Article

Not peer-reviewed version

Structural Ethical Infeasibility in AI-Enabled Infrastructure Systems: A Constraint-Based Diagnostic Framework

Sudipta Chowdhury * , Md Abdul Quddus , Ammar Alzarrad

Posted Date: 21 May 2026

doi: 10.20944/preprints202605.1432.v1

Keywords: ethical AI; infrastructure systems; constraint programming; algorithmic fairness; irre-ducible infeasible subsystem; emergency medical services; equity; demographic parity

Article

Structural Ethical Infeasibility in AI-Enabled Infrastructure Systems: A Constraint-Based Diagnostic Framework

Sudipta Chowdhury [1,*], Md Abdul Quddus [2] and Ammar Alzarrad [3]

[1] Assistant Professor, Mechanical and Industrial Engineering, Marshall University, Huntington, WV 25755
[2] Assistant Professor, Textile Engineering, Chemistry and Science, NC State University, Raleigh, NC 27606
[3] Associate Professor, Civil Engineering, Marshall University, Huntington, WV 25755
* Correspondence: chowdhurys@marshall.edu

Abstract

AI-enabled infrastructure systems increasingly govern access to emergency services, dis-aster relief, and utility restoration, yet they routinely produce inequitable outcomes even when allocation algorithms apply procedurally neutral rules. The standard explanation locates the cause inside the algorithm. This paper argues instead that inequity arises from the interaction between the algorithm and the physical environment in which it operates: network topology, resource locations, and demand distribution jointly constrain what any policy can achieve, and when those constraints are sufficiently binding, ethical infeasibility is structural rather than algorithmic. We introduce a constraint-based formulation that embeds ethical requirements into the feasible region, and a hierarchical Irreducible Infea-sible Subsystem (IIS) procedure that attributes infeasibility to rule design, algorithmic choice, or physical infrastructure. We further establish the Structural Infeasibility Theorem, deriving closed-form bounds on inter-group disparity across all feasible policies in zone-decomposable problems. Applied to a metropolitan ambulance-dispatch instance, the framework certifies minimum-service infeasibility as infrastructural, shows the efficiency–equity trade-off to be an artifact of constrained infrastructure, identifies pre-investment equity gains as harm redistribution rather than harm reduction, and converts the IIS certificate into a quantified capital-investment specification.

Keywords: ethical AI; infrastructure systems; constraint programming; algorithmic fairness; irreducible infeasible subsystem; emergency medical services; equity; demographic parity

---

1. Introduction

AI-enabled optimization models now routinely govern the allocation of critical infrastructure services, including emergency medical dispatch [1,21,32], post-disaster shelter assignment [2,39], and utility restoration [3,10,22]. Concurrent with their operational adoption, these systems have become subject to ethical-governance requirements mandating fairness, accountability, and human oversight [14,15,18,28,35,36]. A persistent empirical finding, however, complicates the implementation of such requirements: infrastructure service allocation outcomes exhibit systematic disparities across racial and socioeconomic groups, and these disparities persist even when the governing decision rules apply identical criteria to every demand unit [11,19,23,29,33,38–40]. That is, unequal outcomes arise not because the algorithm treats different groups differently, but because neutral rules interact with a physical and historical infrastructure that was itself built unequally [4,5,16,34,43].

The majority of the existing technical literature has responded to fairness conflict in AI systems along two largely independent lines. The algorithmic-fairness literature formalizes fairness as a property of decision rules and establishes impossibility results showing that several natural fairness criteria cannot be simultaneously satisfied within a classifier [13,17,27,30,41,42]. These results characterize fairness conflict as an internal property of the rule set: criteria are jointly infeasible because they are logically inconsistent with one another. The operations research literature independently develops equity-oriented optimization formulations [6,7,20,26,31,37], but these typically encode equity as a single scalar metric and do not provide structured procedures for attributing the source of infeasibility when ethical requirements cannot be jointly satisfied. Both literatures share a common implicit assumption: the physical infrastructure is treated as a fixed background parameter, and the locus of ethical failure is located in the algorithm or the rule set operating on it.

This paper relaxes that assumption and develops its consequences. The two literatures reviewed above share a common diagnostic orientation: when an infrastructure AI system produces inequitable outcomes, the search for the cause begins and ends with the algorithm. We argue that this orientation is structurally incomplete. In infrastructure allocation systems, the physical configuration of the network, including the location of facilities, the topology of service coverage, and the spatial distribution of demand across demographic groups, jointly determines what outcomes are achievable under any allocation policy whatsoever. These constraints are not properties of the algorithm; they are properties of the built environment in which the algorithm operates. An allocation algorithm that is fairness-compliant in one infrastructure configuration may be structurally incapable of achieving compliance in another, not because it has been poorly designed, but because the physical system does not possess the capacity to support the ethical commitments imposed upon it.

We formalize this observation as a distinct mechanism of ethical infeasibility. In zone-decomposable allocation problems, the joint structure of network topology, resource geography, and group-level demand composition imposes hard bounds on the inter-group disparity attainable across all feasible allocation policies. These bounds are closed-form functions of infrastructure parameters alone and hold regardless of algorithmic design. When they preclude compliance with a normatively specified fairness tolerance, the system is ethically infeasible in a strong sense: no policy revision, no matter how carefully constructed, can restore compliance. The source of that infeasibility is invisible to any audit confined to the allocation algorithm, because the binding constraints lie not in the decision rule but in the physical system on which the rule operates. Diagnosing it correctly requires a framework capable of attributing infeasibility to its actual source, which may be the rule set, the algorithm, or the infrastructure itself, and of translating that attribution into the appropriate class of intervention.

We also argue that the central implication extends beyond methodology. If infrastructure topology can render ethical requirements structurally infeasible, then the standard governance model, in which ethical compliance is achieved by constraining or auditing the allocation algorithm, is insufficient as a general approach to ethical infrastructure AI. The physical infrastructure on which the algorithm operates must itself be evaluated for its capacity to support the ethical commitments imposed upon it, and when that capacity is found wanting, the appropriate response is not algorithmic revision but investment in the physical system.

Overall, the paper develops three connected contributions, each oriented to the question of when ethical infeasibility is algorithmic versus structural.

First, we introduce a minimal constraint-based formulation that augments a standard ethically-agnostic allocation problem with hard ethical constraints, soft Pareto objectives, and state-dependent rules. We then specify a hierarchical infeasibility-diagnosis procedure based on Irreducible Infeasible Subsystem (IIS) analysis [8,12,24] that distinguishes four conflict types and attributes each to one of three sources: rule design, algorithmic choice, or physical infrastructure. The three-way attribution is the diagnostic core of the framework, because each source requires a categorically different intervention.

Second, we establish the Structural Infeasibility Theorem, which derives closed-form bounds on the inter-group disparity attainable across all feasible allocation policies in any zone-decomposable allocation problem. The bounds are determined entirely by the infrastructure topology, i.e., the travel-time matrix, the spatial distribution of demand, and the group-composition profile; and not by any algorithmic design choice. An explicit corollary identifies the one-sided-interval condition under

 doi:10.20944/preprints202605.1432.v1

which these bounds imply demographic-parity infeasibility below a critical tolerance regardless of the dispatch policy. The corollary recovers, as a special case, the parity-infeasibility result that obtains in stylized two-group, two-depot configurations, while clarifying that the structural-floor mechanism does not generalize automatically to multi-group, multi-depot settings.

Third, We demonstrate the framework on a metropolitan ambulance-dispatch instance with eight demand zones, three depots, and three protected groups. The case study produces four substantive findings.

(i) The legal minimum response time standard (NFPA 1710) is certified as structurally infeasible [9]: the IIS procedure establishes that no dispatch policy can meet the standard for two southern zones because no existing depot is physically close enough to serve them within the required threshold, and no algorithmic adjustment can change that fact.

(ii) Demographic parity at the oversight board's tolerance is nonetheless feasible on this instance, because the geometric conditions required for a structural parity floor do not hold in the multi-group, multi-depot setting, a result that distinguishes this regime from stylized two-group, two-depot configurations and that the Structural Infeasibility Theorem makes precise.

