# The Allocation of Uncertainty

Continuation, Control, and Evidence in the Politics of Loss

Flyxion
Independent Researcher

September 2026

**Abstract**

Any claim that a loss was necessary is a counterfactual claim: it asserts that no feasible alternative would have prevented it. Such claims are made constantly, and they are rarely settled by argument alone. Whether they hold depends on who can observe the loss, who specifies the alternatives against which it is judged, who can audit the account, and whose interpretation is accepted. This essay treats those capacities as a distribution. It distinguishes three capacities that any politics of loss allocates: *continuation capacity*, the ability to keep reachable futures as common conditions degrade; *control capacity*, the ability to shape the rules that govern those conditions; and *epistemic capacity*, the ability to establish why a given loss occurred. The three are analytically separable but causally coupled. When they align across a population, selected loss becomes easier to stabilize as necessary loss. Epistemic capacity is treated as a vector with ordered components, and the essay argues that the most consequential of them is specification: power over the categories that determine what evidence must exist at all. A relation of interested epistemic dependence is defined, and shown in one case to run upward through the nominal hierarchy of regulation, from community to regulator to regulated party. Three case studies test four mechanisms: classification-conditioned disclosure and counter-observation at xAI's turbine sites in Memphis and Southaven, counter-measurement in Flint, and counter-specification in pandemic triage protocols. The essay's conclusion is that the allocation of uncertainty is part of the allocation of loss.

## 1 There Was No Alternative

The sentence "there was no alternative" appears wherever losses are distributed: in budget hearings, triage protocols, permit decisions, restructuring plans, and emergency declarations. It is usually spoken as a report about the world. It is in fact a claim about a space of possibilities. To say that a loss was necessary is to say that no feasible alternative allocation of the available resources would have preserved what was lost. Writing $\mathcal{F}(B)$ for the allocations feasible under a genuine constraint $B$, and $A$ for the allocation actually chosen, the claim is

$$\nexists\, A' \in \mathcal{F}(B) : \ \Gamma_{\mathrm{excl}}(A') \subsetneq \Gamma_{\mathrm{excl}}(A),$$

where $\Gamma_{\mathrm{excl}}(A)$ is the set of continuations excluded under $A$. That is a strong statement over an entire feasible space. A companion essay, *The Value of Floods*, stated the principle in a single line: necessity is a counterfactual claim. It then divided any lost region of continuation into three parts,

$$\Gamma_{\mathrm{lost}} = \Gamma_{\mathrm{F}} \sqcup \Gamma_{\mathrm{S}} \sqcup \Gamma_{\mathrm{U}},$$

losses forced by a genuine constraint, losses selected by an allocation rule, and losses whose status the available evidence cannot determine. It proposed tests for moving losses out of the third set, and observed in passing that uncertainty can be allocated as well as resolved.

This essay develops that observation. Its starting point is that a claim of necessity has two lives. It is first *made*, and anyone can make it. It is then *maintained*, against challenge, over time, in institutions that decide which evidence exists, who can see it, what alternatives it is compared against, and whose reading of it is accepted. The first life is rhetorical. The second is institutional, and it is where the politics of loss is mostly decided.

The argument is organized around three capacities that any such politics allocates. Continuation capacity is the ability to keep reachable futures as common conditions degrade. Control capacity is the ability to shape the rules that govern those conditions. Epistemic capacity is the ability to establish why a given loss occurred. Collect them for each person or institution $i$ in a vector

$$\mathbf{z}_i = (c_i, \kappa_i, \boldsymbol{\varepsilon}_i).$$

The thesis concerns their alignment, not their inequality. Each of the three is unequally distributed in every society, and much of that inequality is ordinary. The danger arises when they travel together: when the actors least exposed to a loss are also the actors who shape the rules producing it and the actors whose account of it prevails. Under that alignment, selected loss becomes easier to stabilize as necessary loss. The word *stabilize* carries the claim. Concentrated epistemic capacity does not create the initial assertion of necessity. It determines whose assertion survives contact with counterevidence.

Sections 2 through 4 define the three capacities, with most of the attention given to the third. Section 5 states how they couple and proposes three evidentiary tiers for claims about that coupling. Section 6 defines a relation of interested epistemic dependence. Section 7 examines three cases, which together test four mechanisms. Sections 8 and 9 state what follows for remedies and where the framework stops.

## 2 Continuation Capacity

Let $Q$ denote the quality of common conditions: water, air, power, public health systems, the shared infrastructure on which ordinary life depends. For each person $i$, let $\Gamma_i(Q)$ be the set of futures reachable by $i$ at that level of common quality, decomposed as in the companion essay into a part supplied in common and a part reachable only through private capacity,

$$\Gamma_i(Q) = \Gamma^{\text{common}}(Q) \cup \Gamma_i^{\text{priv}}.$$

**Definition 1** (Continuation capacity)**.** *The continuation capacity of $i$ at common quality $Q$ is*

$$c_i(Q) = \Phi_i\big(\Gamma_i(Q)\big),$$

*where $\Phi_i$ is an evaluation functional that is monotone under inclusion: $\Gamma \subseteq \Gamma'$ implies $\Phi_i(\Gamma) \leq \Phi_i(\Gamma')$.*

The functional form is deliberately weak. It does not assume that heterogeneous futures are commensurable, or that more of them in some cardinal sense is always better. A plain measure $\mu$ over reachable futures is one special case of $\Phi_i$, and a convenient one for illustration, but nothing in the argument requires it. What the definition requires is only that losing reachable futures cannot increase continuation capacity.

**Definition 2** (Exposure)**.** *The exposure of $i$ to common conditions is*

$$e_i(Q) = \frac{\partial c_i(Q)}{\partial Q}.$$

The sign convention follows the companion essay and needs one sentence of explanation, since the word *exposure* can suggest harm. A high $e_i$ means that $i$'s reachable futures depend heavily on common quality: when $Q$ falls, high-exposure actors lose the most. An actor with $e_i \approx 0$ is insulated. Its continuation capacity does not move when the common world deteriorates, because its reachable futures are supplied privately.

A shock to $Q$ therefore divides the affected population. Some can replace what the commons supplied with privately acquired substitutes; call them $P_{\text{replace}}$. Others cannot, and for them a contraction of $\Gamma^{\text{common}}$ is simply a contraction of $\Gamma_i$; call them $P_{\text{loss}}$. The distinction is not one of expenditure but of capacity to replace.

