engineering proceedings

MDPI

Proceeding Paper

Neutrosophic Uncertainty Set Construction for Adaptive Robust Pre-Disaster Evacuation Planning Under Indeterminacy †

Yasmany Fernández-Fernández [1,2,*], Sira M. Allende-Alonso [2], Ridelio Miranda-Pérez [3] and Gemayqzel Bouza-Allende [2]

[1] Facultad de Industrias Agropecuarias y Ciencias Ambientales, Universidad Politécnica Estatal del Carchi (UPEC), Tulcán 040101, Ecuador
[2] Facultad de Matemática y Computación, Universidad de La Habana, La Habana 10400, Cuba; gema@matcom.uh.cu (G.B.-A.)
[3] Facultad de Ciencias Económicas y Empresariales, Universidad de Cienfuegos, Cienfuegos 55100, Cuba; rmiranda@ucf.edu.cu
* Correspondence: yfernandezf@upec.edu.ec
† Presented at the Sixth International Conference on Information Systems and Software Technologies (ICI2ST 2026), Quito, Ecuador, 23–25 September 2026.

Abstract

Pre-disaster evacuation planning under hurricane threats faces two intertwined uncertainty sources: stochastic demand variability across storm trajectories and epistemic temporal indeterminacy in travel and service durations. Classical robust optimization addresses the latter through worst-case box sets, yet the construction of these sets—their size, shape, and coupling to the scenario-driven demand model—remains largely ad hoc. We propose a neutrosophic-informed framework for the systematic construction of uncertainty sets in adaptive robust vehicle scheduling. Each route parameter is endowed with a truth degree (confidence in nominal estimates), an indeterminacy degree (plausible but unresolvable deviations) and a falsity degree (physically excluded extremes), and these are propagated into a scenario-anchored budget set whose geometry is calibrated via neutrosophic indeterminacy measures. The adaptive robust counterpart is derived in closed form by linear-programming duality and embedded in a four-stage hierarchical metamodel for NP-hard fleet scheduling. The framework is validated on a georeferenced Caribbean network for Hurricane Matthew (Santiago de Cuba, 2016; 19 nodes, 57 arcs, 20 vehicles, up to 8196 persons). Neutrosophic calibration yields $\Gamma^{n} = 13.53$ over 145 route segments, and the affine decision rules deliver makespans of 470–545 time units under fixed first-stage assignments, with the full pipeline running in under four minutes.

Keywords: neutrosophic uncertainty sets; adaptive robust optimization; hurricane evacuation; affine decision rules; budget uncertainty; hierarchical metamodel; stochastic–robust anchoring; NP-hard scheduling

Check for updates

Academic Editor: Erick Herrera

Published: 18 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Hurricane evacuation planning requires committing to resource allocation decisions hours or days before landfall, under two fundamentally distinct uncertainties: probabilistic demand variability across plausible storm trajectories and epistemic temporal perturbations in travel and service durations. Stochastic programming [1] addresses the first by optimising over a discrete scenario ensemble; robust optimisation [2] handles the second by immunising decisions against worst-case realisations within prescribed uncertainty sets. However, the construction of these uncertainty sets—their shape, size and relationship to

the stochastic scenarios already in the model—remains largely ad hoc in the humanitarian logistics literature, despite its direct impact on solution quality and conservatism [3].

The classical approach, introduced by Soyster [4] and refined by Ben-Tal and Nemirovski [5] and Bertsimas and Sim [6], defines uncertainty sets as boxes or budgeted polyhedra around nominal parameter values. In evacuation planning, however, the nominal values themselves depend on which hurricane scenario materialises: a Category 4 direct landfall produces different road conditions than a weakened trajectory. Ignoring this dependence forces the planner either to inflate the uncertainty set conservatively—protecting against temporal deviations on top of the worst-case demand scenario—or to underestimate it by using nominal parameters that may not apply under extreme conditions. Recent advances in disaster logistics have begun to address this gap. Sun et al. [7] proposed scenario-based robust models for emergency facility location under disruption uncertainty. Siddig and Song [8] incorporated evacuation-related uncertainties through event-based and box representations. Bayram and Yaman [9] combined staging with dynamic resource allocation under demand uncertainty. Eriskin and Karatas [10] applied robust optimisation to shelter location-allocation with budgeted uncertainty. In the broader robust-optimisation literature, Bertsimas et al. [11] introduced adaptive distributionally robust optimisation, while Lorca and Sun [12] developed multistage adaptive robust models with affine decision rules for power systems—a setting with structural parallels to fleet scheduling under temporal uncertainty.