(iii) Along the pre-investment Pareto frontier, equity gains are achieved entirely through harm redistribution rather than harm reduction: reducing inter-group disparity within the existing infrastructure requires slowing the historically faster-served group rather than improving service to the slower-served group, so that all population groups experience worse absolute outcomes under the equity-compliant policy than under the operationally optimal baseline.

(iv) The IIS certificate is converted into a quantified capital-investment specification identifying the minimum siting requirement for a fourth depot that simultaneously resolves the minimum-service violation, satisfies the demographic-parity constraint, and meets the disaster-state response standard, while also improving mean system-wide response time. This finding carries a direct implication for finding (iii): the efficiency-equity trade-off observed along the pre-investment frontier is not a fundamental property of equitable allocation but an artifact of insufficient infrastructure capacity that disappears once the binding physical constraint is removed.

Ultimately, these results imply a methodological shift in how ethical AI for infrastructure should be evaluated. Algorithm-only auditing is structurally incapable of identifying infrastructure-induced infeasibility; the infrastructure on which the algorithm operates must itself be subjected to formal ethical-feasibility analysis. The exposition of this paper is as follows: Section 2 introduces the minimal Ethical constraint programming (ECP) formulation, the three-tier rule structure, and the four conflict types. Section 3 develops the hierarchical IIS-based diagnosis procedure and the three-way attribution mechanism. Section 4 states and proves the Structural Infeasibility Theorem and its parity-infeasibility corollary. Section 5 applies the framework to the metropolitan ambulance-dispatch instance, presenting the IIS attribution, Pareto analysis, harm-redistribution finding, and capital-investment recommendation. Section 6 discusses implications for the auditing of AI-enabled infrastructure systems and outlines extensions. Section 7 concludes.

2. A Minimal Constraint-Based Formulation

This section introduces the minimal formulation needed to state and prove the structural-infeasibility result. More specifically, we develop the constructs required for the diagnostic procedure (Section 3) and the theorem (Section 4).

2.1. Notation

Let $\mathcal{I}$ denote the set of individuals or demand units and $\mathcal{J}$ the set of infrastructure locations (depots, shelters, or service facilities). A decision $\mathbf{x}$ drawn from a feasible decision space $\mathcal{X}$ determines, for each individual $i \in \mathcal{I}$, a quantifiable service outcome $y_i(\mathbf{x}) \in \mathbb{R}_{\geq 0}$. The population is partitioned into protected groups indexed by $\mathcal{G}$, with $\bigcup_{g \in \mathcal{G}} \mathcal{I}_g = \mathcal{I}$ and $\mathcal{I}_g \cap \mathcal{I}_{g'} = \emptyset$ for $g \neq g'$. The partition will be normatively specified by stakeholders prior to deployment, i.e., the framework does not infer it from data. The group mean outcome can be represented by:

$$\bar{y}_g(\mathbf{x}) = \frac{1}{|\mathcal{I}_g|} \sum_{i \in \mathcal{I}_g} y_i(\mathbf{x}) \tag{1}$$

which is the primary quantity through which distributional inequity is detected. Operational performance is captured by an objective vector $\mathbf{F} : \mathcal{X} \to \mathbb{R}^K$, with components $f_k(\mathbf{x})$ representing metrics such as mean response time, restoration speed, or aggregate cost. A system state $s \in \mathcal{S}$ encodes the current operational context (e.g., NORMAL, SURGE, DISASTER).

2.2. The Standard Allocation Problem and the ECP-Augmented Form

Let $\mathcal{G}_{\text{op}} \subseteq \mathcal{X}$ denote the set of decisions satisfying the operational constraints (capacity, coverage, routing, scheduling, or conservation). The standard ethically-agnostic allocation problem can be mathematically represented as:

$$\min_{\mathbf{x} \in \mathcal{G}_{\text{op}} \cap \mathcal{X}} \mathbf{F}(\mathbf{x}), \tag{P0}$$

where minimization is in the Pareto sense when $K \geq 2$. Problem (P0) is the formal object implemented by current optimization-based infrastructure AI: it is silent on the distribution of $y_i(\mathbf{x})$ across the population and therefore cannot certify ethical compliance.

The ECP form augments (P0) with a normatively specified ethical rule set

$$\mathcal{R} = \mathcal{H} \cup \dot{\mathcal{S}} \cup \mathcal{D},$$

partitioned into hard constraints, soft Pareto objectives, and dynamic state-dependent rules respectively. Hard rules $r \in \mathcal{H}$ contribute constraints $e_r(\mathbf{x}) \leq 0$ that any ethically admissible decision must satisfy. Soft rules $r \in \dot{\mathcal{S}}$ contribute objectives $e_r^S(\mathbf{x})$ to be minimized on a Pareto frontier alongside $\mathbf{F}$. Dynamic rules $r \in \mathcal{D}$ contribute state-conditional constraints $e_r(\mathbf{x}, s) \leq 0$ that activate only when $s \in \mathcal{S}_r$. The ECP problem hence can be mathematically formulated as:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \big(\mathbf{F}(\mathbf{x}),\, \mathbf{E}^S(\mathbf{x})\big) \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{G}_{\text{op}} \cap \mathcal{X}, \\ & e_r(\mathbf{x}) \leq 0, \quad \forall r \in \mathcal{H}, \\ & e_r(\mathbf{x}, s) \leq 0, \quad \forall r \in \mathcal{D}(s), \end{aligned} \tag{ECP}$$

where $\mathcal{D}(s) = \{r \in \mathcal{D} : s \in \mathcal{S}_r\}$ and $\mathbf{E}^S$ stacks the soft objectives into a vector. The feasible region under state s is

$$\Omega(s) = \big\{\mathbf{x} \in \mathcal{G}_{\text{op}} \cap \mathcal{X} : e_r(\mathbf{x}) \leq 0 \; \forall r \in \mathcal{H},\; e_r(\mathbf{x}, s) \leq 0 \; \forall r \in \mathcal{D}(s)\big\}. \tag{2}$$

The central diagnostic question of this paper is to determine whether $\Omega(s)$ is empty, and if so, which constraints are responsible.

2.3. Three-Tier Rule Set

The hard tier $\mathcal{H}$ contains rules whose violation renders a decision ethically inadmissible regardless of operational performance. We primarily focus on two such rules. The demographic-parity constraint requires inter-group mean outcomes to agree within a tolerance δ:

$$|\bar{y}_g(\mathbf{x}) - \bar{y}_{g'}(\mathbf{x})| \leq \delta, \quad \forall g, g' \in \mathcal{G}. \tag{HC-DP}$$

 doi:10.20944/preprints202605.1432.v1

The minimum-service guarantee requires each individual to receive a baseline service level (e.g., NFPA 1710 mandates Advanced Life Support response within $\theta = 480$ seconds in urban areas):

$$y_i(\mathbf{x}) \geq \theta_i, \quad \forall i \in \mathcal{I}. \tag{HC-MS}$$

The soft tier $\dot{\mathcal{S}}$ contains objectives admitting legitimate trade-offs against operational efficiency. We primarily focus on Gini-type dispersion (SO-EQ), a vulnerability-weighted welfare $E_{\text{vw}}(\mathbf{x}) = -\sum_i v_i\, y_i(\mathbf{x})$ (SO-VW), and a Rawlsian welfare $E_{\text{rw}}(\mathbf{x}) = -\min_i y_i(\mathbf{x})$ (SO-RW). The dynamic tier $\mathcal{D}$ contains state-conditional constraints implemented via the big-M formulation

$$e_r(\mathbf{x}, s) \leq M_r\big(1 - \Phi_r(s)\big), \quad \forall r \in \mathcal{D}, \quad \Phi_r(s) = \mathbf{1}[s \in \mathcal{S}_r], \tag{3}$$

which becomes binding when $\Phi_r(s) = 1$ and is otherwise non-restrictive.