The concept has an obvious relative in the capability approach of Amartya Sen and Martha Nussbaum, which evaluates wellbeing by the substantive freedoms people have to do and be what they have reason to value. The capability approach is not static, and it would be a mistake to contrast the two by saying that capabilities concern the present while continuation concerns the future. The difference is narrower and more specific. Continuation capacity is *degradation-indexed*. It asks which substantive opportunities remain reachable as a specified common substrate deteriorates, and exposure measures the sensitivity of that answer to the substrate. The novelty is not futurity. It is attention to the dependence of reachable futures on common conditions, and to the unequal distribution of that dependence.

## 3 Control Capacity

Let $a_i$ stand for $i$'s actions on the rules, budgets, siting decisions, standards, and exceptions that determine the condition of common systems.

**Definition 3** (Control capacity)**.** *The control capacity of $i$ is the sensitivity of common continuation to $i$'s actions,*

$$\kappa_i = \left\| \frac{\partial \Gamma^{\text{common}}}{\partial a_i} \right\|.$$

High control capacity is not itself domination. Legislators, regulators, engineers, and utility managers all have large $\kappa_i$, and societies need them to. The relevant distinction is Philip Pettit's. In his republican account, domination is exposure to another's capacity for arbitrary or uncontrolled interference, and it can exist even where no interference occurs, because the dominated party must live at the discretion of the dominating one. Control capacity becomes domination when its exercise is insufficiently constrained by those whose continuation it affects.

The companion essay identified a configuration it called *retained voice with withdrawn dependence*: actors who have ceased to rely on common systems but keep influence over them. In the present notation this is

$$e_i \to 0, \qquad \kappa_i \gg 0.$$

The actor is insulated from the consequences of common deterioration while remaining able to shape its causes. Albert Hirschman's analysis of exit and voice treats members who leave a declining organization and so stop pressing it to repair itself. The case here is different and, in one respect, worse. The insulated actor does not leave the jurisdiction. It remains a citizen, a

taxpayer, a donor, or a regulated party, and keeps its voice over systems on which it no longer depends.

# 4 Epistemic Capacity

The first two capacities describe who can bear a loss and who can shape it. The third describes who can establish what the loss was and why it happened. It requires more care than the others, because it is the least visible and because it is internally divided.

## 4.1 A vector, not a scalar

**Definition 4** (Epistemic capacity)**.** *The epistemic capacity of $i$ with respect to a class of losses is the vector*

$$\boldsymbol{\varepsilon}_i = \left(\varepsilon_i^S,\ \varepsilon_i^D,\ \varepsilon_i^A,\ \varepsilon_i^I\right),$$

*whose components are:*

- $\varepsilon^S$ specification capacity*: power over the categories, models, feasible sets, population boundaries, and definitions that determine which evidence is relevant, and which evidence is obligatory;*
- $\varepsilon^D$ data capacity*: the ability to produce, access, and retain evidence;*
- $\varepsilon^A$ audit capacity*: the ability to test, contest, and independently verify an account;*
- $\varepsilon^I$ interpretive authority*: the degree to which one's account of a loss is accepted by the institutions that act on it.*

A scalar summary $\varepsilon_i = \sum_k w_k \varepsilon_i^k$ may be useful for some purposes, but it should not be introduced first, because the components come apart in ways that matter. A community can produce abundant local evidence and have no say over the categories against which that evidence is judged. An agency can hold data and audit powers without interpretive legitimacy. A firm can own the data about its own operations and still lose in court. Calling all of these "high epistemic capacity" would hide exactly the internal allocation the concept exists to describe.

Specification capacity is listed first because it frequently comes first causally. Whoever fixes the category into which an activity falls may thereby fix what must be measured about it, what must be reported, what must be retained, and therefore what can later be known. The cases below show this ordering in practice:

$$\varepsilon^S \longrightarrow \varepsilon^D \longrightarrow \varepsilon^A.$$

A classification settles a reporting obligation; the reporting obligation settles what data exist in the regulatory record; the data in the record settle what can be audited. This ordering is the strongest argument for treating epistemic capacity as a vector. Its components are not merely different; in the cases that matter most they are sequential, and control over the first propagates through the rest.

The definition of specification includes population boundaries for a reason that the cases make concrete. Before any test of necessity is run, someone has decided whose losses count as losses in the relevant proceeding. A boundary drawn by jurisdiction rather than exposure can place the people who bear a harm outside the forum that evaluates it.

### 4.2 Observability

Possessing evidence and being obliged to supply it are different conditions, and the difference is central to what follows.

**Definition 5** (Observability)**.** *An institution $R$* observes *a class of changes in a regulated state when there is a reliable institutional path by which evidence of those changes must become available to $R$, with a delay bounded by $R$'s own rules rather than by the discretion of the party whose state is changing.*

On this definition an agency can possess a great deal of information at particular moments and still fail to observe the activity it regulates, if the arrival of that information depends on the regulated party's choice to send it. Voluntary or episodic disclosure is not observability. The distinction will turn out to be the key to the anchor case.

### 4.3 The underdetermined region as an institutional product

In the companion essay the underdetermined region $\Gamma_{\mathrm{U}}$ appeared as a residue: the part of a loss that the tests could not yet classify. Its size is better understood as a function of institutions,

$$\Gamma_{\mathrm{U}} = \Gamma_{\mathrm{U}}\big(E, \mathcal{F}, \mathcal{A}, \mathcal{P}\big),$$

depending on the evidence $E$ that exists, the feasible alternatives $\mathcal{F}$ that are represented, the audit institutions $\mathcal{A}$ available to test claims, and the procedures $\mathcal{P}$ for adjudicating between competing accounts. Each argument is itself allocated. Evidence exists where someone was obliged or able to produce it; alternatives are represented where someone was able to specify them; audits happen where someone has the standing and resources to run them.

### 4.4 Interventions as operators on the underdetermined region

An epistemic intervention, whether a monitoring program, an audit, a lawsuit, a records request, or an independent study, can be represented as an operator

$$T : \Gamma_{\mathrm{U}} \longrightarrow \Gamma_{\mathrm{F}} \cup \Gamma_{\mathrm{S}} \cup \Gamma'_{\mathrm{U}}, \qquad \Phi(\Gamma'_{\mathrm{U}}) < \Phi(\Gamma_{\mathrm{U}}).$$

A successful intervention reduces underdetermination. It does not presuppose the direction of the result. Some losses will move into $\Gamma_{\mathrm{S}}$, revealed as selected; others will move into $\Gamma_{\mathrm{F}}$, confirmed as forced. That symmetry is essential. Without it, "redistributing epistemic capacity" would quietly become a synonym for proving wrongdoing, and the framework would be unable to recognize a genuine constraint when it found one. The normative criterion is therefore stated in terms of resolvability:

| epistemic redistribution should increase resolvability, not predetermine classification. |
|---|

### 4.5 Three bodies of work, three locations in the model

Three literatures bear directly on this account, and each belongs at a distinct point in it rather than in a general review.