A distinctive feature of evacuation parameter uncertainty is the coexistence of three epistemological states for any given travel time τ_r on route r: (i) a degree of truth T associated with the nominal estimate, derived from geographic data and average conditions; (ii) a degree of indeterminacy I reflecting plausible deviations that cannot be resolved before the operation begins (congestion effects, weather-induced degradation, evacuee behavioural variability); and (iii) a degree of falsity F corresponding to extreme values excluded by physical constraints. This tripartite structure maps onto the neutrosophic-set framework introduced by Smarandache [13], where every element is characterised by truth-, indeterminacy-, and falsity-memberships, each independently valued in $]^{-}0, 1^{+}[$. The key departure from intuitionistic fuzzy sets is the independence of I: the planner's inability to resolve temporal indeterminacy is not the complement of what is known and what is excluded, but carries its own informational content that should shape the uncertainty-set geometry.

Contributions. The paper makes four contributions: (i) we construct the temporal uncertainty set $\mathcal{U}$ by deriving route-level deviation bounds $\sigma(\lambda)$ directly from the stochastic program's scenario ensemble Ξ, establishing a formal link between demand-driven scenarios and temporal protection (Section 3); (ii) we introduce a principled methodology for calibrating the budget parameter Γ and the set geometry using neutrosophic indeterminacy measures, replacing the ad-hoc $\Gamma = \sqrt{|\Lambda|}$ heuristic by a data-driven Γ^{n} grounded in the (T, I, F) characterisation of each parameter (Section 4); (iii) we formulate the robust vehicle-scheduling problem with affine recourse $t(\xi) = t^{0} + \sum_{r} t^{r}\xi_{r}$, derive its tractable robust counterpart via linear-programming duality, and prove that the adaptive formulation dominates the static counterpart for every realisation within $\mathcal{U}$ (Section 5); (iv) we validate the framework on a georeferenced Hurricane Matthew instance, showing that neutrosophic-calibrated scenario-anchored sets ($\Gamma^{n} = 13.53$ on 145 route segments) yield adaptive schedules with scenario-dependent makespans between 470 and 545 ut, while revealing the precise conditions under which adaptive flexibility justifies its cost premium over static robustness (Section 6). Section 2 reviews the neutrosophic–robust interface; Section 7 concludes.

https://doi.org/10.3390/engproc2026153041

2. Preliminaries and the Neutrosophic–Robust Interface

For ease of reference, Table 1 collects the principal symbols, grouped by the layer of the framework in which they appear.

Table 1. Principal notation, grouped by framework layer.

Symbol	Meaning
Network and fleet	
$\mathcal{E}, \Lambda, \mathcal{V}, K$	arcs; routes/segments; vehicles; trip slots
$\bar{d}_{ij}$, v_{avg}, τ_{serv}	nominal arc distance; average speed; per-person service time
P_λ, Cap_v, $T_{\max}$	route demand; vehicle capacity; planning horizon
Scenarios and anchoring	
Ξ, ω, p_ω	scenario set; scenario; probability
$\phi^\omega_{h,i}$, δ^ω	scenario demand; demand ratio $\sum \phi^\omega / \sum \phi^0$
κ, ζ^ω	demand–distance coupling; distance (travel-condition) factor
$\bar{\tau}_\lambda$, τ^ω_λ, $\sigma(\lambda)$	nominal duration; scenario duration; anchored deviation
Neutrosophic characterisation	
$\langle T, I, F\rangle$, $S(\cdot)$	truth/indeterminacy/falsity degrees; score function
$\bar{I}$, $\text{cv}^2(I)$	aggregate indeterminacy; squared coefficient of variation
$\Lambda^* = \{\lambda : F_\lambda = 0\}$	falsity-filtered (physically feasible) routes
Robust counterpart	
$\mathcal{U}$, $\mathcal{U}^\Xi_\Gamma$, Γ, Γ^n	uncertainty set; anchored set; budget; neutrosophic budget
$\xi_\lambda \in [0,1]$	standardised load $(\tau_\lambda - \bar{\tau}_\lambda)/\sigma(\lambda)$
$t(\boldsymbol{\xi}) = t^0 + \sum_r t^r \xi_r$	affine decision rule (nominal + sensitivity terms)
$W, F, U, \eta, \mu_\lambda$	assignment; families; activation; makespan; demand slack
$\pi, \mu^\pm$	dual variables of the robust counterpart

2.1. Neutrosophic Sets: Definitions and Notation

Definition 1 (Neutrosophic Set [13]). *Let X be a universe of discourse. A neutrosophic set A in X is characterised by three membership functions $T_A, I_A, F_A : X \to \,]^-0, 1^+[$, called the truth-, indeterminacy- and falsity-membership functions. There is no restriction on the sum $T_A(x) + I_A(x) + F_A(x)$; it may range anywhere in $]^-0, 3^+[$.*

For computational purposes we work with single-valued neutrosophic numbers (SVNNs), where $T, I, F \in [0,1]$.

