

RESERVE-SYSTEM PLANNING FOR ETHICAL ALLOCATION OF SCARCE HEALTHCARE RESOURCES DURING PUBLIC HEALTH EMERGENCIES

Mick Duncan

Public health emergencies expose a recurring management problem: institutions must allocate scarce life-saving resources quickly while honoring ethical commitments that cannot be collapsed into a single administrative score. This article examines reserve systems as a planning framework for the allocation of vaccines, ventilators, intensive-care capacity, and antiviral therapies during crisis conditions. The analysis uses a structured conceptual and comparative reading of the formal institutional framework and documented COVID-19 policy applications reported by Pathak, Sönmez, Ünver, and Yenmez. A reserve system is defined by three managerial levers: the division of units into categories, the number of units assigned to each category, and the priority rule applied within each category. The article's contribution is to translate those levers into operational design choices, compare how they function across documented policy settings, and identify the implementation trade-offs that arise when resources differ in urgency, durability, and monitoring requirements. The documented experience of U.S. states and health systems that used reserve-based policies during the pandemic, including Pennsylvania, Massachusetts, Tennessee, Connecticut, California, Richmond and Henrico (Virginia), and Washington, DC, is synthesized to test feasibility rather than to claim causal superiority. The analysis shows that reserve systems are best understood not merely as ethical devices, but as operational planning instruments whose value depends on explicit category design, sequencing, transparency, and ongoing adjustment under uncertainty.

© The author(s) 2024. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license (<http://creativecommons.org/licenses/by/4.0/>).

INTRODUCTION

Scarcity during a public health emergency is not only a clinical problem; it is an institutional design problem. When demand for vaccines, ventilators, intensive-care beds, or antiviral treatments exceeds supply, institutions must allocate resources under time pressure, informational limits, and intense public scrutiny while preserving procedural legitimacy. These decisions are rarely governed by a single objective. Pandemic guidance commonly invokes multiple ethical principles at once, including equity, welfare maximization, reciprocity, solidarity, and nondiscrimination [4, 13, 7]. The resulting administrative challenge is therefore structural: policy makers must translate plural principles into a workable allocation procedure without obscuring the trade-offs embedded in that procedure.

A traditional response has been the *priority system*, in which all claimants are placed into a single priority order and units are assigned from the top downward. In many jurisdictions, especially for ventilator and ICU triage, this has taken the form of a priority point system tied to clinical indicators such as the Sequential Organ Failure Assessment (SOFA) score [9, 7]. The appeal of this structure is administrative simplicity. Its weakness is that it compresses ethically distinct claims into one ranking.

This article develops a planning-oriented account of *reserve systems* as an alternative institutional design. A reserve system divides units into categories, assigns a quantity of units to each category, and applies category-specific priority rules. In management terms, it replaces one monolithic queue with a structured allocation architecture. This is not simply a normative refinement. It is a practical planning framework that gives administrators visible policy levers: they can define categories, adjust reserve sizes, publish cutoffs, and tailor implementation to the characteristics of the resource being rationed. The article makes three linked contributions: it translates reserve design into operational language for management audiences, it compares documented pandemic applications through a common institutional lens, and it clarifies the implementation contingencies that arise when the underlying resource differs in durability and urgency.

The article is organized as follows. First, it explains why single-list priority systems encounter organizational and ethical limits in crisis settings. Second, it presents the reserve system as a planning framework, emphasizing the administrative levers it creates. Third, it synthesizes documented COVID-19 policy applications drawn directly from the published record. Fourth, it examines how implementation should differ across vaccines, ventilators and ICU beds, and antiviral therapies. The article concludes with implications for management and planning practice. The aim is not to claim that reserve systems dominate every alternative in every setting, but to show why they materially widen the feasible design space for emergency allocation.

ANALYTICAL APPROACH AND EVIDENCE BASE

This article is a conceptual and comparative policy analysis. Its evidentiary base is the published formal framework and documented case material reported by [7], supplemented by the cited bioethics and emergency-allocation literature. No hypothetical scenario panel, synthetic patient-level dataset, or simulated performance metrics are introduced. The analysis therefore does not offer outcome estimation; instead, it uses source-grounded comparison to test whether reserve systems are institutionally feasible, administratively intelligible, and adaptable across resource types.