Robert Proctor's program of *agnotology*, the study of the cultural and institutional production of ignorance, describes the mechanisms that enlarge $\Gamma_{\mathrm{U}}$: doubt manufactured deliberately,

as in the tobacco industry's campaigns, and ignorance produced structurally, by research that is never funded or data that are never collected. Its location in the model is the growth of the underdetermined region.

Miranda Fricker's account of *epistemic injustice* describes deficits in interpretive authority. Testimonial injustice occurs when a speaker's credibility is discounted because of prejudice; hermeneutical injustice occurs when a group lacks shared concepts adequate to render its experience intelligible, because it has been excluded from the practices that generate such concepts. Its location is $\varepsilon^I$: people can suffer a loss without the resources to make its causal structure publicly credible.

Sheila Jasanoff's idiom of *co-production* holds that the ways a society knows the world and the ways it governs itself are made together, each shaping the other. Its location is the coupling between $\kappa$ and $\boldsymbol{\varepsilon}$, which is the subject of the next section. Jasanoff's work supplies the reason the three capacities, though analytically separable, cannot be treated as causally independent.

# 5 How the Capacities Couple

## 5.1 Alignment

The three capacities can be measured, at least in principle, across a population. Positive pairwise covariances,

$$\mathrm{Cov}(c, \kappa) > 0, \qquad \mathrm{Cov}(\kappa, \varepsilon) > 0, \qquad \mathrm{Cov}(c, \varepsilon) > 0,$$

describe an institutional system in which protection, rule-setting power, and evidentiary authority tend to travel together. They are useful population diagnostics, but they do not establish that the same actors occupy the high tail of all three distributions at once. A stronger object does. For thresholds $c^*, \kappa^*, \varepsilon^*$, define the joint tail

$$A^* = \{\, i : c_i > c^*,\ \kappa_i > \kappa^*,\ \varepsilon_i > \varepsilon^* \,\}.$$

The size of $A^*$, and the share of each capacity concentrated within it, measure alignment directly. The configuration the companion essay called complete impunity is membership in $A^*$ together with the insulation condition:

$$e_i \to 0, \qquad \kappa_i \gg 0, \qquad \varepsilon_i \gg 0.$$

Such an actor is insulated from the consequence, influential over its production, and dominant over its explanation.

## 5.2 The persistence of accounts

It would be a mistake to say that epistemic capacity as such supports necessity claims. High epistemic capacity can undermine a necessity claim as easily as defend it, and in several of the cases below it did. What matters is the alignment of epistemic capacity with an interested party. Let $L_i(N)$ denote the institutional persistence of actor $i$'s necessity account $N$: the degree to which that account remains the operative classification of a loss in the institutions that act on it. The coupling hypothesis can then be stated as

$$\kappa_{i,t} \uparrow \longrightarrow\ \boldsymbol{\varepsilon}_{i,t+1} \uparrow \longrightarrow\ L_i(N)_{t+1} \uparrow \longrightarrow\ \kappa_{i,t+2} \uparrow,$$

conditional on the relevant institutions allowing $i$'s epistemic capacities to translate into authoritative classification. The first arrow holds because control decides who gathers data, who funds expertise, and who writes procedures. The third holds because accepted knowledge decides which interventions appear legitimate. This is Jasanoff's co-production stated as a hypothesis about a particular actor over time.

### 5.3 Three evidentiary tiers

The loop is a temporal causal claim, and it should not be asserted on the strength of an illustration. Claims about coupling fall into three tiers of increasing strength:

$$\text{co-location} \;<\; \text{control-conditioned persistence} \;<\; \text{feedback}.$$

*Co-location* means that the same actor holds control and epistemic advantages. *Control-conditioned epistemic persistence* means that the exercise of control helps preserve the epistemic advantage over time. *Feedback* requires, in addition, that the resulting epistemic advantage enlarges or preserves the actor's control. The cases below document the first two tiers clearly. They also document the coupling $\varepsilon \to \kappa$, but running against the incumbent: independent epistemic capacity displacing an official account and bringing outside oversight to bear on the actors who had maintained it. None of them documents the full loop for a single actor. That remains a hypothesis, and the essay treats it as one.

### 5.4 Borrowed necessity

One further mechanism appears in the anchor case and has no place in the loop as stated. A necessity claim can be lent. A third party with large control capacity can adopt an operator's claim that a loss is unavoidable and move it into a domain where the counterfactual test is hardest to run. The clearest such domain is national security, in which the feasible set and the evidence needed to contest it can be withdrawn from public view by classification. The sequence is

$$\begin{aligned}&\text{necessity claim} \to \text{restricted-evidence domain}\\ &\quad\to \text{counterfactual test unrunnable} \to \text{loss held in } \Gamma_{\mathrm{U}}.\end{aligned}$$

Borrowed necessity is not the feedback loop, since the operator's own epistemic advantage does not grow. It is a way of closing a forum without resolving anything in it. Three forms of the claim should be kept apart:

| | |
|---|---|
| ordinary necessity: | no feasible alternative exists; |
| borrowed necessity: | another authoritative institution establishes that none exists; |
| audit closure: | this forum cannot test whether one exists. |

The third is not evidence that the necessity claim is false. It changes the claim's contestability. In the notation of Section 4, closing an audit path does not move a loss into $\Gamma_{\mathrm{F}}$ or $\Gamma_{\mathrm{S}}$; it prevents the operator $T$ from acting on $\Gamma_{\mathrm{U}}$ at all. The epistemic problem is then not insufficient evidence but the absence of an institution authorized to say what the evidence establishes. The general proposition is compact:

$$\boxed{\text{authority can travel farther than auditability.}}$$

A necessity claim can cross institutional boundaries, from a military office to a court, or from a state agency to a federal filing, while the mechanisms for contesting it stay behind.

# 6 Configurations and Epistemic Dependence

## 6.1 Snapshots

Treating each capacity as high or low yields eight configurations. These are best understood as snapshots of the dynamic system just described rather than as types of political order. Four are informative.

*High on all three*: complete impunity, as defined above.

*High continuation and control, low epistemic capacity*: selection that can be exposed. The actor can bear and shape the loss but cannot control its account. This is where audits, litigation, and investigative reporting bite.

*Low continuation and control, high epistemic capacity*: affected parties who can document their own losses. This is the counter-configuration, and several of the cases turn on it.