Definition 2 (Single-Valued Neutrosophic Number). *An SVNN is a triplet $\langle T, I, F\rangle$ with $T, I, F \in [0,1]$ and $0 \le T + I + F \le 3$.*

Definition 3 (Score Function [14]). *The score function of an SVNN $\alpha = \langle T, I, F\rangle$ is*

$$S(\alpha) \;=\; \frac{T + (1-I) + (1-F)}{3} \;=\; \frac{2 + T - I - F}{3} \;\in\; [0,1]. \tag{1}$$

2.2. Classical Robust Optimisation Framework

Consider a generic optimisation problem with uncertain parameter $\boldsymbol{\xi} \in \mathbb{R}^n$:

$$\min_{\boldsymbol{x}} \{ f(\boldsymbol{x}) \;:\; g(\boldsymbol{x}, \boldsymbol{\xi}) \le 0,\ \forall \boldsymbol{\xi} \in \mathcal{U} \}, \tag{2}$$

where $\mathcal{U} \subseteq \mathbb{R}^n$ is the uncertainty set. The three classical constructions are the box set [4] $\mathcal{U}_{\text{box}} = \{\boldsymbol{\xi} : |\xi_r - \bar{\xi}_r| \le \hat{\xi}_r,\ \forall r\}$, the ellipsoidal set [5] $\mathcal{U}_{\text{ell}} = \{\boldsymbol{\xi} : \|\boldsymbol{\xi} - \bar{\boldsymbol{\xi}}\|_2 \le \Omega\}$ and the budget set [6] $\mathcal{U}_\Gamma = \{\boldsymbol{\xi} : |\xi_r - \bar{\xi}_r| \le \hat{\xi}_r,\ \sum_r |\xi_r - \bar{\xi}_r|/\hat{\xi}_r \le \Gamma\}$. The budget parameter $\Gamma \in [0, n]$

controls conservatism: $\Gamma = 0$ recovers the nominal problem and $\Gamma = n$ recovers the box. Bertsimas and Sim [6] showed that the violation probability under the budget set satisfies

$$\Pr\left[\sum_r \frac{|\xi_r - \bar{\xi}_r|}{\hat{\xi}_r} > \Gamma\right] \leq \exp\left(-\frac{\Gamma^2}{2n}\right), \tag{3}$$

providing a principled, though conservative, relationship between Γ and the probability guarantee.

2.3. The Neutrosophic–Robust Interface

The connection between neutrosophic theory and robust optimisation is direct: the uncertainty set $\mathcal{U}$ encodes the indeterminacy degree I of the parameter space.

Definition 4 (Neutrosophic Characterisation of an Uncertain Parameter)**.** *Let τ_r be an uncertain duration parameter for route r with nominal value $\bar{\tau}_r$, scenario realisations $\{\tau_r^\omega\}_{\omega\in\Xi}$ with probabilities $\{p_\omega\}$, maximum deviation $\hat{\tau}_r = \max_\omega |\tau_r^\omega - \bar{\tau}_r|$, and planning horizon $T_{\max}$. Its neutrosophic characterisation is the SVNN*

$$\mathfrak{n}(\tau_r) = \langle T_r, I_r, F_r \rangle, \tag{4}$$

where the three degrees are read off three operationally distinct statistics of the scenario ensemble:

$$\underbrace{T_r = 1 - \frac{\sum_\omega p_\omega |\tau_r^\omega - \bar{\tau}_r|}{\hat{\tau}_r}}_{\text{probability concentration}}, \qquad \underbrace{I_r = \frac{\hat{\tau}_r}{\bar{\tau}_r}}_{\text{worst-case spread}}, \qquad \underbrace{F_r = \frac{\max(0,\, \bar{\tau}_r + \hat{\tau}_r - T_{\max})}{T_{\max}}}_{\text{horizon exceedance}}. \tag{5}$$

The truth degree T_r is one minus the ratio of the mean absolute deviation (a measure of where probability mass sits) to the maximum deviation (the support radius): $T_r = 1$ when all mass concentrates at the nominal value, and $T_r \to 0$ when mass piles up at the extremes. The indeterminacy I_r measures the normalised worst-case spread, and the falsity F_r the normalised amount by which the worst case overshoots the horizon.

The reviewer rightly asks in what precise sense this construction is neutrosophic rather than a re-parameterised box. The earlier parameterisation $T_r = 1 - \hat{\tau}_r/\bar{\tau}_r$ forced $T_r + I_r = 1$ and was, indeed, reducible. Definition 4 removes that artefact: T_r is now a concentration statistic (a probability-weighted mean), I_r a range statistic (a supremum), and F_r a location statistic (proximity to the horizon). These three functionals of the ensemble are algebraically free, as the following proposition shows.