2.4. Four Conflict Types

Infeasibility of $\Omega(s)$ may arise through four mechanisms. A *Type I (rule–rule)* conflict arises when two hard rules are internally inconsistent: $\Omega(\eta_r) \cap \Omega(\eta_{r'}) = \emptyset$ for some $(r, r') \in \mathcal{H} \times \mathcal{H}$, where $\Omega(\eta_r) = \{\mathbf{x} \in \mathcal{X} : e_r(\mathbf{x}) \leq 0\}$. This is the formal structure of the classifier-fairness impossibility applied to current problem setting [13,30]. A *Type II (rule–operational)* conflict arises when the hard rules cannot be jointly satisfied within the operational and infrastructural region: $\bigcap_{r\in\mathcal{H}} \Omega(\eta_r) \cap \mathcal{G}_{\text{op}} = \emptyset$. The infeasibility in this case is not a property of the rule set but of its interaction with the physical system. A *Type III (soft–soft)* conflict is not a failure: it is the expected Pareto trade-off among $(\mathbf{F}, \mathbf{E}^{S})$ on a non-empty $\Omega(s)$. A *Type IV (state-conditional)* conflict arises when $\Omega(s_1) \neq \emptyset$ but $\Omega(s_2) = \emptyset$ for some $s_1, s_2 \in \mathcal{S}$, capturing latent infeasibilities that activate only under specific operational states.

3. Hierarchical IIS Diagnosis and Three-Way Attribution

When $\Omega(s) = \emptyset$, the framework will execute a three-stage hierarchical procedure based on IIS analysis. An IIS is a minimal subset of constraints that is jointly infeasible but becomes feasible upon removal of any single constraint [12,25]; modern mixed-integer solvers (Gurobi, CPLEX, SCIP) extract IISs by deletion-filtering heuristics that terminate in a number of feasibility solves linear in the constraint count.

Stage 1 will test pairwise hard-rule feasibility: for each $(r, r') \in \mathcal{H} \times \mathcal{H}$, solve the feasibility problem $\mathbf{x} \in \mathcal{G}_{\text{op}} \cap \mathcal{X}$, $e_r(\mathbf{x}) \leq 0$, $e_{r'}(\mathbf{x}) \leq 0$. A positive infeasibility finding identifies a Type I conflict: the rule set is internally inconsistent and must be revised through governance. Stage 2 will test joint hard-rule feasibility: solve $\mathbf{x} \in \Omega(s)$ under all of $\mathcal{H} \cup \mathcal{D}(s)$. If infeasible, extract the IIS by deletion-filtering. Stage 3 will test state-conditional feasibility for each anticipated state $s' \in \mathcal{S}$, identifying Type IV exposures.

Once an IIS $\text{IIS}^* \subseteq \mathcal{H} \cup \mathcal{D}(s) \cup \mathcal{G}_{\text{op}}$ is extracted, its content is partitioned by origin into constraints drawn from the ethical rule set $\mathcal{H} \cup \mathcal{D}(s)$ and constraints drawn from the operational and infrastructure set $\mathcal{G}_{\text{op}}$, with the latter further distinguished by whether the binding parameter reflects an algorithmic restriction or a physical system property.

- *Rule-induced* infeasibility, when IIS^* contains multiple constraints from $\mathcal{H}$ alone, indicating a Type I inconsistency. The intervention is rule revision through governance.
- *Algorithm-induced* infeasibility, when IIS^* contains constraints arising from algorithmic restrictions on $\mathcal{X}$ (e.g., fixed assignment heuristics, nearest-depot policies). The intervention is algorithmic relaxation or reformulation.
- *Infrastructure-induced* infeasibility, when IIS^* consists primarily of operational and physical-system parameters in $\mathcal{G}_{\text{op}}$ (capacities, distances, network topology) jointly incompatible with $\mathcal{H}$. The intervention is capital investment in the physical infrastructure.

The three-way attribution is the most consequential analytical output of the diagnosis, because each cause requires a categorically different intervention. Algorithm-only auditing of an AI-enabled

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 21 May 2026 doi:10.20944/preprints202605.1432.v1

infrastructure system can detect Type I and algorithm-induced Type II conflicts, but is structurally incapable of identifying infrastructure-induced infeasibility: the IIS for such a conflict contains no algorithmic constraint that any auditor of the algorithm alone is empowered to examine. The remedy when IIS* attributes infeasibility to physical infrastructure is therefore not algorithmic but constructive: ECP converts the IIS into a quantified infrastructure-investment specification, identifying the threshold change in physical configuration required to restore $\Omega(s) \neq \emptyset$.

Note that the partition of constraints into the categories of rule design, algorithmic choice, and physical infrastructure is not derived mechanically from the IIS itself. The IIS returns a minimal jointly-infeasible subsystem; determining which constraints within that subsystem reflect a policy choice, an algorithmic restriction, or a physical system property requires contextual knowledge of how the ECP instance was constructed. This interpretive step is explicit and auditable: the governance log preserves both the IIS certificate and the partition applied to it, so that the attribution can be reviewed, contested, and revised as institutional knowledge of the system develops. Different plausible partitions of the same IIS can yield different attributions and therefore different intervention recommendations, which is precisely why the partition must be recorded and not treated as an automatic output of the diagnostic procedure.

4. The Structural Infeasibility Theorem

The IIS procedure of Section 3 identifies which constraints are jointly responsible for infeasibility on a given instance, but does not characterize the structure of that infeasibility across all feasible policies. For the class of zone-decomposable allocation problems, the Type II infeasibility mechanism admits a closed-form characterization that holds simultaneously for every allocation policy in $\mathcal{X}$, with bounds determined entirely by infrastructure topology rather than by algorithmic design. This characterization matters for two reasons. First, it allows a governing institution to determine whether compliance with a fairness tolerance is structurally achievable across all possible policies, prior to any optimization. Second, it identifies precisely which infrastructure parameters drive the disparity bounds, providing a principled basis for determining where physical investment would be most effective in restoring feasibility. This section states and proves that result and develops its parity-infeasibility corollary.

4.1. Setting

Consider an ECP instance in which (i) the decision space encodes a zone-to-facility assignment, with each zone z in a discrete index set $\mathcal{Z}$ assigned to a single facility $d \in \mathcal{J}$ via binary indicators $a_{dz} \in \{0,1\}$ satisfying $\sum_{d\in\mathcal{J}} a_{dz} = 1$; (ii) individual outcomes are response or service times $r_z(\mathbf{x}) = \sum_d a_{dz}\, T[d,z]$, where $T[d,z] \in \mathbb{R}_{\geq 0}$ is the travel time from facility d to zone z; (iii) the group-composition matrix $\phi_{g,z} \in [0,1]$ specifies the fraction of zone-z demand attributable to group g, with $\sum_g \phi_{g,z} = 1$ for all z; (iv) zone demand intensities are $\lambda_z \geq 0$ with group totals $\lambda_g = \sum_z \phi_{g,z}\,\lambda_z$; and (v) the operational set $\mathcal{G}_{\text{op}}$ imposes only the per-zone assignment condition. Group mean outcomes follow

$$\bar{y}_g(\mathbf{x}) \;=\; \frac{1}{\lambda_g} \sum_{z\in\mathcal{Z}} \phi_{g,z}\,\lambda_z\, r_z(\mathbf{x}). \tag{4}$$

4.2. Statement and Proof

Theorem 1 (Structural Infeasibility Theorem: signed-disparity bounds)**.** *Under the setting of Section 4.1, define for each pair $(g,g') \in \mathcal{G}\times\mathcal{G}$ the disparity coefficient*

$$c_{g,g',z} \;=\; \left(\frac{\phi_{g,z}}{\lambda_g} - \frac{\phi_{g',z}}{\lambda_{g'}}\right)\lambda_z, \tag{5}$$

 doi:10.20944/preprints202605.1432.v1

and the achievable signed-disparity bounds

$$L_{g,g'} = \sum_{z \in \mathcal{Z}} \begin{cases} c_{g,g',z} \min_d T[d,z], & c_{g,g',z} \geq 0, \\ c_{g,g',z} \max_d T[d,z], & c_{g,g',z} < 0, \end{cases} \tag{6}$$

$$U_{g,g'} = \sum_{z \in \mathcal{Z}} \begin{cases} c_{g,g',z} \max_d T[d,z], & c_{g,g',z} \geq 0, \\ c_{g,g',z} \min_d T[d,z], & c_{g,g',z} < 0. \end{cases} \tag{7}$$

Then for every $\mathbf{x} \in \mathcal{X}$*,*

$$L_{g,g'} \leq \bar{y}_g(\mathbf{x}) - \bar{y}_{g'}(\mathbf{x}) \leq U_{g,g'}, \tag{8}$$

and both bounds are tight: each is attained by the assignment that selects, in every zone, the facility achieving the relevant per-zone extremum.