*Low on all three*: causal dependence. This is not merely deprivation. The group loses continuation, has no voice in the allocation, and cannot establish why it lost. Its characteristic condition is reliance on an account of its loss supplied by whoever made the allocation.

## 6.2 Interested epistemic dependence

The last configuration motivates a relational definition, since the problem it names is not a shortfall in some absolute quantity of knowledge.

**Definition 6** (Epistemic dependence)**.** *$i \prec_E j$ when establishing the causal classification of a loss suffered by $i$ depends materially on data, categories, models, procedures, or evidentiary institutions substantially controlled by $j$.*

**Definition 7** (Interested epistemic dependence)**.** *$i \prec_E^* j$ when $i \prec_E j$ and $j$ has a stake in how $i$'s loss is classified.*

Modern societies depend on other people's expertise constantly, and the first relation is therefore ubiquitous and usually benign. Patients depend on laboratories, travelers on aviation regulators, citizens on statistical agencies. The second relation is the one that matters: dependence on an epistemic gatekeeper who has an interest in what the gate returns. One consequence follows directly from the definitions:

$$\text{specialization} \neq \text{epistemic domination.}$$

Specialization produces $\prec_E$. Epistemic domination requires $\prec_E^*$.

The relation need not follow the lines of formal authority. In the anchor case below it runs upward through the nominal hierarchy of regulation:

$$\text{community} \ \prec_E \ \text{regulator} \ \prec_E^* \ \text{regulated party.}$$

The regulator depends on the regulated party for timely knowledge of changes in the regulated state, and the regulated party has an obvious stake in how its activity is classified. The dependence is partial, since the regulator retains substantial independent capacity, but it is precisely located: it concerns whether evidence of change was obliged to reach the regulator at all.

# 7 Three Cases, Four Mechanisms

The cases were chosen to test different mechanisms, not to illustrate the same one repeatedly. The anchor case, xAI's turbine operations in Memphis and Southaven, tests classification-conditioned disclosure and counter-observation. Flint tests counter-measurement and is the historical case in which the counter-path succeeded. Pandemic triage tests counter-specification, and it is the case in which the method most clearly returns results in both directions.

## 7.1 Classification and disclosure: xAI in Memphis and Southaven

**Colossus 1 and Boxtown.** In June 2024 xAI began operating its Colossus supercomputer in South Memphis, beside Boxtown, a working-class, majority-Black neighborhood. The facility arrived in an already burdened airshed. As Klein and Taylor document, poverty rates in Boxtown run double the city's average, and the surrounding county leads Tennessee in childhood asthma-related emergency visits. An air-toxics study conducted in southwest Memphis between 2008 and 2010 estimated a cumulative cancer risk from ambient carcinogens about four times the national average, with benzene, formaldehyde, and acrylonitrile as the principal drivers. These figures describe the conditions into which Colossus was introduced; they do not by themselves establish any health effect attributable to it. The distinction matters here, because the four-times figure is a modeled risk from ambient concentrations, not an observed cancer rate, and is sometimes repeated as the latter. To power the facility, the company installed some thirty-five methane gas turbines without air permits.

The first evidentiary move came from outside the regulatory system. After residents reported turbines at the site, the Southern Environmental Law Center (SELC) filed a public-records request with the Shelby County Health Department for any permits; the department reported no results. When the mayor said that only fifteen turbines were actually running, SELC commissioned aerial surveys with thermal imaging and found thirty-three turbines giving off significant heat. This is counter-observation in its simplest form: an affected party, lacking access to the operator's records and receiving no account from the regulator, built its own observational record of the activity.

The regulatory dispute that followed turned on classification rather than observation. The health department took the position that the temporary turbines were "nonroad engines," exempt from permitting. An appeal filed in July 2025 by the NAACP and Young, Gifted & Green contested that determination, arguing that it would allow operators to install and run any number of new turbines without written approval, public notice, or pollution controls. A second dispute concerned the scope of measurement: city officials reported that independent tests had found no dangerous levels of air pollutants near the site, and SELC responded that the tests had not measured ozone in an area whose ozone averages had exceeded federal limits.

Both disputes are exercises of specification capacity. The first decides whether any evidence of the activity was required at all; the second decides which pollutants count as the relevant evidence of harm. The health department granted xAI a permit for fifteen turbines in July 2025. The company removed its other unpermitted turbines after SELC, on behalf of the NAACP, sent a notice of intent to sue under the Clean Air Act. That is a documented instance of epistemic capacity changing conduct, achieved outside the formal appeal.

The appeal itself ended without a ruling on the merits. On December 15, 2025, after a seven-hour hearing, the Memphis and Shelby County Air Pollution Control Board voted six to one to dismiss it as moot. Department attorneys told the board that misinformation surrounded

the case. Because the dismissal rested on mootness, the question whether a year of unpermitted operation had been lawful was never decided. The loss remained in $\Gamma_{\mathrm{U}}$, with no forum left in which to resolve it. That is the clearest illustration in this essay of the claim that the underdetermined region is not a neutral residue: a procedural disposition closed the forum and left the classification exactly where it was.

**Colossus 2 and Southaven.** The same practice then appeared across the state line. In July 2025 an xAI affiliate, MZX Tech, acquired a former power plant in Southaven, Mississippi, a few miles from the company's second Memphis data center. The Mississippi Department of Environmental Quality (MDEQ) approved an exemption for eighteen unpermitted turbines in August 2025. The count rose to twenty-seven by December and thirty-three by March 2026. MDEQ maintained that these engines did not require permits because they fell into a "temporary-mobile" category limited to use of under a year.

On March 10, 2026, three weeks after the state's only public hearing, the Mississippi Environmental Quality Permit Board approved a construction permit for forty-one permanent turbines. That permit covered turbines not yet installed and excluded the unpermitted temporary units already running on the site. The NAACP, represented by SELC and Earthjustice, sued in federal court in April 2026 over the unpermitted turbines, and separately appealed the permit on the ground that it relied on faulty data that minimized pollution impacts.

The decisive evidentiary fact concerns disclosure. MDEQ told Mississippi Today that xAI had added nineteen temporary-mobile turbines between March 25 and May 2, that the company had notified the agency of the additions only that week, and that xAI was not required to tell the agency when it brought additional generators on site. When asked about the combined emissions of the growing cluster, the agency's spokesperson answered that, as indicated by the facility, all the temporary turbines were equipped with emissions controls. The regulator's knowledge of the scale of the activity, and its account of the activity's emissions, both rested on the regulated party's voluntary statements.