First, it uses the article's formal institutional characterization of reserve systems. Second, it draws on the source article's documented policy examples from the COVID-19 pandemic, including reserve-based allocation plans for vaccines and therapies in multiple jurisdictions. Third, it incorporates the source article's resource-specific implementation discussion for vaccines, ventilators and ICU beds, and antiviral treatments. Across those components, the manuscript evaluates reserve systems on four planning criteria: representational

Table 1: Multiprinciple Point System for Ventilator Allocation

Principle	Specification	1	2	3	4
Save the most lives	Prognosis for short-term survival (SOFA score)	SOFA < 6	SOFA 6–9	SOFA 10–12	SOFA > 12
Save the most years	Prognosis for long-term survival (medical assessment of comorbidities)	No comorbidities	Minor comorbidities	Major comorbidities	Severe comorbidities; likely death within 1 year
Life-cycle principle	Prioritize those who have had the least chance to live through life's stages	Age 12–40	Age 41–60	Age 61–74	Age > 74

Adapted from the multiprinciple point system discussed by [14] and reproduced in [7]. Lower point totals indicate higher priority.

capacity, transparency, operational adaptability, and implementation burden.

The analytical objective is not to estimate causal effects. It is to clarify which institutional features matter for planning, why reserve systems offer greater managerial flexibility than single-list systems, where that flexibility creates additional administrative demands, and how documented pandemic applications illuminate design choices relevant to emergency governance. The resulting claims are intentionally bounded: the article argues for practical value and institutional plausibility, not automatic superiority in every emergency setting.

PRIORITY SYSTEMS AND THEIR ORGANIZATIONAL LIMITS

Before reserve systems can be assessed, the managerial logic of the conventional priority system must be stated clearly. Under a priority system, all eligible claimants are ordered according to one common rule, and every unit is allocated according to that same order [7]. In practice, the priority order may be determined by a single clinical criterion or by a multiprinciple point system that aggregates several criteria into one score. The model is administratively familiar, but its apparent neatness depends on collapsing heterogeneous claims into a single decision channel.

For ventilator allocation, a widely discussed multiprinciple design is summarized in Table 1. It combines short-term survival, long-term survival, and life-cycle considerations in a single point structure. The design is understandable, but its core limitation is that incommensurable ethical claims must be forced into a common metric. Once this occurs, the institution loses the ability to preserve distinct ethical channels at the level of implementation.

The administrative problem is not that such systems are irrational or unusable. The problem is that a single priority order becomes too restrictive when policy makers seek to protect several ethical values simultaneously. Equity claims raised by disadvantaged communities, reciprocity-based claims for essential personnel, and urgency-based clinical claims do not always map cleanly onto one scale [11, 7]. In organizational terms, the priority system offers only one instrument for multiple objectives, so any gain in simplicity may come at the cost of ethical compression and fragile legitimacy.

This constraint became especially visible during COVID-19. [7] note that, for ventilator allocation, more than half of U.S. states with guidelines adopted either single-principle or multiprinciple priority point systems, and many of these relied heavily on SOFA-based scoring. The same article also highlights criticism that such

systems can reproduce pre-existing inequities when clinical measures proxy unequal prior access to health care. A management system that appears neutral at the point of application may therefore generate contested distributional effects at the policy level.

RESERVE SYSTEMS AS A PLANNING FRAMEWORK

A reserve system addresses the foregoing limitation by replacing one all-purpose queue with a structured allocation design. In the formulation summarized by [7], a reserve system has three core parameters:

1. the division of all available units into reserve categories;
2. the number of units assigned to each category; and
3. a priority order for each category.

This structure matters because it transforms ethical commitments into administrative levers and makes trade-offs auditable. Instead of asking one ranking to do all the work, planners can assign particular shares of scarce supply to distinct objectives and define the priority rule appropriate to each one. An open category can preserve broad access, while one or more protected categories can be used for specific beneficiary groups or policy goals. The gain is not that disagreement disappears, but that the location of disagreement becomes explicit and therefore governable.