Proof. By (4),

$$\bar{y}_g(\mathbf{x}) - \bar{y}_{g'}(\mathbf{x}) = \sum_{z \in \mathcal{Z}} c_{g,g',z}\, r_z(\mathbf{x}).$$

Because $\mathcal{G}_{\text{op}}$ imposes no cross-zone coupling beyond the per-zone assignment condition, each summand is a function of an independent zone-level decision: zone z selects a single $d_z \in \mathcal{J}$ and contributes $c_{g,g',z}\, T[d_z, z]$. The minimum of the sum equals the sum of per-zone minima, and similarly for the maximum. For each zone, the per-zone minimum of $c_{g,g',z}\, T[d,z]$ over $d \in \mathcal{J}$ is attained at $\arg\min_d T[d,z]$ when $c_{g,g',z} \geq 0$ and at $\arg\max_d T[d,z]$ when $c_{g,g',z} < 0$; the maximum case is symmetric. Summing per-zone extrema yields $L_{g,g'}$ and $U_{g,g'}$ as in (6)–(7), and tightness follows from constructive achievability. □

Corollary 1 (Parity infeasibility under one-sided intervals)**.** *If for some pair* (g, g') *the achievable signed-disparity interval* $[L_{g,g'}, U_{g,g'}]$ *does not contain zero, then for every* $\mathbf{x} \in \mathcal{X}$*,*

$$\left|\bar{y}_g(\mathbf{x}) - \bar{y}_{g'}(\mathbf{x})\right| \geq \min\left(|L_{g,g'}|, |U_{g,g'}|\right).$$

Consequently, (HC-DP) *at any tolerance* $\delta < \min(|L_{g,g'}|, |U_{g,g'}|)$ *is infeasible across all allocation policies, and the infeasibility is attributable to the travel-time matrix* T*, not to the dispatch algorithm. If* $[L_{g,g'}, U_{g,g'}]$ *contains zero, no such infeasibility implication follows from Theorem 1 alone.*

Note that the theorem addresses HC-DP specifically because demographic parity is the only hard constraint in the ECP rule set whose feasibility depends on the global interaction of network topology, demand distribution, and group composition across all policies simultaneously. Whether two group means can be brought within a tolerance δ of each other is not a local question: it depends on how group membership is distributed across every zone, how much demand each zone generates, and what response times are achievable from each facility, all interacting at once. This global structure is precisely what the disparity coefficient $c_{g,g',z}$ captures and what makes a closed-form bound both non-trivial and necessary.

The remaining hard constraint has a different character and is handled by the IIS procedure directly. HC-MS infeasibility reduces to a local per-zone check: zone z is infeasible if and only if $\min_d T[d,z] > \theta$, a condition that can be verified independently for each zone without reference to group composition or cross-zone interactions. The soft objectives SO-EQ, SO-RW, and SO-VW occupy a different position in the framework. Because they are Tier II objectives rather than hard constraints, their conflict structure is Type III: they generate Pareto trade-offs against operational performance rather than infeasibility. The theorem does not apply to them directly. However, the disparity bounds $[L_{g,g'}, U_{g,g'}]$ established by the theorem do constrain the achievable Pareto frontier over these objectives: no policy can achieve a combination of SO-EQ, SO-RW, or SO-VW values that would require group mean outcomes outside the bounds guaranteed by Theorem 1.

4.3. Two Regimes

The practical implication of Theorem 1 depends on whether the achievable signed-disparity interval $[L_{g,g'}, U_{g,g'}]$ contains zero. This is determined by the signs of the disparity coefficients $c_{g,g',z}$ across zones, which in turn reflect the joint structure of group composition, demand intensity, and network topology. Two qualitatively distinct regimes arise.

In the first regime, which corresponds to stylized configurations with two groups, two facilities, and binary group composition $\phi_{g,z} \in \{0,1\}$, the disparity coefficients $c_{g,g',z}$ carry the same sign across all zones. Intuitively, this occurs when one group is spatially concentrated in zones that are systematically closer to one set of facilities, and the other group is concentrated in zones that are systematically closer to another. When all coefficients share a sign, the signed-disparity interval $[L_{g,g'}, U_{g,g'}]$ lies entirely on one side of zero, and Corollary 1 applies: the absolute disparity between the two groups is bounded away from zero across every feasible policy, and demographic parity below the implied floor is structurally infeasible regardless of the dispatch algorithm. This is the regime in which infrastructure topology directly and unavoidably produces parity infeasibility.

In the second regime, which corresponds to settings with three or more groups, three or more facilities, and mixed group composition $0 < \phi_{g,z} < 1$, the disparity coefficients $c_{g,g',z}$ typically take both positive and negative values across zones. This occurs because each zone serves multiple groups simultaneously, so the assignment of a zone to a nearby facility improves outcomes for all groups present in that zone rather than benefiting one group at the expense of another. When the coefficients take mixed signs, the signed-disparity interval $[L_{g,g'}, U_{g,g'}]$ contains zero, and Corollary 1 delivers no infeasibility implication: demographic parity is not ruled out by the theorem, and may be achievable despite severe operational asymmetries, though potentially at significant operational cost.

The distinction between these two regimes is substantive and not obvious without the theorem. A planner or auditor working with a multi-group, multi-depot system might incorrectly assume that the parity-infeasibility result familiar from stylized two-group configurations carries over to their setting. The theorem establishes that it does not, and provides the analytical conditions under which parity remains structurally achievable. The case study of Section 5 demonstrates this distinction concretely: the metropolitan ambulance instance falls into the second regime, where minimum-service is structurally infeasible but demographic parity at the regulator's tolerance remains feasible, a finding that the theorem makes precise and that algorithm-only auditing cannot produce.

5. Case Study: Metropolitan Ambulance Dispatch

The case study exercises the diagnostic procedure of Section 3 and the theorem of Section 4 on a metropolitan emergency medical services (EMS) instance. The substantive findings are: a Type II infeasibility on the minimum-service rule (HC-MS) certified as infrastructural; a substantive failure of Corollary 1's one-sided-interval condition that renders demographic parity feasible despite operational asymmetry; a Pareto frontier exhibiting a harm-redistribution pattern; and a quantified capital-investment specification that simultaneously resolves four binding constraints while improving operational performance.

Every numerical claim below is produced along an analytical path. It uses the framework's prescribed procedures: per-zone feasibility checks, Stage 2 IIS attribution, ε-constraint Pareto computation, and the closed-form bounds of Theorem 1. The verification path enumerates the 3^8 assignments in $\mathcal{X}$ (and, for the post-investment analysis, the $4^8 = 65{,}536$ assignments in the extended space) and evaluates rule violations and operational objectives at each point.

5.1. Instance Specification

A stylized, hypothetical metropolitan EMS system serving approximately 280,000 residents was considered. The service area was partitioned into eight demand zones along a north–south spatial gradient, $\mathcal{Z} = \{z_1, \ldots, z_8\}$, served by three depots $\mathcal{J} = \{d_1, d_2, d_3\}$. Depot d_1 (north) and d_2 (central) represented historically established service centers, while d_3 was located in the southwest. Demand

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 21 May 2026 doi:10.20944/preprints202605.1432.v1

was concentrated in southern zones, reflecting population growth and demographic change; no additional depot construction was assumed. Three protected groups $\mathcal{G} = \{\text{Maj}, \text{Black}, \text{Hisp}\}$ and three operational states $\mathcal{S} = \{\text{NORMAL}, \text{SURGE}, \text{DISASTER}\}$ were defined, with the group partition normatively specified by the oversight board.

The decision variable was $\mathbf{x} = (a_{dz}) \in \{0,1\}^{|\mathcal{J}|\cdot|\mathcal{Z}|}$ with $\mathcal{G}_{\text{op}}$ imposing $\sum_d a_{dz} = 1$ for each zone, yielding $|\mathcal{J}|^{|\mathcal{Z}|} = 3^8 = 6{,}561$ admissible assignments. The instance was evaluated at $s = \text{NORMAL}$, with state transitions considered through the dynamic-rule analysis below.

Travel times $T[d,z]$ in minutes are reported in Table 1. The topology drives much of the subsequent analysis: northern and central zones are efficiently covered by d_1 and d_2, while the easternmost southern zones z_7, z_8 are relatively distant from all depots, including d_3.