Proposition 1 (Genuine independence of the three degrees)**.** *For every admissible triple (t, i, f) with $t, i, f \in [0, 1]$ and $i > 0$ there exists a route ensemble $(\bar{\tau}, \{\tau^\omega\}_\omega, \{p_\omega\}_\omega, T_{\max})$ whose neutrosophic characterisation (5) equals exactly $\langle t, i, f \rangle$. In particular, no algebraic identity (such as $T + I = 1$ or $T + I + F = 1$) constrains the degrees; the relation $T_r + I_r = 1$ holds only in the degenerate case in which all probability mass lies at the two extremes.*

Proof. Fix (t, i, f). Choose three free design knobs and match each degree in turn. *(i) Spread.* Set $\bar{\tau} = 1$ and $\hat{\tau} = i$, so $I = \hat{\tau}/\bar{\tau} = i$. *(ii) Concentration.* Use three scenarios at deviations $\{-i, 0, +i\}$ with symmetric probabilities $\{q/2,\, 1 - q,\, q/2\}$, $q \in [0, 1]$. Then the mean absolute deviation is $q\, i$ and $T = 1 - (q\, i)/i = 1 - q$, which sweeps the whole interval $[0, 1]$ as q ranges over $[0, 1]$; set $q = 1 - t$ to obtain $T = t$. This choice does not affect I (still i, the support radius). *(iii) Horizon.* The degree $F = \max(0, \bar{\tau} + \hat{\tau} - T_{\max})/T_{\max} = \max(0, 1 + i - T_{\max})/T_{\max}$ is a strictly decreasing bijection of $T_{\max}$ on $(0, 1 + i]$ onto $[0, 1)$; solving $1 + i - T_{\max} = f\, T_{\max}$ gives $T_{\max} = (1 + i)/(1 + f)$, which fixes $F = f$ without altering T or I (neither depends on $T_{\max}$). The three steps act on disjoint parameters (support

radius, mass allocation, horizon), so the map (ensemble) $\mapsto (T, I, F)$ is surjective onto the open admissible region; hence the degrees are independent. Finally, $T + I = 1$ requires $q = t$ only when $t = 1 - q$ and i cancels, i.e., when $q = 1$ (all mass at the extremes), the degenerate case. □

Remark 1 (Operational irreducibility to a box)**.** *Independence is not merely formal: each degree drives a distinct decision. The indeterminacy I_r sizes the uncertainty set (Section 3); the falsity F_r filters routes out of the budget, defining the admissible set $\Lambda^* = \{r : F_r = 0\}$ (Remark 5); and the truth T_r reports reliability and modulates the calibration surcharge through the indeterminacy dispersion (Section 4). A plain interval $[\bar{\tau}_r, \bar{\tau}_r + \hat{\tau}_r]$ encodes only I_r: it cannot distinguish a wide-but-concentrated profile (high I, high T) from a wide-and-diffuse one (high I, low T), and it carries no horizon-exclusion information (F). It is precisely this three-way separation—confidence in the nominal, irreducible spread, physical exclusion—that an intuitionistic fuzzy model would collapse by forcing $I = 1 - T - F$, conflating "we do not know" with "we know it is impossible".*

Proposition 2 (Uncertainty Set as Indeterminacy Envelope)**.** *The budget uncertainty set $\mathcal{U}_\Gamma$ for temporal parameters $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_{|\Lambda|})$ can be written as*

$$\mathcal{U}_\Gamma = \left\{ \boldsymbol{\tau} \in \mathbb{R}_+^{|\Lambda|} : \tau_r \in [\bar{\tau}_r,\ \bar{\tau}_r(1 + I_r)],\ \sum_r \frac{\tau_r - \bar{\tau}_r}{\bar{\tau}_r \cdot I_r} \le \Gamma,\ \bar{\tau}_r + \hat{\tau}_r \le T_{\max} \right\}, \tag{6}$$

where I_r is the indeterminacy degree from Definition 4 and the last condition enforces $F_r = 0$ (physical feasibility).

Proof. Substituting $\hat{\tau}_r = \bar{\tau}_r \cdot I_r$ into the Bertsimas–Sim budget set yields $[\bar{\tau}_r, \bar{\tau}_r(1 + I_r)]$; the budget constraint becomes $\sum_r (\tau_r - \bar{\tau}_r)/(\bar{\tau}_r I_r) \le \Gamma$. The condition $F_r = 0$ amounts to $\bar{\tau}_r + \hat{\tau}_r \le T_{\max}$. □

3. Scenario-Anchored Uncertainty Set Construction

3.1. The Anchoring Principle

The standard approach to constructing temporal uncertainty sets for robust vehicle scheduling proceeds independently of the stochastic demand model: the planner specifies a relative deviation ρ (e.g., 25%) and applies it uniformly to all route durations. This has two deficiencies. First, it ignores the correlation between demand scenarios and travel conditions—a pessimistic hurricane trajectory producing high demand also degrades the road conditions, increasing travel times. Second, it provides no mechanism for calibrating ρ from data. We address both through scenario anchoring: the temporal uncertainty set is constructed by propagating the stochastic program's scenario ensemble $\Xi = \{\omega_1, \ldots, \omega_S\}$ into the temporal parameter space.