Court filings then showed how far the regulatory record could lag the physical state of the site. SELC reported being notified that the facility had forty-six turbines on May 6 and fifty-seven on May 15, and a filing stated that the May 6 disclosure was outdated the moment it was sent. The information reached the challengers by a long path: from the operator, through its consultant, to the agency, and from the agency to SELC through a public-records request. In July 2026 Reuters counted fifty-nine unpermitted turbines, roughly twice the number the company had acknowledged, a figure the plaintiffs' counsel then repeated. MDEQ, for its part, had said in July that the plant was protected by the Clean Air Act's mobile-source provisions because the turbines were portable units. At the end of July, MDEQ announced a schedule to retire the site's temporary generators, by then numbering sixty-nine, while allowing some to run as much as five months past their original deadline.

**The federal category.** The classification question has a federal layer, and the federal text bears directly on it. In January 2026 EPA finalized revised performance standards for stationary combustion turbines. The preamble states the governing distinction plainly: a turbine is stationary because it is not self-propelled or intended to be propelled while performing its function, and turbines may be mounted on a vehicle or trailer for portability and still be considered stationary. The rule exempted from the standards only turbines subject to applicable Title II mobile-source standards, and EPA added that, to its knowledge, no such turbines were then being used in

temporary applications. The route out of the stationary category, in other words, existed on paper and was empty in practice. The rule added a subcategory for stationary turbines used in temporary applications and exempted certain sources from Title V permitting. The subcategory defines temporariness by residence rather than by the individual machine: a small or medium turbine may remain at the same source for up to twenty-four consecutive months, the clock runs from the start of operation whether or not the unit operates throughout, relocation within the same source does not restart it, and any temporary turbine that replaces another in the same function counts toward the same period. A July 2026 correction clarified that owners or operators may *elect* to be a temporary combustion turbine, as an alternative to the otherwise applicable emission standards and monitoring, recordkeeping, and reporting requirements. The same rule can be described as tightening or loosening, depending on the baseline chosen. Against the alternative on which EPA had sought comment in its proposal, an outright exemption for turbines used in temporary applications, it chose a regulated subcategory instead. Against full stationary permitting, a law-firm analysis described it as significantly reducing permit obligations for data-center turbines. Both descriptions are accurate, and the choice between baselines is itself an act of specification. What the rule plainly does is place part of the classification decision in the regulated party's hands, and attach to that decision the party's obligations to monitor, record, and report.

**Population and forum.** Two further features complete the case. The Southaven turbines sit a few minutes from the state line, and residents on the Tennessee side, who live downwind, have no standing in the Mississippi permitting process. The population entitled to take part in evaluating the harm was defined by jurisdiction rather than by exposure. And on June 15, 2026, the United States moved to intervene in the NAACP's suit as a plaintiff and to dismiss it with prejudice. The government's memorandum rests on three supports. The first is a classification: it recites that MDEQ had determined that the state's implementation plan did not require permits for the turbines, and that the state had declined enforcement. The second is a necessity claim, supported by a declaration from the Department of War's chief digital and artificial intelligence officer describing xAI's model as one of only four frontier systems able to support national security applications, and by a letter from Mississippi's governor describing the economic disruption a shutdown would cause. The memorandum concludes that the suit threatens American national, economic, and energy security. The third support is procedural and, for this essay, the most consequential. The government argues that the executive's enforcement discretion is preclusive; that under the rule it invokes the court has no adjudicatory role in its dismissal; and that its reasons for seeking dismissal do not dictate the outcome and are, like other decisions not to enforce, unsuitable for judicial review.

The legal merits of that argument, which draws on a long line of cases about prosecutorial discretion and on judicial opinions questioning citizen suits under Article II, are not this essay's subject. Its structure is. The filing performs, in a single document, three operations the framework distinguishes. It takes a contested state classification as its premise, which is specification capacity carried forward. It supplies a necessity claim from a third party with large control capacity, which is borrowed necessity. And it seeks to close the audit path through which the claim could be tested: the citizen suit is the mechanism by which affected parties exercise $\varepsilon^A$ against both the operator and the regulator, and the government's position is that the reasons offered for closing it need not be examined. The counterfactual test is not failed. It is declared outside the forum.

**What the case shows.** The anchor case documents the full chain from classification to contestability:

$$\text{classification} \rightarrow \text{reporting obligation} \rightarrow \text{regulatory observability} \rightarrow \text{public contestability}.$$

The category did not merely determine what the activity was for regulatory purposes. It determined whether evidence of changes in the activity was obliged to travel to the regulator, which in turn determined what the public could later obtain and contest. The relation regulator $\prec_E^*$ operator holds in the precise sense defined above: the agency depended on the operator for timely knowledge of changes in the regulated state, and the operator had a direct stake in the classification. In terms of the evidentiary tiers, the case is strong evidence for control-conditioned epistemic persistence: the classification, maintained across two jurisdictions, preserved an information asymmetry over time. It is not evidence for the full feedback loop, since nothing in the record shows the operator's epistemic advantage enlarging its own control. What the federal filing adds is a different mechanism: a third party's control capacity deployed to remove the audit path, so that the classification persists not because it has been defended but because the forum that could test it is asked to stand down. The invariant across Memphis and Southaven is not the particular legal category, nonroad in one place and temporary-mobile in the other, but the transformation

$$\boxed{\text{classification of activity} \longrightarrow \text{classification of evidence obligations.}}$$

### 7.2 Counter-measurement: Flint

Flint is the historical case in which the counter-path succeeded, and it documents each link of the coupling more completely than any other.

**Control over evidence.** In April 2014 Flint began supplying treated Flint River water under a state-appointed emergency manager who had replaced local representative decision-making. The Flint Water Advisory Task Force, appointed by the governor, later concluded that the causes of the crisis lay primarily with the state, through its agencies' failures and its appointed emergency managers' misjudgments. The same state agency that decided the treatment question also administered compliance. MDEQ incorrectly determined that corrosion control, required under EPA's Lead and Copper Rule, was not needed immediately, and allowed the city to complete two six-month monitoring periods before deciding whether it was. Control over the system and control over the evidence about the system sat in the same hands.

**Asymmetric evidence.** Records later obtained by the ACLU of Michigan and by Virginia Tech's Marc Edwards indicated that state officials had encouraged the city to find samples with low lead levels and to disqualify two samples with high readings. The records reached the public only through public-records requests. The protocol that produced official compliance could not be audited by the people whose water it measured until they obtained the documents that described it.