A priority system can therefore be understood as a limiting case: a single-category reserve system. This continuity is important for implementation. Administrators, clinicians, and public officials familiar with priority systems need not abandon the general logic of ordered allocation. They instead operate a richer version of it, albeit one that requires clearer eligibility definitions, more monitoring, and stronger coordination.

Planning Advantages

From a management perspective, reserve systems offer four conditional advantages.

First, they improve *goal specification*. Ethical values are not hidden inside one score. They are represented through identifiable categories and visible quantities, which makes disagreements easier to diagnose and revise.

Second, they improve *policy transparency*. The size of a reserve, the identity of beneficiary groups, and the within-category priority rule can all be publicly described. That benefit is realized only when eligibility definitions are stable enough to be explained and audited in practice.

Third, they improve *institutional flexibility*. Different resources can be governed by the same broad architecture while using different operational rules, which matters when the form of scarcity changes faster than the institution's ethical commitments.

Fourth, they improve *committee-level coordination*. Because reserve systems allow multiple groups to see their priorities reflected in separate policy channels, they can facilitate compromise in task forces and emergency governance settings [7]. They do not eliminate conflict, but they can move conflict into an explicit design stage where trade-offs are easier to manage.

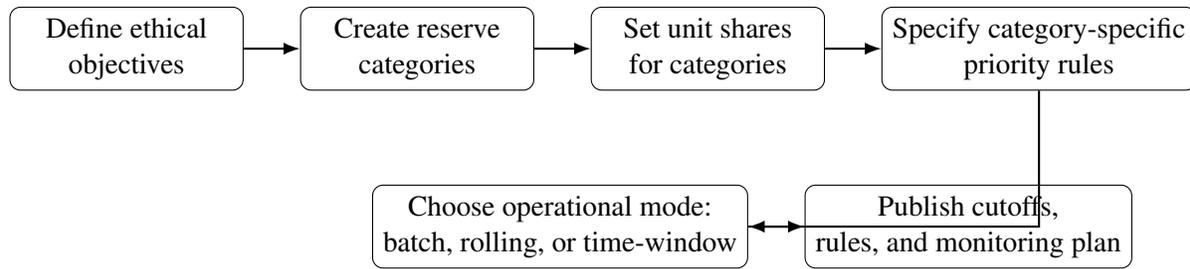


Figure 1: Operational planning sequence for reserve-system implementation in emergency allocation.

Sequential and Smart Processing

In many practical applications, reserve categories are processed sequentially: units associated with one category are allocated before moving to the next. This is administratively intuitive, but it can be inefficient when beneficiary groups overlap. If a claimant eligible for a protected category is assigned first through the open category, a protected unit may later remain idle.

The source article discusses this problem directly in its appendix on smart reserve matching [7]. The planning implication is straightforward: sequence matters. When categories overlap, administrators should not treat reserve ordering as a clerical detail. It is a substantive design choice that can affect whether resources are fully utilized, whether protected categories achieve their stated purpose, and whether the system's promised fairness survives real-time implementation.

A PRACTICAL PLANNING SEQUENCE

Figure 1 presents a practical planning sequence for reserve-system deployment. The sequence is consistent with the institutional logic documented in the published literature and is designed for emergency planning committees, hospital systems, and public health agencies.

The sequence highlights that reserve-system design is not exhausted by moral reasoning. It requires explicit planning choices: the institution must determine its categories, reserve sizes, priority rules, and operating cadence before implementation begins. The practical contribution of the framework lies precisely in forcing those operational choices to be specified in advance rather than improvised at the point of scarcity.

DOCUMENTED RESERVE-SYSTEM USE DURING COVID-19

The most important validation available to this analysis is that reserve systems were not merely theoretical during COVID-19. They were recommended or adopted in multiple jurisdictions, and the published record provides concrete examples of how the framework was operationalized [7]. Table 2 summarizes selected cases directly reported in the source article.

Several planning lessons emerge from these cases.

First, reserve systems were used in both *over-and-above* and *minimum-guarantee* forms. This matters because the same broad institutional model can support different political and administrative choices. An over-and-above reserve adds extra units above a baseline distribution, whereas a minimum guarantee protects a floor within the existing allotment.