3.2. Formal Construction

Let $\mathrm{P}_\mathrm{I}^\mathrm{S}$ be the two-stage stochastic program with scenario set Ξ, where each $\omega \in \Xi$ specifies demands $\phi_{h,i}^\omega$ and has probability $p_\omega > 0$ with $\sum_\omega p_\omega = 1$.

Definition 5 (Scenario Distance Factor)**.** *For each scenario $\omega \in \Xi$, the distance factor ζ^ω models the correlation between demand intensity and travel-condition degradation:*

$$\zeta^\omega = 1 + \left(\frac{\sum_{h,i} \phi_{h,i}^\omega}{\sum_{h,i} \phi_{h,i}^0} - 1 \right) \cdot \kappa, \tag{7}$$

where $\phi_{h,i}^0$ is the nominal demand and $\kappa \in [0, 1]$ is the demand–distance coupling parameter. When $\kappa = 0$, distances are scenario-independent; when $\kappa = 1$, the relative distance deviation equals the relative demand deviation.

https://doi.org/10.3390/engproc2026153041

Definition 6 (Scenario-Specific Arc Distance and Route Duration). *For each arc* $(i,j) \in \mathcal{E}$ *and scenario* $\omega \in \Xi$, $d_{ij}^{\omega} = \bar{d}_{ij} \cdot \zeta^{\omega}$. *For each route* $\lambda \in \Lambda$,

$$\tau_{\lambda}^{\omega} = \frac{d_{\lambda}^{\omega}}{v_{\text{avg}}} + \min(P_{\lambda}, \text{Cap}_{\max}) \cdot \tau_{\text{serv}}, \tag{8}$$

where $d_{\lambda}^{\omega} = \sum_{(i,j)\in\lambda} d_{ij}^{\omega}$, v_{avg} *is the average speed,* $P_{\lambda} = \sum_{r} f_r h_r$ *is the total route demand,* $\text{Cap}_{\max} = \max_{v} \text{Cap}_{v}$ *and* τ_{serv} *is the per-person service time.*

Theorem 1 (Anchored Deviation Bound). *Define the anchored deviation for route* λ *as*

$$\sigma(\lambda) = \max_{\omega\in\Xi} \left| \tau_{\lambda}^{\omega} - \bar{\tau}_{\lambda} \right|, \tag{9}$$

with $\bar{\tau}_{\lambda} = \tau_{\lambda}^{0}$ *the nominal duration. Then*

$$\sigma(\lambda) = \bar{\tau}_{\lambda}^{\text{tp}} \cdot \max_{\omega\in\Xi} |\zeta^{\omega} - 1|, \tag{10}$$

where $\bar{\tau}_{\lambda}^{\text{tp}} = \bar{d}_{\lambda}/v_{\text{avg}}$ *is the nominal transport component; the service component cancels because it is scenario-independent.*

Proof. From (8), $\tau_{\lambda}^{\omega} - \bar{\tau}_{\lambda} = \bar{\tau}_{\lambda}^{\text{tp}}(\zeta^{\omega} - 1)$. Taking absolute values and maximising over ω yields (10). □

Corollary 1 (Effective Relative Deviation). $\sigma(\lambda)/\bar{\tau}_{\lambda} \le \max_{\omega} |\zeta^{\omega} - 1|$, *with equality when the service component vanishes. For routes with significant service time, the effective relative deviation is smaller than the demand-induced distance variation, reflecting the stabilising effect of the deterministic service component.*

Remark 2 (The anchoring chain, step by step). *The construction maps demand scenarios to a temporal indeterminacy degree through four explicit links, each isolating one modelling decision:*

$$\underbrace{\delta^{\omega} = \frac{\sum_{h,i} \phi_{h,i}^{\omega}}{\sum_{h,i} \phi_{h,i}^{0}}}_{(1)\ \textit{demand ratio}} \xrightarrow{\kappa} \underbrace{\zeta^{\omega} = 1 + (\delta^{\omega} - 1)\kappa}_{(2)\ \textit{distance factor}} \xrightarrow{d^{\omega} = \bar{d}\zeta^{\omega}} \underbrace{\sigma(\lambda) = \bar{\tau}_{\lambda}^{\text{tp}} \max_{\omega} |\zeta^{\omega} - 1|}_{(3)\ \textit{deviation}} \xrightarrow{/\,\bar{\tau}_{\lambda}} \underbrace{I_{\lambda} = \sigma(\lambda)/\bar{\tau}_{\lambda}}_{(4)\ \textit{indeterminacy}}. \tag{11}$$

The coupling κ *is the only free modelling parameter: it is the elasticity of travel-condition degradation to demand intensity. At* $\kappa = 0$ *the chain is severed* ($\zeta^{\omega} \equiv 1$, $\sigma \equiv 0$, *the nominal problem); at* $\kappa = 1$ *a relative demand surge passes undamped into a relative distance inflation; intermediate values encode the empirical fact that storm-driven demand and road degradation are correlated but not identical. The quantity* ζ^{ω} *is, thus, a dimensionless travel-condition multiplier applied uniformly to every arc length under scenario* ω.

