they're shortsighted; they only look at the current state, which can lead to repeatedly favoring high-demand nodes and hurting fairness over time.

To handle uncertainty, researchers have turned to dynamic programming and adaptive scheduling. Dynamic programming can find globally optimal solutions, but it's computationally heavy. Adaptive methods adjust to demand and supply changes, boosting robustness, but many require repeated optimization, which limits real-time use.

Fairness-aware allocation is another key thread. Recent work embeds fairness into objective functions via waiting-time penalties, proportional sharing, or utility models [8]. These help prevent long-term resource hoarding, but many approaches are computationally expensive. To get around the limits of static greedy rules, some have proposed adaptive or distributed

algorithms that react to system changes in real time. Distributed methods, in particular, improve scalability and robustness by letting multiple agents decide locally [9].

So while progress has been made, there's still a gap: lightweight adaptive algorithms that are fast, responsive, and fair under stochastic outbreak conditions are rare. That's why we propose a Predictive-Fair Adaptive Algorithm-it blends demand forecasting, waiting-time-based fairness adjustment, and dynamic strategy switching to do a better job in epidemic resource allocation.

3. Research Methodology

3.1. System Model

This study set up a simulation network for epidemic supply distribution, with two main players: warehouses and hospitals. Warehouses store and dispatch supplies; hospitals are the ones requesting and receiving them. Each hospital comes with its own set of attributes-case count, current inventory, bed capacity, and local population density.

To get a sense of what's coming down the pipeline, built a lightweight forecasting module that estimates future case numbers:

$$\hat{C}_i(t + \Delta) = C_i(t) + 3 \cdot \left(\frac{\sum_i C_i(t)}{N} \times 0.1 \right) \cdot D_i$$

Here, $C_i(t)$ is the current case load at hospital i , D_i is its population density coefficient, N is the total number of hospitals, and $\hat{C}_i(t+\Delta)$ is the predicted case count for the next time step. This isn't meant to be a crystal ball-it just gives us a rough idea of where demand might be heading, so can factor that into scheduling decisions.

3.2. Static Greedy Algorithm

The static greedy algorithm is pretty straightforward: it always serves the hospital that seems most urgent at that moment. Urgency is calculated as:

$$U_i(t) = \alpha \frac{C_i(t)}{K_i} + \beta \left(1 - \frac{S_i(t)}{2C_i(t) + 1} \right) + \gamma D_i$$

Where K_i is capacity, $S_i(t)$ is current stock, D_i is population density, and α, β, γ are weights we tweak. The system picks the hospital with the highest U value and sends resources there.

Pros? It's fast. Cons? Over time, it tends to pour resources into the same few hotspots, leaving others chronically under-supplied.

3.3. Hybrid Strategy Algorithm

This one tries to mix things up. Most of the time (80% of scheduling rounds), it runs the greedy rule. But in the remaining 20% of rounds, it flips the script and gives priority to hospitals that have been waiting the longest-no matter their urgency score.

This helps a bit with fairness, but the fixed 80/20 split means it can't really adapt when the situation changes rapidly. It's a step forward, but still pretty rigid.

3.4. Predictive-Fair Adaptive Algorithm

Now this is main contribution. Instead of sticking to a fixed rule, we introduced a waiting-time threshold mechanism. If a hospital goes without any allocation for more than a certain number

of consecutive rounds-say, 4 rounds in experiments-the system automatically switches into fairness mode.

The trigger condition:

$$\max_i W_i(t) \geq \tau_w$$

Here, $W_i(t)$ is how many rounds hospital i has been waiting, and τ_w is the threshold (set to 4). Once triggered, the system gives top priority to the hospital with the longest wait, making sure no one gets left behind for too long.

In normal operation, the algorithm uses a mix: 60% greedy and 40% round-robin. This keeps efficiency high while still giving fairness a fighting chance—and crucially, it can adjust on the fly as the epidemic evolves.

4. Experimental Design

4.1. Scenario Setup

To see how well each algorithm holds up, designed three test scenarios with increasing difficulty.

4.1.1. Scenario A – Baseline

This is the "easy" case. Resources are relatively plentiful, case numbers rise slowly, and the warehouse can keep up with demand. Think of it as a best-case scenario-used mainly to get a sense of baseline performance.

4.1.2. Scenario B – Stress Test

Here, turn up the heat. Supply drops, case growth accelerates, and things get tight. This lets us see how each algorithm behaves when resources start running low.