Table 2: Selected Reserve-System Applications Reported During COVID-19

Jurisdiction / Institution	Date	Documented Policy Design
UPMC Medical System, Pennsylvania	May 2020	Remdesivir was allocated through a reserve-based weighted lottery with three special provisions: hardest-hit area (based on the Area Deprivation Index), essential-worker status, and whether the patient was not expected to die within one year. More than 300 patients were processed through the procedure.
National Academies (NASEM) framework	October 2020	The framework recommended a 10% over-and-above reserve for hard-hit areas, defined as the top 25% of the Social Vulnerability Index distribution within a state.
Tennessee	October 2020	The initial vaccine plan allocated 5% of doses equitably among all 95 counties, reserved 10% in an over-and-above form for hard-hit or highly vulnerable counties, and distributed the remaining 85% by population; the vulnerability reserve was later reduced to 5% for some vaccines because of storage-related constraints.
Massachusetts	December 2020	The state vaccine plan reserved 20% of doses in an over-and-above form for communities with disproportionate COVID-19 burden and high social vulnerability.
Pennsylvania	January 2021	A 10% minimum-guarantee reserve was used for vaccination of Phase 1A healthcare personnel who were not affiliated with any hospital system.
Connecticut	February 2021	A 25% community-level reserve was announced in minimum-guarantee form for communities with high social vulnerability.
California	March 2021	The state reserved 40% of vaccines for the hardest-hit or most socially vulnerable quartile and 20% for each remaining quartile. Within each quartile, 70% of reserved doses were allocated by age eligibility and 30% by sector eligibility.
Richmond and Henrico, Virginia	March 2021	Phase 1b vaccine allocation used multiple categories, including 50% for people aged 65 and older, 23% for frontline essential personnel, 23% for those aged 16–64 with comorbidities, and 4% for people in congregate-care settings and previously eligible but still unvaccinated individuals; preregistered claimants were prioritized by category-specific rules.
Washington, DC	March 2021	The plan used age, worker type, qualifying medical conditions, and prioritized zip codes as separate category structures, including paired reserves for prioritized zip codes and non-prioritized areas.
Massachusetts (therapeutics)	November 2020	A reserve system for monoclonal antibodies used an over-and-above reserve for patients from hard-hit or socially vulnerable communities, with an even lottery within categories.

Compiled from the cases summarized in [7]. The table reports published policy descriptions rather than new empirical estimates.

Second, jurisdictions varied in the *granularity* of their category design. Some systems created one high-vulnerability reserve; others, such as California, partitioned the allocation architecture across multiple geographic or social-vulnerability strata.

Third, the cases show that reserve systems can use *category-specific priorities*. California separated age-based and sector-based allocation within quartiles. Richmond and Henrico used multiple categories with distinct policy aims and explicit category-specific prioritization. This is especially important for management practice because it demonstrates that reserve systems are not limited to simple quota overlays; they can structure the internal ordering of access as well.

Fourth, the institutional form was portable across *resource types*. The same planning logic was used not only for vaccines but also for Remdesivir and monoclonal antibodies. This cross-resource portability strengthens the argument that reserve systems are a generalizable governance tool. At the same time, the cases do not prove outcome dominance; they provide bounded validation that reserve systems are administratively workable, revisable under changing logistics, and capable of accommodating more than one policy objective at once.

RESOURCE-SPECIFIC IMPLEMENTATION CONSIDERATIONS

A reserve system is not implemented the same way for every resource. One of the most practically useful parts of the published literature is its explicit recognition that operational design depends on the physical and temporal characteristics of the good being allocated [7]. In this article, that cross-resource comparison serves as the main stress test of the framework: a design that works only for one class of goods is less persuasive as a general planning tool. Two questions are especially important:

1. Is the resource consumed upon allocation, or is it durable and potentially reassignable?
2. Is allocation immediately urgent, or can it be batched?

These questions determine the appropriate planning cadence.