Example 1 (Worked anchoring for the Hurricane Matthew instance). *The three scenarios carry total demands* $P^{\omega_1} = 8196$ *(pessimistic),* $P^{0} = 6830$ *(nominal) and* $P^{\omega_3} = 5464$ *(optimistic), so the demand ratios are* $\delta^{\omega_1} = 8196/6830 = 1.20$, $\delta^{0} = 1.00$ *and* $\delta^{\omega_3} = 0.80$. *With the operational coupling* $\kappa = 0.75$, *the distance factors are* $\zeta^{\omega_1} = 1 + (1.20 - 1) \cdot 0.75 = 1.15$, $\zeta^{0} = 1.00$ *and* $\zeta^{\omega_3} = 1 - 0.15 = 0.85$; *hence* $\max_{\omega} |\zeta^{\omega} - 1| = 0.15$. *For route* λ_1 *(Cayo Granma, nominal transport time* $\bar{\tau}_{\lambda_1}^{\text{tp}} = 52.4$ *min), the anchored deviation is* $\sigma(\lambda_1) = 52.4 \times 0.15 = 7.86$ *min and the indeterminacy degree is* $I_{\lambda_1} = 7.86/52.4 = 0.15$. *Because the ensemble is symmetric with probabilities* $(0.25, 0.50, 0.25)$, *the mean absolute deviation is* $0.25(7.86) + 0.50(0) + 0.25(7.86) = 3.93$ *min, giving the concentration truth* $T_{\lambda_1} = 1 - 3.93/7.86 = 0.50$ *and, with* $\bar{\tau}_{\lambda_1} + \sigma(\lambda_1) = 60.3 < T_{\max} = 480$, *falsity* $F_{\lambda_1} = 0$. *The triplet* $\langle 0.50, 0.15, 0.00 \rangle$ *reproduces the first row of Table 3.*

https://doi.org/10.3390/engproc2026153041

Definition 7 (Scenario-Anchored Budget Uncertainty Set)**.**

$$\mathcal{U}_{\Gamma}^{\Xi} = \Big\{ \boldsymbol{\tau} \in \mathbb{R}_{+}^{|\Lambda|} : \tau_\lambda \in [\bar{\tau}_\lambda,\, \bar{\tau}_\lambda + \sigma(\lambda)],\, \sum_{\lambda} \frac{\tau_\lambda - \bar{\tau}_\lambda}{\sigma(\lambda)} \le \Gamma \Big\}. \tag{12}$$

Theorem 2 (Scenario Containment)**.** *$\mathcal{U}_{\Gamma}^{\Xi}$ contains every scenario realisation when $\Gamma \ge |\Lambda|$, i.e., $\{\boldsymbol{\tau}^\omega : \omega \in \Xi\} \subseteq \mathcal{U}_{|\Lambda|}^{\Xi}$. Moreover, for the tighter bound $\Gamma^* = \max_\omega \sum_\lambda |\tau_\lambda^\omega - \bar{\tau}_\lambda|/\sigma(\lambda)$, containment holds with $\Gamma \ge \Gamma^*$.*

Proof. By construction, $|\tau_\lambda^\omega - \bar{\tau}_\lambda| \le \sigma(\lambda)$ for all ω, λ (Definition 6 and Equation (9)), so each term in the budget sum is at most 1; with $|\Lambda|$ such terms, $\Gamma = |\Lambda|$ suffices. The tighter Γ^* follows by direct evaluation. □

Remark 3 (Two guarantees: deterministic coverage versus probabilistic budget)**.** *Theorem 2 is a deterministic statement about the finite ensemble: with $\Gamma = \Gamma^*$ every scenario point $\boldsymbol{\tau}^\omega$ lies inside $\mathcal{U}_{\Gamma}^{\Xi}$, so a schedule robust over the set is feasible for each realised scenario with certainty. This is distinct from the probabilistic budget calibration of Section 4, which controls violation over the continuous box $[\,\bar{\tau}_\lambda, \bar{\tau}_\lambda + \sigma(\lambda)\,]$ under a stochastic perturbation model. The operational reading of Γ^* is a coverage–conservatism dial: $\Gamma = |\Lambda|$ is the full box (maximal conservatism, every route simultaneously at its worst), while $\Gamma^* = \max_\omega \sum_\lambda |\tau_\lambda^\omega - \bar{\tau}_\lambda|/\sigma(\lambda) \le |\Lambda|$ is the smallest budget that still covers the ensemble exactly, removing the slack between "protect against all scenarios" and "protect against the box". When indeterminacy is homogeneous and the worst scenario saturates every route ($|\tau_\lambda^{\omega^*} - \bar{\tau}_\lambda| = \sigma(\lambda)\ \forall \lambda$), one has $\Gamma^* = |\Lambda|$ and the two coincide; heterogeneous ensembles produce $\Gamma^* < |\Lambda|$, which is exactly where the gap—and the value of anchoring—appears.*