**Persistence of the official account.** For roughly eighteen months officials maintained that the water met federal standards. When the Virginia Tech team presented contrary findings in September 2015, officials dismissed the researchers and residents. When Mona Hanna-Attisha,

a pediatrician at Hurley Medical Center, released an analysis showing increased elevated blood lead levels among young children, a state health department spokesperson said that the state's data were not in line with hers and were more comprehensive, and state officials described her data as "spliced and diced." This is $L_i(N)$ at work: an institution with more data used that advantage to keep its classification standing against counterevidence. It is control-conditioned epistemic persistence in the incumbent's direction.

**Counter-capacity.** The account that replaced it was assembled from three sources outside the state's control. A resident, LeeAnne Walters, worked with Miguel Del Toral of EPA Region 5, whose April 2015 correspondence noted that sequential sampling was needed to see the lead levels at her home and that Flint had not been operating any corrosion control treatment since leaving the Detroit system. Edwards's team then organized citizen sampling across the city, distributing some three hundred kits, producing an instructional video, and publishing results as they came in. Finally, Hanna-Attisha obtained and analyzed blood lead data. Her published study found that the incidence of elevated blood lead levels among children under five increased from 2.4% to 4.9% after the water source change, and that neighborhoods with the highest water lead levels experienced a 6.6% increase. In her later congressional statement she reported that in the areas with the highest water lead, rates more than tripled; in one ward they rose from 4.9% to 15.7%.

**Reclassification and federal oversight.** Proof that the water was harming children moved the loss out of $\Gamma_{\mathrm{U}}$. The city returned to Detroit water in October 2015. In January 2016 EPA issued an emergency administrative order, and from January 21, 2016, the agency oversaw its implementation. EPA's Office of Inspector General later found implementation and oversight lapses at every level: EPA, the state, MDEQ, and the city. Among the task force's recommendations was the creation of a culture in state government that is not defensive about concerns and evidence contradicting official positions. That recommendation is, in effect, an institution naming the mechanism this essay describes.

**What the case shows.** Flint documents co-location, control-conditioned persistence of the official account, and, in the end, the coupling $\varepsilon \to \kappa$: independent epistemic capacity displaced the official account and brought the system's operators under federal oversight. The coupling ran against the incumbent, though what shifted was oversight of the system rather than ownership or day-to-day operation of it. That result matters for the remedies discussed below, because it shows the mechanism by which redistributing epistemic capacity can change the allocation of control. The Michigan Civil Rights Commission's 2017 report found that the crisis reflected systemic racism. Its findings can be read through Fricker's account of hermeneutical injustice, which describes how a community's experience can fail to register as credible evidence, though the commission itself does not use that vocabulary.

### 7.3 Counter-specification: crisis standards of care

The triage case differs from the others in one respect that makes it indispensable: the allocation rules were written and published in advance. The procedural test from the companion essay was built into the design. The question the case poses is whether publication removes the concentration of epistemic capacity, or only makes the specification inspectable.

**Explicit criteria.** In March and April 2020, disability rights organizations and individuals filed complaints with the Office for Civil Rights (OCR) of the Department of Health and Human Services, alleging that crisis standards of care in Alabama, Kansas, New York, Pennsylvania, Tennessee, Utah, and Washington discriminated against people with disabilities. Alabama's guidelines, since rescinded, had called for withholding ventilators from patients with severe or profound intellectual disability, moderate to severe dementia, or severe traumatic brain injury. OCR resolved complaints with Alabama on April 8, 2020, Pennsylvania on April 16, Tennessee on June 26, and Utah on August 20, in each case after the state revised its protocol. OCR also expressed concern that many plans still used criteria such as resource intensity or long-term survival that could discriminate on the basis of disability. Publication made the embedded selection contestable, and contest changed the rules. This is counter-specification in the plainest sense.

**The dispute over constraint-relevance.** The case also shows why counter-specification does not simply mean removing criteria. Analysts writing in the *New England Journal of Medicine* argued both that categorical exclusion on the basis of disability is indefensible and that guidelines ignoring near-term prognosis can produce outcomes inconsistent with responsible stewardship of scarce resources. That is the invariance question from the companion essay stated in clinical form. Near-term prognosis is plausibly constraint-relevant: a crisis genuinely limits how many patients can be treated, and the likelihood of surviving treatment bears on how the limit is met. Categorical diagnoses and long-term survival estimates may instead function as proxies for designated categories of person. The contest over which features count as constraint-relevant is a contest over $\varepsilon^S$.

**Implicit criteria.** A second body of evidence goes further, because it concerns a criterion that appeared neutral. Many state standards ranked patients using the Sequential Organ Failure Assessment (SOFA) score, a measure of near-term mortality risk and therefore apparently constraint-relevant. Studies published in 2021 found that the score did not perform equally across racial groups. An analysis of 113,158 patients admitted for sepsis or acute respiratory failure at twenty-seven hospitals found that SOFA overestimated in-hospital mortality among Black patients and underestimated it among White patients; in a simulation, 9.4% of all Black patients were erroneously excluded from the highest priority group. Removing the score's creatinine component reduced the miscalibration. A study at Yale New Haven Health found that Black patients hospitalized with COVID-19 had greater odds than White patients of an elevated SOFA score, with an odds ratio of 1.49. A third, using data from more than 95,000 patients in 233 intensive care units, reached consistent conclusions. In October 2021, after Idaho activated crisis standards that ranked patients by SOFA, Harald Schmidt, Dorothy Roberts, and Nwamaka Eneanya argued that officials were still implementing outdated and more inequitable triage guidelines, with no effective process for updating influential model frameworks. A review of state standards in January 2022, when the Omicron wave pushed twenty states above 85% intensive care capacity, asked whether they still relied on SOFA.

This evidence forces a refinement of the invariance test. It is not enough that a constraint-relevant feature be selected. Its *measurement* must also be invariant across the groups it ranks. A prognostic score miscalibrated by race converts a genuine constraint into structural selection while appearing entirely neutral. The disability complaints were counter-specification of an explicit rule; the SOFA studies were counter-measurement of an implicit one.

**What the case shows.** Triage is the case in which the method most clearly returns results in both directions. Some criteria were validated as constraint-relevant: near-term prognosis survives scrutiny. Some criteria were revealed as selection: categorical exclusions by diagnosis did not. And some measurements of a valid criterion were shown to carry selection within them. The case does not end in a verdict of wrongdoing. It ends with a smaller underdetermined region, which is exactly what the criterion of resolvability asks for.