Table 1. Travel-time matrix $T[d,z]$ in minutes.

Depot	z_1	z_2	z_3	z_4	z_5	z_6	z_7	z_8
d_1 (North)	3	4	5	8	9	18	20	22
d_2 (Central)	7	6	8	3	4	12	14	16
d_3 (Southwest)	14	13	15	9	8	7	9	11

Call rates in calls per hour were $\lambda = (4,3,3,5,4,4,3,2)$ with total $\Lambda = 28$. The group-composition matrix $\phi_{g,z}$ (Table 2) gave the fraction of zone-z calls attributable to each group. Group call rates were $\lambda_{\text{Maj}} = 14.95$, $\lambda_{\text{Black}} = 8.14$, and $\lambda_{\text{Hisp}} = 4.91$. Vulnerability scores $v_z \in [0,1]$ were $v = (0.15, 0.20, 0.25, 0.40, 0.45, 0.75, 0.80, 0.78)$, rising sharply across the north–south gradient.

Table 2. Group composition $\phi_{g,z}$ for the case-study instance. Each column sums to one.

Group	z_1	z_2	z_3	z_4	z_5	z_6	z_7	z_8
Majority	0.80	0.85	0.75	0.60	0.55	0.20	0.15	0.25
Black	0.10	0.08	0.15	0.25	0.20	0.55	0.60	0.50
Hispanic	0.10	0.07	0.10	0.15	0.25	0.25	0.25	0.25

The zone-level outcome is $y_z(\mathbf{x}) = r_z(\mathbf{x}) = \sum_d a_{dz}\, T[d,z]$. Group means follow Eq. (4). The demand-weighted system mean is $\bar{R}(\mathbf{x}) = \Lambda^{-1} \sum_z \lambda_z r_z(\mathbf{x})$, and the operational objective vector is $\mathbf{F}(\mathbf{x}) = (\bar{R}(\mathbf{x}), \max_z r_z(\mathbf{x}))^\top$.

5.2. Baseline P0 Audit

Solving (P0) under $\mathbf{F}(\mathbf{x})$ alone, subject only to operational constraints, yields the unique nearest-depot policy

$$\mathbf{x}_{\text{P0}} = (d_1, d_1, d_1, d_2, d_2, d_3, d_3, d_3), \quad \bar{R}(\mathbf{x}_{\text{P0}}) = 5.25 \text{ min}, \quad \max_z r_z(\mathbf{x}_{\text{P0}}) = 11 \text{ min}.$$

Group means and pairwise disparities under $\mathbf{x}_{\text{P0}}$ are reported in Table 3. The audit findings are immediate: minimum service is violated at z_7 and z_8 (response times 9 and 11 min exceed $\theta = 8$); demographic parity is violated by a small margin (max disparity $2.348 > \delta = 2$); and the soft objectives evaluate to SO-EQ$=$ 0.252, SO-RW$=$ 11, and SO-VW$=$ -80.91. The baseline solution is therefore inadmissible under two hard requirements. Whether either failure can be resolved within $\mathcal{X}$ is the question to which Stages 1–3 of the diagnostic procedure now respond.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 21 May 2026 doi:10.20944/preprints202605.1432.v1

Table 3. Group-level outcomes under the P0 (ethically-agnostic) baseline.

Quantity	Value (minutes)
$\bar{y}_{\text{Maj}}(\mathbf{x}_{\text{P0}})$	4.281
$\bar{y}_{\text{Black}}(\mathbf{x}_{\text{P0}})$	6.629
$\bar{y}_{\text{Hisp}}(\mathbf{x}_{\text{P0}})$	5.915
$\lvert\bar{y}_{\text{Maj}} - \bar{y}_{\text{Black}}\rvert$	2.348
$\lvert\bar{y}_{\text{Maj}} - \bar{y}_{\text{Hisp}}\rvert$	1.634
$\lvert\bar{y}_{\text{Black}} - \bar{y}_{\text{Hisp}}\rvert$	0.715
maximum pairwise disparity	2.348

5.3. Ethical Requirements and Feasibility Diagnosis

The audited rule set was $\mathcal{R} = \mathcal{H} \cup \dot{\mathcal{S}} \cup \mathcal{D}$ with $\mathcal{H} = \{(\text{HC-MS})@\theta = 8, (\text{HC-DP})@\delta = 2\}$, $\dot{\mathcal{S}} = \{\text{SO-EQ}, \text{SO-VW}, \text{SO-RW}\}$, and $\mathcal{D} = \{D_1, D_2\}$. The dynamic rule D_1 (surge-triggered) relaxed the service threshold to 12 minutes when $s \in \{\text{SURGE}, \text{DISASTER}\}$; the dynamic rule D_2 (disaster-triggered) imposed a 10-minute response standard on high-vulnerability zones ($v_z \geq 0.5$, here z_6, z_7, z_8) when $s = \text{DISASTER}$.

Stage 1 (Type I).

Pairwise hard-rule feasibility tests returned no internally infeasible pairs: $(HC\text{-}MS)$ at $\theta = 8$ and $(HC\text{-}DP)$ at $\delta = 2$ each admitted a satisfying assignment in X in isolation. The rule set is therefore internally consistent.

Stage 2 (Type II).

For (HC-MS) to hold, every zone must be reachable from at least one depot within the 8-minute threshold. The per-zone feasibility check in Table 4 shows that for zones z_7 and z_8 no depot satisfies the threshold. Because (HC-MS) is separable across zones, these local infeasibility certificates directly imply global infeasibility: no assignment in $\mathcal{X}$ satisfies (HC-MS), without requiring enumeration. Direct enumeration over $\mathcal{X}$ confirms the result, returning zero of the 6,561 assignments as (HC-MS)-feasible. The IIS certificate is

$$\text{IIS}_{\text{HC-MS}} = \{\, (\text{HC-MS}) \text{ at } \theta = 8;\; T[d_1, z_7] = 20,\; T[d_2, z_7] = 14,\; T[d_3, z_7] = 9; \\ T[d_1, z_8] = 22,\; T[d_2, z_8] = 16,\; T[d_3, z_8] = 11 \,\}, \quad (9)$$

in which all binding conditions are travel-time entries from the infrastructure topology, not algorithmic restrictions on $\mathcal{X}$ or fairness rules. This means the conflict cannot be resolved by any algorithmic adjustment.

Table 4. Per-zone feasibility of (HC-MS) at $\theta = 8$ minutes.

Zone	$\min_d T[d, z]$ (min)	Feasible depot(s)
z_1	3	d_1, d_2
z_2	4	d_1, d_2
z_3	5	d_1, d_2
z_4	3	d_1, d_2
z_5	4	d_2, d_3
z_6	7	d_3
z_7	9	—
z_8	11	—

In contrast, (HC-DP) at $\delta = 2$ remains feasible within $\mathcal{X}$: the lowest-$\bar{R}$ assignment satisfying (HC-DP) is $(d_1, d_2, d_2, d_2, d_2, d_3, d_3, d_3)$ with $\bar{R} = 5.79$ min and max pairwise disparity 1.78 min. The

 doi:10.20944/preprints202605.1432.v1

joint feasibility set (HC-MS)∩(HC-DP) is empty because (HC-MS) alone is infeasible; Stage 2 therefore returns a single binding Type II conflict on (HC-MS).

Stage 3 (Type IV).

Under $s = \text{SURGE}$, activation of D_1 relaxes the service threshold to 12 minutes, restoring feasibility: 144 of the 6,561 assignments satisfy the relaxed standard, and the minimum max-disparity is 0.594 min. Under $s = \text{DISASTER}$, both D_1 and D_2 activate. While z_6 and z_7 satisfy the 10-minute vulnerability-prioritized threshold from d_3, zone z_8 remains infeasible since $\min_d T[d, z_8] = 11 > 10$. The system is therefore infeasible under DISASTER, yielding a Type IV exposure.