The construction is monotone in coupling: for $\kappa_1 \le \kappa_2$, $\mathcal{U}_{\Gamma}^{\Xi}(\kappa_1) \subseteq \mathcal{U}_{\Gamma}^{\Xi}(\kappa_2)$, so stronger demand–distance coupling produces larger uncertainty sets. Algorithm 1 formalises the complete construction. Its complexity is $O(|\Xi| \cdot |\Lambda| \cdot L)$, with $L = \max_\lambda |\lambda|$ the maximum route length; for typical instances ($|\Xi| \le 50$, $|\Lambda| \le 200$, $L \le 10$), it is negligible compared with the MIP solution time.

Algorithm 1 Scenario-Anchored Uncertainty Set Construction.

Require: Scenario set $\Xi = \{\omega_1, \ldots, \omega_S\}$ with probabilities p_ω and demands $\phi_{h,i}^\omega$; nominal arc distances $\bar{d}_{ij}$; route set Λ; coupling $\kappa \in [0,1]$; parameters $v_{\text{avg}}, \tau_{\text{serv}}, T_{\max}$.
Ensure: Anchored set $\{\sigma(\lambda), \bar{\tau}_\lambda\}$ and neutrosophic budget Γ^{n}.
for each $\omega \in \Xi$ **do**
  $\delta^\omega \leftarrow \sum_{h,i} \phi_{h,i}^\omega / \sum_{h,i} \phi_{h,i}^0$; $\zeta^\omega \leftarrow 1 + (\delta^\omega - 1)\kappa$.
end for
for each $\lambda \in \Lambda$ **do**
  $\bar{d}_\lambda \leftarrow \sum_{(i,j)\in\lambda} \bar{d}_{ij}$; $\bar{\tau}_\lambda \leftarrow \bar{d}_\lambda / v_{\text{avg}} + \min(P_\lambda, \text{Cap}_{\max}) \cdot \tau_{\text{serv}}$.
  for each $\omega \in \Xi$ **do**
    $\tau_\lambda^\omega \leftarrow \bar{d}_\lambda \zeta^\omega / v_{\text{avg}} + \min(P_\lambda, \text{Cap}_{\max}) \cdot \tau_{\text{serv}}$.
  end for
  $\sigma(\lambda) \leftarrow \max_\omega |\tau_\lambda^\omega - \bar{\tau}_\lambda|$; $m_\lambda \leftarrow \sum_\omega p_\omega |\tau_\lambda^\omega - \bar{\tau}_\lambda|$.
  $I_\lambda \leftarrow \sigma(\lambda)/\bar{\tau}_\lambda$; $T_\lambda \leftarrow 1 - m_\lambda/\sigma(\lambda)$; $F_\lambda \leftarrow \max(0, \bar{\tau}_\lambda + \sigma(\lambda) - T_{\max})/T_{\max}$.
end for
$\bar{I} \leftarrow \frac{1}{|\Lambda|} \sum_\lambda I_\lambda$; $\text{Var}(I) \leftarrow \frac{1}{|\Lambda|} \sum_\lambda (I_\lambda - \bar{I})^2$.
$\Gamma^{\text{n}} \leftarrow \sqrt{|\Lambda|} \cdot \big(1 + \beta\, \text{Var}(I)/\bar{I}^2\big)$.
return $\{\sigma(\lambda), \bar{\tau}_\lambda, \langle T_\lambda, I_\lambda, F_\lambda \rangle\}_\lambda$, Γ^{n}.

4. Neutrosophic Calibration of the Budget Parameter

The classical heuristic $\Gamma = \sqrt{|\Lambda|}$ (from the probabilistic analysis in [6]) assumes independent, identically distributed perturbations—an assumption violated in evacuation planning, where correlated weather effects perturb multiple routes simultaneously. We propose instead a calibration grounded in the neutrosophic characterisation of each route's temporal uncertainty.