Vaccines

Vaccines are consumed upon allocation and ordinarily do not require immediate bedside triage. This makes batch allocation feasible. As vaccine shipments arrive, units can be distributed through a reserve system on a static or quasi-static basis. For planning purposes, this is the cleanest application of the reserve framework [7]. The principal operational challenge is less bedside triage than distribution discipline: eligibility rosters, shipment forecasting, and reporting systems must still be accurate enough for reserved doses to reach intended groups.

The source article also notes that vaccine allocation can be implemented dynamically on a first-come-first-served basis within reserve categories, especially in settings such as pharmacies or local providers. From a management perspective, the important point is not that one mode is always superior; it is that the reserve structure can be layered onto either batch scheduling or rolling distribution. The better choice depends on data reliability, the speed of replenishment, and the institution's ability to monitor whether reserved channels are being exhausted as intended.

Ventilators and Intensive-Care Beds

Ventilators and ICU beds present a more complex problem. They are durable, used over time, and urgently needed when a patient presents. These features make purely static implementation impractical [7]. Allocation must be dynamic, and the system may need to be reevaluated as new patients arrive and existing patients' conditions change. Here the design burden is higher: reserve rules must function under severe time pressure while remaining clinically coherent and administratively reviewable.

This creates a planning problem that is distinct from vaccine distribution. If withdrawal is allowed under governing policy, reallocation can operate more directly, though it remains ethically and legally controversial. If withdrawal is not allowed, a grandfathering structure is required, giving continued priority to current occupants of units they already hold. Either way, the planning architecture must specify in advance how reevaluation works, how category accounting is handled over time, which legal and ethical constraints govern reassignment, and who has authority to document and review those decisions.

For management and planning practice, the central lesson is that reserve systems for durable emergency resources require not only category design but also occupancy governance.

Antiviral Drugs and Time-Sensitive Therapies

Antiviral drugs and comparable therapies combine features of the other two cases. They are consumed upon use, but allocation is typically urgent. The published literature identifies two practical implementation options: first-come-first-served operation within reserve categories, or assignment in short predefined time windows [7]. The key managerial trade-off is between speed and auditability.

This distinction has important planning consequences. A rolling rule emphasizes speed and bedside usability. A time-window rule emphasizes forecasting, batch coordination, and alignment with expected demand across categories. Institutions should therefore choose an implementation mode that matches the operational environment in which therapies are dispensed, the documentation burden they can sustain, and the degree of variation they expect in short-term demand.

IMPLICATIONS FOR MANAGEMENT AND PLANNING PRACTICE

The reserve-system framework has direct relevance beyond bioethics. It speaks to the design of management systems under severe uncertainty by linking ethical pluralism to explicit administrative choices.

Institutional Design as Capacity Planning

Reserve systems convert broad ethical objectives into administrative parameters. This allows planning bodies to treat fairness not as an ex post aspiration but as an ex ante design problem. In that respect, reserve systems resemble other forms of strategic capacity planning: they require explicit segmentation, target-setting, and rule specification before demand materializes. The gain is strongest when category definitions correspond to data the institution can actually identify and verify in real time.

Transparency and Public Legibility

Emergency allocation systems are judged not only by outcomes but also by their perceived legitimacy. Reserve systems improve legibility because category definitions, reserve quantities, and within-category priority rules can be communicated clearly. This supports accountability, reduces ambiguity in frontline implementation, and can strengthen public trust. Poorly specified categories, however, can undermine trust just as quickly, so transparency depends on disciplined rule drafting as much as on public explanation.

Cross-Functional Coordination

Pandemic allocation involves clinicians, administrators, public health officials, legal advisors, and community representatives. Because reserve systems make policy trade-offs visible and separable, they can facilitate structured deliberation among these actors. Distinct stakeholder concerns can be represented through different categories rather than forced into one contested score. In practice, this often shifts disagreement upstream into committee design, which is preferable to leaving the same conflict unresolved at the bedside.

Monitoring and Adjustment

A reserve system should not be treated as a fixed chart that becomes irrelevant after publication. It requires monitoring: administrators must track category utilization, unfilled units, bottlenecks, and whether protected categories are reaching the intended populations. The documented COVID-19 cases show that reserve sizes were sometimes modified as operational conditions changed [7]. This indicates that reserve design is best understood as a living planning instrument, not a one-time declaration. Without feedback and revision, a reserve can become symbolic rather than effective.