5.4. Application of Theorem 1

The disparity coefficients $c_{g,g',z}$ for the metropolitan instance are reported in Table 5. Applying Theorem 1 yields the signed-disparity bounds

$$\bar{y}_{\text{Maj}} - \bar{y}_{\text{Black}} \in [-7.817, +2.215],$$
$$\bar{y}_{\text{Maj}} - \bar{y}_{\text{Hisp}} \in [-5.669, +2.157],$$
$$\bar{y}_{\text{Black}} - \bar{y}_{\text{Hisp}} \in [-0.346, +2.435]$$

(all in minutes). All three intervals contain zero, so Corollary 1's one-sided condition fails on every group pair, and the corollary delivers no absolute-disparity infeasibility on this instance. Enumeration confirms the implication: the minimum simultaneous max-pairwise disparity attainable across the three pairs is 0.020 min, achieved at $\bar{R} = 9.18$ min and $\max_z r_z = 16$ min. The metropolitan instance therefore exhibits a substantively different behavior from the stylized two-group/two-depot configuration: at this scale, with three groups and three depots, demographic parity at $\delta = 2$ min is feasible, and the binding Type II conflict is on (HC-MS) rather than on (HC-DP). The diagnostic apparatus correctly distinguishes the two regimes.

Table 5. Disparity coefficients $c_{g,g',z}$ from Eq. (5).

z	$c_{\text{Maj,Black},z}$	$c_{\text{Maj,Hisp},z}$	$c_{\text{Black,Hisp},z}$
z_1	+0.165	+0.133	−0.032
z_2	+0.141	+0.128	−0.013
z_3	+0.095	+0.089	−0.006
z_4	+0.047	+0.048	+0.001
z_5	+0.049	−0.057	−0.105
z_6	−0.217	−0.150	+0.067
z_7	−0.191	−0.123	+0.068
z_8	−0.089	−0.068	+0.021

5.5. Pareto Frontier and Harm Redistribution

Because (HC-MS) at $\theta = 8$ is infeasible under NORMAL, the trade-off analysis is presented under a documented relaxation of (HC-MS) to $\theta = 12$ min, corresponding to the dynamic rule D_1 authorized by the oversight board as a temporary operating condition pending capital investment. Under (HC-MS) at $\theta = 12$, 144 assignments are feasible. The Pareto frontier in $(\bar{R}, \max\ \text{disparity})$-space over the full decision space $\mathcal{X}$, computed by enumeration on this small instance, is reported in Table 6.

Table 6. Pareto frontier in $(\bar{R}, \max\ \text{disparity})$-space over the full decision space $\mathcal{X}$.

#	$\bar{R}$ (min)	Max disp. (min)	$\max_z r_z$	Assignment
0	5.250	2.348	11	$(d_1, d_1, d_1, d_2, d_2, d_3, d_3, d_3)$
1	5.464	2.066	11	$(d_1, d_2, d_1, d_2, d_2, d_3, d_3, d_3)$
2	5.571	2.062	11	$(d_1, d_1, d_2, d_2, d_2, d_3, d_3, d_3)$
3	5.786	1.780	11	$(d_1, d_2, d_2, d_2, d_2, d_3, d_3, d_3)$
4	5.821	1.688	11	$(d_2, d_1, d_1, d_2, d_2, d_3, d_3, d_3)$
5	6.036	1.406	11	$(d_2, d_2, d_1, d_2, d_2, d_3, d_3, d_3)$
6	6.143	1.403	11	$(d_2, d_1, d_2, d_2, d_2, d_3, d_3, d_3)$
7	6.214	1.078	13	$(d_1, d_3, d_1, d_2, d_2, d_3, d_3, d_3)$
8	6.536	0.793	13	$(d_1, d_3, d_2, d_2, d_2, d_3, d_3, d_3)$
9	6.786	0.466	13	$(d_2, d_3, d_1, d_2, d_2, d_3, d_3, d_3)$
10	7.036	0.332	14	$(d_3, d_2, d_1, d_2, d_2, d_3, d_3, d_3)$
11	7.357	0.223	13	$(d_2, d_3, d_1, d_2, d_3, d_3, d_3, d_3)$
12	7.607	0.146	14	$(d_3, d_2, d_1, d_2, d_3, d_3, d_3, d_3)$
13	7.679	0.089	13	$(d_2, d_3, d_2, d_2, d_3, d_3, d_3, d_3)$
14	8.250	0.061	13	$(d_2, d_3, d_1, d_1, d_3, d_3, d_3, d_3)$
15	8.714	0.026	15	$(d_2, d_2, d_3, d_1, d_1, d_3, d_3, d_3)$
16	9.179	0.020	16	$(d_3, d_2, d_2, d_1, d_3, d_3, d_3, d_2)$

The lowest-$\bar{R}$ assignment satisfying $\delta = 2$ is assignment #3, $(d_1, d_2, d_2, d_2, d_2, d_3, d_3, d_3)$, with $\bar{R} = 5.79$ min and max disparity 1.78 min, costing +0.54 min (+10.2%) relative to the P0 baseline. The frontier exhibits two regions: assignments #0–#6 maintain $\max_z r_z = 11$ and trade efficiency for parity by adjusting only north-and-central zone assignments, while assignments #7–#16 progressively reassign zones to d_3 in less-direct combinations, incurring larger worst-case zone response in pursuit of further parity reduction. The 0.020-min minimum disparity at point #16 is the smallest attainable in $\mathcal{X}$, consistent with the signed-disparity intervals containing zero (Section 5.4); the minimum absolute disparity reflects the discrete granularity of $\mathcal{X}$, not a structural floor.

The population-impact decomposition of the move from P0 to the incumbent (Table 7) reveals the central substantive feature of equity attained within fixed infrastructure. Equity improves: maximum disparity falls from 2.35 to 1.78 min. *All three groups, however, experience worse absolute response times under the equity-compliant policy than under the operationally-optimal P0 baseline*; no group's service improves. The equity gain is achieved by slowing the Majority group by 18.5%, while the Black and Hispanic groups are also slowed (by 3.4% and 4.5% respectively). The mechanism is explicit in Theorem 1: the achievable group-mean response for the Black group is bounded below by a quantity determined by depot geography, and reducing inter-group disparity within $\mathcal{X}$ requires raising the faster-served group rather than lowering the slower-served group. This finding is precisely the kind of distributional structure that algorithm-only fairness auditing systematically misses, because algorithm-only auditing examines the dispatch policy in isolation from the infrastructure within which the policy operates.

Table 7. Population-level impact of moving from the P0 baseline to the Phase 3 incumbent (assignment #3 of Table 6).

Group	$\bar{y}_g$ at P0 (min)	$\bar{y}_g$ at #3 (min)	Δ (min)	Δ (%)
Majority	4.281	5.074	+0.793	+18.5%
Black	6.629	6.854	+0.225	+3.4%
Hispanic	5.915	6.183	+0.269	+4.5%

5.6. Capital-Investment Specification

The Type II IIS certificate of Eq. (9) contains no rule-set or algorithmic constraints; the binding conditions are travel-time entries to z_7 and z_8. ECP converts this finding into a quantified investment specification. A proposed fourth depot d_4 must satisfy $T[d_4, z_7] \le 8$ and $T[d_4, z_8] \le 8$ (restoring

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 21 May 2026 doi:10.20944/preprints202605.1432.v1

(HC-MS); z_6 is already feasible from d_3 at 7 min), which also implies $T[d_4, z_8] \leq 10$ (restoring D_2 feasibility under DISASTER). The strict-feasibility threshold is

$$\max\big(T[d_4, z_7],\ T[d_4, z_8]\big)\ \leq\ 8\ \text{min}. \tag{10}$$

To verify the recommendation, a hypothetical d_4 with travel-time vector $T[d_4, \cdot] = (25, 24, 26, 14, 13, 5, 4, 6)$—a depot sited near the southern minority cluster and far from the northern zones—was evaluated. The decision space expands to $4^8 = 65{,}536$. Two natural Pareto corners under the joint constraint set (HC-MS) at $\theta = 8$ + (HC-DP) at $\delta = 2 + D_2$ are reported in Table 8: the $\bar{R}$-optimal corner and the parity-optimal corner.

Table 8. Pre- and post-investment feasibility for the case-study instance.