4.1.3. Scenario C – Extreme Challenge

This is the one really care about. On top of shortages and fast case growth, throw in random local outbreaks-think new variants, superspreader events, or sudden case clusters.

Outbreak probability is defined as $P_{\text{outbreak}} = 0.1$ per round. When one hits, a random hospital gets a spike in cases:

$\Delta C \sim U(50, 100)$ where $U(50, 100)$ means a uniform random number between 50 and 100. This adds a lot of uncertainty, making Scenario C the toughest test of adaptability.

4.2. Evaluation Metric System

This study didn't just look at one number. Instead, we evaluated performance across three dimensions:

4.2.1. Clinical Effectiveness

How well did allocation help control the outbreak? Track total cases and death rate—lower is better.

4.2.2. Resource Efficiency

How effectively were supplies used? Metrics include utilization rate and average response time. Higher utilization and shorter response times = better efficiency.

4.2.3. Fairness

This is a big one for this study. We used the Gini coefficient as our main fairness indicator:

$$G = \frac{2 \sum_{k=1}^n kx_{(k)}}{n \sum_{k=1}^n x_{(k)}} - \frac{n+1}{n}$$

where x are resource allocations, sorted ascending, and n is the number of hospitals. A Gini of 0 means perfect equality; the higher it goes, the more skewed the distribution. We also tracked max waiting time and allocation variance to get a fuller picture.

5. Experimental Results and Analysis

5.1. Comprehensive Performance Analysis

The numbers tell a pretty clear story-different algorithms shine in different situations.

In Scenario A (the baseline), the hybrid strategy actually came out on top. With plenty of resources and low pressure, even a fixed 80/20 mix did a decent job balancing efficiency and fairness.

In Scenario B (stress test), predictive-fair adaptive algorithm took the lead, scoring 41.95 overall. As resources got tighter, its ability to adjust dynamically started paying off.

And in Scenario C-the chaotic one with random outbreaks-the adaptive algorithm again outperformed the others. The full breakdown is in the table below.

Table 1. Scenario C Composite Score Table

Algorithm Type	Overall Score	Fairness Ranking	System Stability
Static Greedy Algorithm	37.86	3	Low
Hybrid Strategy Algorithm	36.56	2	Medium
Predictive-Fair Adaptive Algorithm	38.61	1	High

Let's start with the composite scores. In the extreme challenge scenario (Scenario C), the Predictive-Fair Adaptive Algorithm came out on top with 38.61. To be honest, the score gaps between the three algorithms aren't huge-but what's worth noting is that the adaptive algorithm managed to hold its lead consistently throughout the high-uncertainty environment. That alone tells us something: the dynamic fairness compensation mechanism this study introduced isn't just for show-it's actually doing its job.

5.2. Fairness Analysis

Now let's look at fairness. The static greedy algorithm has the same problem we mentioned earlier-it keeps feeding resources to the hospitals with the highest demand, and over time, resources become heavily concentrated in those few hotspots. Some hospitals, because their case counts remain persistently high, keep receiving supply after supply, while others might go multiple rounds without getting a single shipment. As this pattern repeats, waiting times between hospitals diverge further and further, and the inequity in resource distribution grows increasingly severe.

The hybrid strategy does mitigate this to some extent by occasionally switching to a compensation mode. But the catch is that the compensation ratio is fixed in advance-the system can't adjust it based on the current state. Sometimes it compensates more than needed; other times not enough. Flexibility is still a weak spot.

The Predictive-Fair Adaptive Algorithm, in contrast, takes a real-time monitoring and on-demand triggering approach. It continuously tracks each hospital's waiting rounds, and as soon as any hospital's consecutive wait exceeds the preset threshold, the system immediately flips into fair compensation mode and prioritizes that hospital for supply. As a result, resources are distributed much more evenly across hospital nodes, and extreme waiting times are significantly reduced. From a practical standpoint, the real value of this mechanism is that it ensures no hospital gets left behind for too long-which in turn boosts the overall coordination capacity of the entire healthcare system.

5.3. Robustness Analysis

Robustness, plain and simple, is about whether an algorithm can hold its ground when things go sideways-when random disturbances hit, does it still keep its composure?