### 7.4 What the cases show together

| Case | Mechanism | Tier documented | Resolution |
|---|---|---|---|
| xAI: Southaven | classification-conditioned disclosure | control-conditioned persistence | open; lawsuit and appeal pending |
| xAI: Boxtown | counter-observation; contested classification | persistence; conduct changed by notice of suit | legal question left in $\Gamma_U$ by mootness |
| Flint | counter-measurement | persistence; $\varepsilon \to \kappa$ against the incumbent | official account displaced; federal oversight imposed |
| Triage | counter-specification; counter-measurement | co-location; revision under contest | split; rules revised, score use persisted |

Three conclusions follow. First, epistemic capacity is not a single thing, and its components are ordered: in the anchor case, classification determined reporting, reporting determined observability, and observability determined what could be contested. Second, the underdetermined region is shaped by institutions, and procedural dispositions such as mootness can hold a loss in it indefinitely. Third, the coupling between evidence and control is real and runs in both directions. The same mechanism that lets an incumbent stabilize its account lets independent evidence dislodge it, which is why the distribution of epistemic capacity is a political question and not only a technical one.

## 8 Remedies as Increases in Resolvability

If the problem is the alignment of epistemic capacity with interested parties, the remedies are ways of creating independent paths out of $\Gamma_U$. Each is justified by the resolvability it adds, not by any presumption about where the resolved losses will land.

*Pre-registered counterfactuals.* The counterfactual test asks whether some feasible alternative allocation would have excluded less. It is the test most exposed to motivated reasoning, because feasibility is itself contestable. Alternatives to be tested should be stated before outcomes are known, and assessed by parties who do not benefit from the allocation under review.

*Observability by rule.* Where an activity's classification determines its reporting obligations, the classification should not be allowed to switch off the regulator's knowledge of changes in the activity. A reporting duty attached to the activity rather than to the category would have prevented the Southaven record from lagging the site by eleven turbines in nine days. The federal temporary-turbine rule already contains the germ of this principle: by counting replacement units toward a single residence period, it attaches temporariness to the continuing function at a site rather than to whichever machine happens to be performing it.

*Preservation of audit paths.* Where affected parties hold a legal route to test a claim of necessity, as citizen suits provide under the Clean Air Act, the closing of that route should require reasons that can themselves be examined. A forum that may be dismissed without examined reasons is an audit path that exists at the discretion of the party it audits.

*Audit independence.* Audits should not be funded, appointed, or scoped by the audited party. The EPA's Office of Inspector General, which examined the agency's own conduct in Flint, notes that its funding comes from Congress separately from the agency for exactly this reason.

*Access to the evidence behind necessity claims.* In Flint, Boxtown, and Southaven alike, the decisive documents reached the public through public-records requests. Rights of access are therefore not a procedural nicety. They were the actual route by which losses left $\Gamma_{\mathrm{U}}$.

*Community monitoring with standing.* Thermal imaging flown by volunteer pilots and citizen sampling kits both produced evidence that official channels had not. Such evidence should carry standing to trigger review, so that its production is not wasted when no forum is obliged to receive it.

*Public specification of the population and the threshold.* Who counts as affected, and what counts as a livable outcome, should be decided openly and with reference to exposure rather than jurisdiction alone. The Southaven case shows how a boundary drawn by a state line can remove the most exposed residents from the forum that evaluates their exposure.

Some of these remedies will move losses into $\Gamma_{\mathrm{F}}$, confirming that a constraint was genuine. That is a success of the method, not a failure.

# 9 Limits

*Expertise is real.* Concentrating epistemic capacity is often efficient, and most epistemic dependence is benign. The target of the framework is interested dependence, $\prec_E^*$, not specialization.

*Epistemic capacity works in both directions.* Manufactured doubt can keep genuinely forced losses underdetermined, just as it can protect selected ones. A community, a firm, or an advocacy organization can each enlarge $\Gamma_{\mathrm{U}}$ strategically. The framework's criterion of resolvability applies to all of them.

*The components are hard to measure.* Specification, data, audit, and interpretive capacity are easier to recognize in a case than to quantify across a population. In practice they will often be inferred from proxies such as data-access regimes, audit rates, disclosure obligations, and litigation outcomes.

*The configurations simplify.* The eight high and low configurations compress continuous, coupled quantities into binary snapshots. Population-level covariances can conceal configurations specific to particular groups, which is why the joint tail $A^*$ is the better measure where data allow.

*The feedback loop is unproven.* None of the cases documents the full loop in which an actor's epistemic advantage enlarges its own control. The anchor case documents control-conditioned persistence, and Flint documents the coupling running against the incumbent. The strongest form of the thesis remains a hypothesis.

*Several facts in the anchor case are still developing.* The Southaven litigation and permit appeal were pending as of this writing, and an evidentiary hearing on the preliminary injunction had been scheduled for August 24, 2026; this essay does not report its outcome. The account of the federal government's position is taken from its filed motion and memorandum, but the

declaration and gubernatorial letter attached to it are described only as the memorandum characterizes them. The turbine counts rest on the parties' disclosures and on press counts.

## 10 Conclusion

"There was no alternative" is a claim about a space of possibilities, and its credibility depends on the distribution of the capacity to check it. That distribution has components. Someone decides what category an activity belongs to, and thereby decides what must be reported about it. Someone holds the data that the reporting produces, or does not produce. Someone has the standing to audit the account, and someone's reading of it is accepted by the institutions that act.

When those capacities belong to the parties least exposed to a loss and most able to shape its causes, a selected loss can be held as a necessary one for a long time. It does not need to be defended as necessary on the merits. It only needs to remain unresolved, in a region the institutions never obliged anyone to clear. The cases show that such regions can also be cleared: by residents with sampling kits, by pediatricians with blood lead data, by pilots with thermal cameras, by advocates with records requests and published protocols. What clears them is not better argument alone. It is the redistribution of the capacity to know.

| The allocation of uncertainty is part of the allocation of loss. |
|---|

The question is not only who survives the flood and who controls the levee, but who is believed about why the levee failed.

## References


[1] Flyxion, *The Value of Floods: Private Protection and Differential Survivability*, 2026.

[2] Flyxion, *The Forgotten Specification: Human Imitation, Continuation Gain, and the Geometry of Technological Progress*, 2026.

[3] A. Sen, *Development as Freedom*. New York: Knopf, 1999.

[4] M. C. Nussbaum, *Creating Capabilities: The Human Development Approach*. Cambridge, MA: Harvard University Press, 2011.

[5] P. Pettit, *Republicanism: A Theory of Freedom and Government*. Oxford: Oxford University Press, 1997.

[6] A. O. Hirschman, *Exit, Voice, and Loyalty*. Cambridge, MA: Harvard University Press, 1970.