CONCLUSION

The allocation of scarce healthcare resources in a public health emergency is a problem of institutional planning under constraint. A single priority list can be simple, but it is often too blunt to represent the plurality of ethical objectives that emergency policies are expected to serve. Reserve systems provide a more versatile governance structure by dividing scarce units into categories, assigning category sizes, and using category-specific priority rules.

The documented COVID-19 experience provides real-world validation that reserve systems were usable across multiple resource domains, including vaccines, antiviral drugs, and monoclonal antibody therapies. The cases also show that reserve systems can be configured in different ways—over-and-above, minimum-guarantee, geographically targeted, occupation-based, or clinically structured—without abandoning a common institutional framework.

The evidence assembled here supports a narrower but stronger conclusion. Reserve systems do not eliminate scarcity, political disagreement, or administrative burden. They do, however, expand the design space available to planners by making ethical trade-offs explicit, revisable, and operationally tractable. For management and planning research, their significance lies in showing how institutional design can translate plural goals into auditable allocation rules. In crisis conditions, that is not a marginal advantage. It is a core feature of competent institutional design.

REFERENCES

- [1] Commonwealth of Pennsylvania (2020) *Interim Pennsylvania Crisis Standards of Care for Pandemic Guidelines*. Harrisburg, PA: Commonwealth of Pennsylvania.
- [2] Dur U, Kominers SD, Pathak PA, Sönmez T (2018) Reserve design: Unintended consequences and the demise of Boston's walk zones. *Journal of Political Economy* 126(6):2457–2479.
- [3] Dur U, Pathak PA, Sönmez T (2020) Explicit versus statistical preferential treatment in affirmative action: Theory and evidence from Chicago's exam schools. *Journal of Economic Theory* 187:104996.
- [4] Emanuel EJ, Persad G, Upshur R, et al. (2020) Fair allocation of scarce medical resources in the time of Covid-19. *New England Journal of Medicine* 382(21):2049–2055.
- [5] Hafalir IE, Yenmez MB, Yildirim MA (2013) Effective affirmative action in school choice. *Theoretical Economics* 8(2):325–363.
- [6] National Academies of Sciences, Engineering, and Medicine (2020) *Framework for Equitable Allocation of COVID-19 Vaccine*. Washington, DC: The National Academies Press.
- [7] Pathak PA, Sönmez T, Ünver MU, Yenmez MB (2024) Fair allocation of vaccines, ventilators and antiviral treatments: Leaving no ethical value behind in healthcare rationing. *Management Science* 70(6):3999–4036.
- [8] Persad G, Peek ME, Emanuel EJ (2020) Fairly prioritizing groups for access to COVID-19 vaccines. *JAMA* 324(16):1601–1602.
- [9] Piscitello GM, Kapania EM, Miller WD, Rojas JC, Siegler M, Parker WF (2020) Variation in ventilator allocation guidelines by U.S. state during the coronavirus disease 2019 pandemic: A systematic review. *JAMA Network Open* 3(6):e2012606.
- [10] Rubin EB, Persad G, Huang J, et al. (2021) Prioritizing equity in allocation of scarce critical care resources. *Chest* 160(2):621–631.
- [11] Schmidt H (2020) The way we ration ventilators is biased. *The New York Times*. April 15.
- [12] Sönmez T (2023) *Minimalist Market Design*. Working paper.
- [13] Truog RD, Mitchell C, Daley GQ (2020) The toughest triage: Allocating ventilators in a pandemic. *New England Journal of Medicine* 382(21):1973–1975.
- [14] White DB, Katz MH, Luce JM, Lo B (2009) Who should receive life support during a public health emergency? Using ethical principles to improve allocation decisions. *Annals of Internal Medicine* 150(2):132–138.
- [15] White DB, Lo B, Mitra N, et al. (2022) Allocation of scarce therapies for coronavirus disease 2019: Development of a weighted lottery system. *Chest* 161(2):565–574.

Mick Duncan, Scottish architect associated with the Scottish Parliament project.

Manuscript Published; 30 November 2024.