In Scenario C, we introduced random outbreak events. Each outbreak causes a sudden spike in cases at some hospital, which puts considerable pressure on the scheduling strategy. The static greedy algorithm reacts quickly-when an outbreak hits, it immediately throws resources that way-but the trade-off is that other hospitals in need might get sidelined. The hybrid strategy can ease this tension in some cases, but because its compensation ratio is fixed, it struggles to cope with the complexity of rapidly changing situations.

The Predictive-Fair Adaptive Algorithm handles this more gracefully: it adjusts its strategy dynamically according to the system's current operational state. When a specific region experiences a resource crunch, the system rebalances the efficiency-fairness trade-off on the fly. Results from multiple repeated runs also show that this algorithm exhibits less performance fluctuation and greater stability-the robustness advantage is fairly clear.

5.4. Discussion of Results

Putting the three algorithms side by side, their distinct profiles become quite clear.

The static greedy algorithm is fast and simple-it works well in resource-rich environments where speed is the top priority, but its fairness shortcomings are hard to ignore. The hybrid strategy performs reasonably well in relatively stable settings, achieving a basic balance between efficiency and fairness, but its fixed-ratio mechanism limits its ability to adapt when conditions change sharply. The Predictive-Fair Adaptive Algorithm is structurally more complex and costlier to implement, but in resource-constrained environments with random outbreak shocks, its overall performance is noticeably better.

One thing worth emphasizing is that the advantage of this study proposed algorithm doesn't lie in dominating any single metric-it lies in lifting the system's overall performance. Especially in the extreme challenge scenario, it manages to strike a decent balance across efficiency, fairness, and stability simultaneously-and that's where its real value resides.

6. Conclusion

This study examines how to balance efficiency and fairness in medical supply allocation during pandemics, particularly when static strategies cannot respond adequately to changing conditions. We developed a discrete-event simulation system for resource scheduling and proposed an adaptive allocation algorithm that combines demand forecasting with a waiting-time-based compensation mechanism.

The results show that differences among algorithms remain limited when resources are relatively sufficient. As resource pressure increases, however, the adaptive algorithm performs more strongly. In the extreme scenario with random outbreak events, it achieved the highest overall score of 38.61 and maintained favorable levels of fairness and system stability. These

results indicate that waiting-time-triggered compensation can reduce persistent allocation disparities and improve system responsiveness under uncertain epidemic conditions.

Future research could improve demand prediction through more advanced machine-learning models, extend the framework to multi-warehouse coordination, and test the model using real-world pandemic data. These extensions would strengthen its practical relevance and decision-support value.

References

- [1] World Health Organization. (2020). COVID-19 strategic preparedness and response plan: Operational planning guidelines to support country preparedness and response. World Health Organization.
- [2] Dai, Z. X., Lan, H. J., Hai, N., et al. (2025). Balancing fairness and efficiency in dynamic vaccine allocation during major infectious disease outbreaks. *Scientific Reports*, 15(1), Article 1371.
- [3] Huang, S., Zhu, H., Wang, H., et al. (2025). Balancing the trade-off between efficiency and equity in a stochastic emergency supplies allocation problem. *Applied Mathematical Modelling*, Article 116242.
- [4] Zhang, X., Zhao, K., Gao, S., et al. (2025). Optimization of urban fire emergency resource allocation based on a pre-allocated swarm algorithm. *Fire*, 8(1), Article 27.
- [5] Wright, A. M., Snowden, A., Saunders, M., et al. (2024). The necessity of healthcare supply chain resilience for crisis preparedness. *Healthcare Management Forum*, 37(Suppl. 1).
- [6] Ghahremani-Nahr, J., Nozari, H., & Szmelter-Jarosz, A. (2024). Designing a humanitarian relief logistics network considering the cost of deprivation using a robust-fuzzy-probabilistic planning method. *Journal of International Humanitarian Action*, 9(1), Article 19.
- [7] Huang, B., Zhou, M., Lu, X. S., & Abusorrah, A. (2023). Scheduling of resource allocation systems with timed Petri nets: A survey. *ACM Computing Surveys*, 55(11), 1–27.
- [8] Chen, V. X., & Hooker, J. N. (2023). A guide to formulating fairness in an optimization model. *Annals of Operations Research*, 326(1), 581–619.
- [9] Doostmohammadian, M., Aghasi, A., Pirani, M., Nekouei, E., Zarrabi, H., Keypour, R., ... Johansson, K. H. (2025). Survey of distributed algorithms for resource allocation over multi-agent systems. *Annual Reviews in Control*, 59, Article 100983.