	Pre (3 depots)	**Post: $\bar{R}$-optimal (4 depots)**	**Post: parity-optimal (4 depots)**
Joint feasibility (HC-MS@8 + HC-DP@2 + D_2)	infeasible	feasible	feasible
Assignments satisfying joint constraints	0 of 6,561	64 of 65,536	
$\bar{R}$ at the corner	—	4.07 min	5.16 min
max disparity at the corner	—	0.51 min	0.11 min
$\bar{y}_{\text{Maj}}$	—	3.86	5.16
$\bar{y}_{\text{Black}}$	—	4.37	5.16
$\bar{y}_{\text{Hisp}}$	—	4.23	5.05
$\max_z r_z$	—	6 min	7 min

The capital investment simultaneously resolves (HC-MS), (HC-DP) at $\delta = 2$, the D_2 disaster-state requirement, and the SO-RW worst-case bound at every joint-feasible assignment. At the $\bar{R}$-optimal corner, mean response time improves to 4.07 min relative to the 5.25-min P0 baseline, all three group means fall below 4.4 min, and the worst-case zone response time drops from 11 to 6 min. The parity-optimal corner trades some of this efficiency gain for a tighter equity profile (max disparity 0.11 min vs. 0.51 min) at $\bar{R} = 5.16$ min, still below the pre-investment baseline. Both corners satisfy the board's $\delta = 2$ target with substantial margin.

The implication is the principal substantive finding of the case study. The apparent efficiency–equity trade-off documented along the pre-investment Pareto frontier of Section 5.5 is itself an artifact of constrained infrastructure rather than a fundamental property of the allocation problem. Once the binding infrastructural constraint is removed, both objectives improve simultaneously. An auditor evaluating only the dispatch algorithm under the three-depot configuration would observe a real but infrastructure-induced trade-off and might conclude that equity gains require operational sacrifice; the IIS-based diagnosis instead identifies the trade-off as removable and produces the specification of how to remove it.

6. Discussion

The findings advance the core argument of the paper: ethical infeasibility in AI-enabled infrastructure systems is not a single phenomenon with a single cause. The dominant algorithmic-fairness literature characterizes one mechanism, namely the internal inconsistency of fairness criteria within a classifier, and proves that several natural fairness objectives cannot be simultaneously satisfied. The Structural Infeasibility Theorem identifies a categorically different mechanism. The bounds $L_{g,g'}$ and $U_{g,g'}$ in Eqs. (6)–(7) are functions only of the travel-time matrix, the spatial distribution of demand, and the group-composition profile, all of which are properties of the physical and demographic environment rather than of any algorithmic design choice. When the interval $[L_{g,g'}, U_{g,g'}]$ does not contain zero and the implied parity floor exceeds the regulator's tolerance, no allocation policy in $\mathcal{X}$ can satisfy

 doi:10.20944/preprints202605.1432.v1

demographic parity, and the conflict cannot be resolved by algorithmic adjustment. The source of the failure lies in the infrastructure, and it is invisible to any audit that examines only the algorithm.

The case study makes this distinction operationally concrete in two complementary ways. First, the IIS certificate for the HC-MS infeasibility contains exclusively travel-time entries from the depot-zone topology. No algorithmic constraint, no fairness rule, and no rule-set inconsistency appears in the certificate. Algorithm-only auditing of the dispatch system under the three-depot configuration would correctly identify that the system fails to meet the NFPA 1710 standard, but would have no mechanism to identify the cause as infrastructural and no basis for producing a constructive remedy. The IIS procedure provides both: a formal attribution of cause and a quantified specification of the physical investment required to restore compliance.

Second, the harm-redistribution finding documented in Table 7 demonstrates that equity attained within fixed infrastructure has a fundamentally different welfare character from equity attained through capital investment. Under the Pareto incumbent, all three groups experience worse absolute outcomes relative to the operationally optimal baseline; under the post-investment corner of Table 8, all three groups experience better outcomes. Identical disparity values can therefore correspond to fundamentally different distributional structures, and equity audits that report only disparity statistics cannot distinguish between them. The Structural Infeasibility Theorem clarifies the underlying mechanism: when the achievable signed-disparity interval lies entirely on one side of zero, disparity reduction within fixed infrastructure requires slowing the faster-served group rather than improving service to the slower-served group, so that equity is achieved by leveling down rather than by lifting up.

The findings have three implications for the ethical governance of AI-enabled infrastructure systems. First, audit scope must extend beyond the allocation algorithm to the physical system on which it operates. The IIS certificate is the formal object that determines whether observed inequity is attributable to the algorithm, to the infrastructure, or to a rule conflict, and the appropriate intervention is determined by that attribution. Second, equity statistics evaluated on a fixed infrastructure can mask harm-redistribution patterns that are ethically consequential. The population-impact decomposition of Section 5.5 represents a minimum standard for the distributional analysis of equity-driven policy changes: it is not sufficient to report that disparity fell; it is necessary to report which groups were slowed and which were speeded in order to characterize the welfare structure of the change. Third, when the IIS certificate attributes infeasibility to physical infrastructure, the framework does not return a refused deployment but a quantified capital-investment specification. The locus of decision shifts from the algorithmic-design layer to the capital-planning layer of the governing institution, and the framework provides the technical basis for that shift.

The case study is deliberately deterministic to permit exact verification of the framework's analytical properties. Stochastic extensions including uncertain travel times, stochastic call arrivals, and time-varying demand preserve the qualitative content of Theorem 1 in the sense that the signed-disparity bounds hold in expectation under standard conditions, but would refine the specific numerical thresholds derived above. The framework assumes that ethical rule sets are specified externally by governance institutions acting in good faith; it operationalizes ethical commitments but does not determine which commitments are normatively appropriate. Computational scaling beyond the analytically tractable case-study size, together with the algorithmic specification, the five-phase governance workflow, and cross-domain application templates within which the diagnostic procedure of this paper is operationally embedded, are developed in the companion paper currently being developed by the research team.

7. Conclusions

This paper has identified and formalized a third source of ethical conflict in AI-enabled infrastructure systems, distinct from the internal-inconsistency mechanism of the classifier-fairness literature and from algorithmic design choices: structural infeasibility induced by the physical configuration

 doi:10.20944/preprints202605.1432.v1

of the infrastructure system itself. The Structural Infeasibility Theorem derives closed-form bounds on inter-group disparity attainable across all feasible allocation policies, with the bounds determined by infrastructure topology rather than by algorithm design. The hierarchical IIS-based diagnostic procedure attributes infeasibility to one of three sources—rule, algorithm, or infrastructure—and produces categorically different interventions for each. The metropolitan ambulance-dispatch case study demonstrates that the structural mechanism is empirically substantive: minimum-service infeasibility is certified as infrastructural and resolvable only through capital investment; demographic parity is feasible despite operational asymmetry only because the multi-group, multi-depot regime escapes the one-sided-interval condition; the apparent efficiency–equity trade-off is an artifact of constrained infrastructure that disappears when the binding infrastructural constraint is removed; and equity attained within fixed infrastructure is harm redistribution, not harm reduction.

The principal methodological implication is that ethical auditing of AI-enabled infrastructure systems cannot be confined to the allocation algorithm. The infrastructure on which the algorithm operates must itself be subjected to formal ethical-feasibility analysis, because algorithm-only audits are structurally incapable of identifying infrastructure-induced infeasibility. The constructive counterpart of this critique is that when infeasibility is attributable to infrastructure, the IIS certificate is straightforwardly convertible into a quantified capital-planning specification; the framework converts an algorithmic-fairness audit finding into actionable infrastructure investment guidance.

Author Contributions: Conceptualization, S.C.; methodology, S.C.; formal analysis, S.C.; investigation, S.C., A.A.; Data Curation, S.C., A.A.; writing—original draft preparation, S.C., A.A.; M.A.Q.; writing—review and editing, S.C., A.A., M.A.Q. The authors have read and agreed to the published version of the manuscript.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on reasonable request.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ECP	Ethical Constraint Programming
EMS	Emergency Medical Services
HC-DP	Hard Constraint – Demographic Parity
HC-MS	Hard Constraint – Minimum Service
IIS	Irreducible Infeasible Subsystem
NFPA	National Fire Protection Association
SO-EQ	Soft Objective – Equity (Gini-type)
SO-RW	Soft Objective – Rawlsian Welfare
SO-VW	Soft Objective – Vulnerability-Weighted Welfare

References

1. Bajwa, A. AI-based emergency response systems: A systematic literature review on smart infrastructure safety. *SSRN Electronic Journal* **2025**, Article 5171521.
2. Rehan, H. Enhancing disaster response systems: Predicting and mitigating the impact of natural disasters using AI. *J. Artif. Intell. Res.* **2022**, *2*, 501.
3. Olteanu, A.-D.; Gheorghe, Ș.; Chicco, G. Artificial Intelligence-Driven Strategies for Management in Modern Utility Infrastructures: The Role of AI on Power Quality Management. In *Proceedings of the 2025 25th International Conference on Control Systems and Computer Science (CSCS)*; IEEE: Piscataway, NJ, USA, 2025; pp. 494–501.
4. Bircan, T.; Özbilgin, M. F. Unmasking inequalities of the code: Disentangling the nexus of AI and inequality. *Technol. Forecast. Soc. Change* **2025**, *211*, 123925.