[7] R. N. Proctor and L. Schiebinger (eds.), *Agnotology: The Making and Unmaking of Ignorance*. Stanford: Stanford University Press, 2008.

[8] M. Fricker, *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.

[9] S. Jasanoff (ed.), *States of Knowledge: The Co-Production of Science and Social Order*. London: Routledge, 2004.

[10] N. Klein and A. Taylor, *End Times Fascism: And the Fight for the Living World*. New York: Farrar, Straus and Giroux, 2026.

[11] Southern Environmental Law Center, "Inside Memphis' fight against xAI," updated June 2026. `https://www.selc.org/news/inside-memphis-fight-against-xai/`

[12] Southern Environmental Law Center, "Memphis health leaders grant air permit for xAI data center," July 3, 2025.

[13] Tennessee Lookout, "NAACP, others appeal xAI turbine permits for Memphis data center," July 17, 2025.

[14] Local Memphis, "Memphis and Shelby County Air Pollution Control Board dismisses air quality challenge to xAI project," December 15, 2025.

[15] Action News 5, "Shelby Co. Air Pollution Control Board denies appeal over xAI's turbine use, citing 'mootness'," December 15, 2025.

[16] Mississippi Today, "xAI now has 46 gas turbines without air permits," May 11, 2026.

[17] Mississippi Today, "State sets dates to retire temporary xAI turbines but allows some to go past original deadline," July 31, 2026.

[18] Mississippi Free Press, "Mississippi Permit Board Grants xAI's Request for 41 Southaven Gas Turbines," March 2026.

[19] The Commercial Appeal, "xAI Southaven power plant operating 46 unpermitted turbines, MDEQ says," May 13, 2026.

[20] Memphis Flyer, "xAI Adds Additional Turbines to Southaven Site," June 2026.

[21] NAACP, "NAACP Sues xAI for Illegal Pollution from Data Center Power Plant," April 14, 2026.

[22] The Next Web, "xAI is running 59 unpermitted gas turbines, twice what it has admitted," July 14, 2026 (citing Reuters).

[23] United States, Motion for Intervention and Dismissal and Memorandum in Support, *NAACP v. X.AI*, No. 3:26-cv-00074 (N.D. Miss. filed June 15, 2026), Docs. 58–59.

[24] U.S. Department of Justice, "Justice Department Files to Intervene and Dismiss Lawsuit that Would Hamper America's AI Innovation and Security," press release no. 26-659, June 16, 2026.

[25] 40 C.F.R. § 60.4331a, "What are the requirements for operating a stationary temporary combustion turbine?"

[26] Earthjustice, "Trump Administration Attempts Massive Power Grab in Defense of Musk's Company, xAI," July 20, 2026.

[27] TechCrunch, "DOJ claims xAI's unpermitted gas turbines are a matter of 'national, economic, and energy security'," June 16, 2026.

[28] Utility Dive, "DOJ intervenes on behalf of xAI in data center gas turbine lawsuit," June 18, 2026.

[29] U.S. Environmental Protection Agency, "New Source Performance Standards Review for Stationary Combustion Turbines and Stationary Gas Turbines," *Federal Register* 91 (January 15, 2026): 1910; correction, July 16, 2026.

[30] Clark Hill, "EPA's New Turbine Rules Ease Air Permits for Data Centers," May 2026.

[31] Flint Water Advisory Task Force, *Final Report*. Lansing: Office of the Governor, March 21, 2016.

[32] U.S. EPA Office of Inspector General, *Management Weaknesses Delayed Response to Flint Water Crisis*, Report No. 18-P-0221, July 19, 2018.

[33] M. Del Toral, email correspondence, April 27, 2015, released by the U.S. House Committee on Oversight and Government Reform, 2016.

[34] M. Hanna-Attisha, J. LaChance, R. C. Sadler, and A. C. Schnepp, "Elevated Blood Lead Levels in Children Associated with the Flint Drinking Water Crisis," *American Journal of Public Health* 106(2), 2016, 283–290.

[35] M. Hanna-Attisha, statement to the U.S. House Committee on Oversight and Government Reform, February 10, 2016.

[36] Review of A. Clark, *The Poisoned City*, and M. Hanna-Attisha, *What the Eyes Don't See*, *Nature*, 2018.

[37] Patch, "State Delayed Action on Lead in Flint Water: Records," 2015 (reporting records obtained by the ACLU of Michigan and M. Edwards).

[38] Michigan Civil Rights Commission, *The Flint Water Crisis: Systemic Racism Through the Lens of Flint*, February 17, 2017.

[39] M. M. Mello, G. Persad, and D. B. White, "Respecting Disability Rights: Toward Improved Crisis Standards of Care," *New England Journal of Medicine* 383(5), 2020, e26. doi:10.1056/NEJMp2011997.

[40] U.S. Department of Health and Human Services, Office for Civil Rights, resolutions with Alabama (April 8, 2020), Pennsylvania (April 16, 2020), Tennessee (June 26, 2020), and Utah (August 20, 2020).

[41] D. C. Ashana, G. L. Anesi, V. X. Liu, et al., "Equitably Allocating Resources during Crises: Racial Differences in Mortality Prediction Models," *American Journal of Respiratory and Critical Care Medicine* 204(2), 2021, 178–186.

[42] B. Tolchin et al., "Racial Disparities in the SOFA Score Among Patients Hospitalized with COVID-19," *PLOS ONE* 16(9), 2021, e0257608.

[43] W. D. Miller, X. Han, M. E. Peek, D. C. Ashana, and W. F. Parker, "Accuracy of the Sequential Organ Failure Assessment Score for In-Hospital Mortality by Race and Relevance to Crisis Standards of Care," *JAMA Network Open* 4, 2021.

[44] C. Jia and J. Foran, "Air Toxics Concentrations, Source Identification, and Health Risks: An Air Pollution Hot Spot in Southwest Memphis, TN," *Atmospheric Environment* 81, 2013, 112–116.

[45] H. Schmidt, D. E. Roberts, and N. D. Eneanya, "Sequential Organ Failure Assessment, Ventilator Rationing and Evolving Triage Guidance: New Evidence Underlines the Need to Recognise and Revise, Unjust Allocation Frameworks," *Journal of Medical Ethics* 48(2), 2022, 136–138.

[46] H. Schmidt, D. E. Roberts, and A. D. Eneanya, "The Sequential Organ Failure Assessment (SOFA) and Revised Ventilator Triage Guidance," *Journal of Medical Ethics* blog, October 12, 2021.

[47] Abstract on state crisis standards of care and SOFA use in January 2022, *BMJ Global Health* 7, Suppl. 2, 2022, A16.