5. Yu, P. K. The algorithmic divide and equality in the age of artificial intelligence. *Florida Law Rev.* **2020**, *72*, 331.
6. Chowdhury, S.; Shahvari, O.; Marufuzzaman, M.; Li, X.; Bian, L. Drone routing and optimization for post-disaster inspection. *Comput. Ind. Eng.* **2021**, *159*, 107495.
7. Chowdhury, S.; Zhu, J.; Zhang, W. Optimized restoration planning of infrastructure system-of-systems using heterogeneous network flow simulation. *J. Comput. Civ. Eng.* **2020**, *34*, 04020032.
8. Amaldi, E.; Pfetsch, M. E.; Trotter, L. E., Jr. On the maximum feasible subsystem problem, IISs and IIS-hypergraphs. *Math. Program.* **2003**, *95*, 533–554.
9. Journigan, W. R., Sr. *Virginia Beach Fire Department and NFPA 1710: Where Are We?*; Executive Fire Officer Program Applied Research Project; U.S. Fire Administration, National Fire Academy: Emmitsburg, MD, USA, 2008.
10. Arab, A.; Khodaei, A.; Khator, S. K.; Ding, K.; Emesih, V. A.; Han, Z. Stochastic pre-hurricane restoration planning for electric power systems infrastructure. *IEEE Trans. Smart Grid* **2015**, *6*, 1046–1054.
11. Berber, T.; Mostafavi, A. Equity in infrastructure restoration after disasters: A systems analysis of disparities. *Nat. Commun.* **2023**, *14*, 4617.
12. Chinneck, J. W. *Feasibility and Infeasibility in Optimization: Algorithms and Computational Methods*; Springer: New York, NY, USA, 2007.
13. Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* **2017**, *5*, 153–163.
14. Cybersecurity and Infrastructure Security Agency. *Artificial Intelligence in Critical Infrastructure: Principles and Framework*; Technical Report; U.S. Department of Homeland Security: Washington, DC, USA, 2023.
15. U.S. Department of Homeland Security. *Artificial Intelligence Roadmap*; Technical Report; U.S. Department of Homeland Security: Washington, DC, USA, 2024.
16. Dong, S.; Mostafavi, A. Measuring emergent social infrastructure for community resilience to disasters. *Nat. Hazards* **2022**, *111*, 1479–1506.
17. Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; Zemel, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*; ACM: New York, NY, USA, 2012; pp. 214–226.
18. European Commission. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Off. J. Eur. Union* **2024**.
19. Elliott, J. R.; Howell, J. Beyond disasters: A longitudinal analysis of natural hazards' unequal impacts on residential instability. *Soc. Forces* **2017**, *95*, 1181–1207.
20. Erkut, E. Inequality measures for location problems. *Locat. Sci.* **1993**, *1*, 199–217.
21. Erkut, E.; Ingolfsson, A.; Erdogan, G. Ambulance location for maximum survival. *Nav. Res. Logist.* **2008**, *55*, 42–58.
22. Fang, Y.-P.; Pedroni, N.; Zio, E. Resilience-based component importance measures for critical infrastructure network systems. *IEEE Trans. Reliab.* **2016**, *65*, 502–512.
23. Fussell, E. Leaving New Orleans: Social stratification, networks, and hurricane evacuation. In *Understanding Katrina: Perspectives from the Social Sciences*; Social Science Research Council: Brooklyn, NY, USA, 2006.
24. Gleeson, J.; Ryan, J. Identifying minimally infeasible subsystems of inequalities. *ORSA J. Comput.* **1990**, *2*, 61–63.
25. Guieu, O.; Chinneck, J. W. Analyzing infeasible mixed-integer and integer linear programs. *INFORMS J. Comput.* **1999**, *11*, 63–77.
26. Gutjahr, W. J.; Nolz, P. C. Multicriteria optimization in humanitarian aid. *Eur. J. Oper. Res.* **2016**, *252*, 351–366.
27. Hardt, M.; Price, E.; Srebro, N. Equality of opportunity in supervised learning. *Adv. Neural Inf. Process. Syst.* **2016**, *29*, 3315–3323.
28. High-Level Expert Group on AI. *Ethics Guidelines for Trustworthy AI*; Technical Report; European Commission: Brussels, Belgium, 2019.
29. Hsia, R. Y.; Shen, Y.-C.; Kanzaria, H. K.; Sarkar, U. Disparities in the management of acute myocardial infarction among racial and ethnic groups. *Med. Care* **2011**, *49*, 578–587.
30. Kleinberg, J.; Mullainathan, S.; Raghavan, M. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*; 2016; arXiv:1609.05807.
31. Marsh, M. T.; Schilling, D. A. Equity measurement in facility location analysis: A review and framework. *Eur. J. Oper. Res.* **1994**, *74*, 1–17.
32. McLay, L. A.; Mayorga, M. E. A dispatching model for server-to-customer systems that balances efficiency and equity. *Manuf. Serv. Oper. Manag.* **2013**, *15*, 205–220.

 doi:10.20944/preprints202605.1432.v1

33. Mostafavi, A.; Ganapati, N. E.; Nazarnia, H.; Pradhananga, N.; Khanal, R. Adaptive capacity under a dynamic resilience framework: Modeling the recovery of infrastructure service provision during flood events. *Nat. Hazards Rev.* **2021**, *22*, 04020051.
34. Nejat, A.; Ghosh, S. LASSO model of postdisaster housing recovery: Case study of Hurricane Sandy. *Nat. Hazards Rev.* **2016**, *17*, 04016007.
35. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*; Technical Report; U.S. Department of Commerce: Washington, DC, USA, 2023.
36. Organisation for Economic Co-operation and Development. *Recommendation of the Council on Artificial Intelligence*; Technical Report OECD/LEGAL/0449; OECD: Paris, France, 2019.
37. Ogryczak, W.; Śliwiński, T. On solving linear programs with the ordered weighted averaging objective. *Eur. J. Oper. Res.* **2003**, *148*, 80–91.
38. Pell, J. P.; Sirel, J. M.; Marsden, A. K.; Ford, I.; Cobbe, S. M. Effect of reducing ambulance response times on deaths from out of hospital cardiac arrest: Cohort study. *BMJ* **2001**, *322*, 1385–1388.
39. Santos, J. R.; van der Linden, J. Understanding household displacement patterns following Hurricane Sandy: A spatial analysis approach. *Nat. Hazards Rev.* **2016**, *17*, 04016008.
40. Sasser, S. M.; Hunt, R. C.; Faul, M.; Sugerman, D.; Pearson, W. S.; Dulski, T.; Wald, M. M.; Jurkovich, G. J.; Newgard, C. D.; Lerner, E. B. Guidelines for field triage of injured patients: Recommendations of the national expert panel on field triage, 2011. *Morb. Mortal. Wkly. Rep. Recomm. Rep.* **2013**, *61*, 1–20.
41. Mitchell, S.; Potash, E.; Barocas, S.; D'Amour, A.; Lum, K. Algorithmic fairness: Choices, assumptions, and definitions. *Annu. Rev. Stat. Appl.* **2021**, *8*, 141–163.
42. Green, B. Escaping the impossibility of fairness: From formal to substantive algorithmic fairness. *Philos. Technol.* **2022**, *35*, 90.
43. Miller, M.; Chowdhury, S.; Alzarrad, A.; Hossain, N. U. I. Synergetic Decision-Making: Analyzing the Interplay of Human Behavior and Physical Infrastructure in Emergency Evacuations via an Analytical Approach. *Glob. J. Flex. Syst. Manag.* **2024**, *25*, 785–803.