Definition 8 (Route Indeterminacy Profile and Aggregate)**.** *For each route $\lambda \in \Lambda$, the neutrosophic triplet $\mathfrak{n}(\lambda) = \langle T_\lambda, I_\lambda, F_\lambda \rangle$ is obtained from the scenario ensemble through (5), with $\hat{\tau}_\lambda = \sigma(\lambda)$: the concentration truth $T_\lambda = 1 - (\sum_\omega p_\omega |\tau_\lambda^\omega - \bar{\tau}_\lambda|)/\sigma(\lambda)$, the spread indeterminacy $I_\lambda = \sigma(\lambda)/\bar{\tau}_\lambda$, and the horizon falsity $F_\lambda = \max(0, \bar{\tau}_\lambda + \sigma(\lambda) - T_{\max})/T_{\max}$. The aggregate indeterminacy of the route network is $\bar{I} = \frac{1}{|\Lambda|}\sum_\lambda I_\lambda$ and its dispersion is measured by the squared coefficient of variation $\mathrm{cv}^2(I) = \mathrm{Var}(I_\lambda)/\bar{I}^2$.*

We build the calibration in three transparent steps: a rigorous independent baseline, a rigorous correlated extension, and an observable surrogate for the (latent) correlation. Throughout, $\xi_\lambda = (\tau_\lambda - \bar{\tau}_\lambda)/\sigma(\lambda) \in [0, 1]$ denotes the standardised load of route λ, so the budget constraint reads $\sum_\lambda \xi_\lambda \le \Gamma$.

Lemma 1 (Independent baseline)**.** *Assume* (A1) *the standardised loads $\{\xi_\lambda\}_{\lambda\in\Lambda}$ are independent and* (A2) *centred and bounded, $\mathbb{E}[\xi_\lambda - \bar{I}] = 0$, $\xi_\lambda - \bar{I} \in [-1, 1]$. Then the probability that the budget is exceeded obeys the Bertsimas–Sim bound*

$$\Pr\Big[\textstyle\sum_\lambda \xi_\lambda > \Gamma\Big] \;\le\; \exp\Big(-\frac{(\Gamma - |\Lambda|\bar{I})^2}{2|\Lambda|}\Big). \tag{13}$$

In particular, taking $\Gamma_0 = \sqrt{|\Lambda|}$ above the mean load makes the right-hand side equal to the fixed level $e^{-1/2} \approx 0.607$, which is the precise probabilistic meaning of the classical heuristic $\Gamma = \sqrt{|\Lambda|}$.

Proof. Hoeffding's inequality applied to the bounded independent centred summands $\xi_\lambda - \bar{I}$ gives (13); this is the form used by Bertsimas and Sim [6] to derive (3). Setting $\Gamma - |\Lambda|\bar{I} = \sqrt{|\Lambda|}$ yields exponent $-1/2$. □

Proposition 3 (Heterogeneity surcharge under correlation)**.** *Replace* (A1) *by* (A1′) *the loads are equicorrelated with $\mathrm{Corr}(\xi_\lambda, \xi_{\lambda'}) = \rho \ge 0$ for $\lambda \ne \lambda'$ and common variance v. Then $\mathrm{Var}(\sum_\lambda \xi_\lambda) = |\Lambda| v\,[1 + (|\Lambda| - 1)\rho]$, and the Gaussian (CLT) budget that controls the exceedance probability at a target level α is*

$$\Gamma_\rho = |\Lambda|\bar{I} + z_\alpha \sqrt{|\Lambda| v}\,\sqrt{1 + (|\Lambda| - 1)\rho}\,, \tag{14}$$

which inflates the independent budget by the factor $\sqrt{1 + (|\Lambda| - 1)\rho}$.

Proof. The variance of an equicorrelated sum is standard; the CLT/Lyapunov approximation of the budget at level α is $\mathbb{E}[\sum \xi_\lambda] + z_\alpha\,\mathrm{sd}(\sum \xi_\lambda)$. □

The correlation ρ is not directly observable, but in evacuation networks it is produced by the same mechanism that makes indeterminacy heterogeneous: routes exposed to the common weather front simultaneously exhibit larger and co-moving deviations. We, therefore, use the empirical dispersion of the indeterminacy profile as an observable surrogate.

Definition 9 (Neutrosophic budget parameter)**.** *Let $\mathrm{cv}^2(I) = \mathrm{Var}(I_\lambda)/\bar{I}^2$ be the squared coefficient of variation of the indeterminacy profile. The neutrosophic budget is the surrogate of* (14) *obtained by replacing the latent correlation surcharge $\sqrt{1 + (|\Lambda| - 1)\rho}$ by its first-order proxy $1 + \beta\,\mathrm{cv}^2(I)$:*

$$\Gamma^{\mathfrak{n}} \;=\; \sqrt{|\Lambda|}\,\big(1 + \beta\,\mathrm{cv}^2(I)\big), \qquad \beta \ge 0. \tag{15}$$

When the network is homogeneous ($\mathrm{Var}(I_\lambda) = 0$) the surrogate is exact in the sense that it returns the independent baseline $\sqrt{|\Lambda|}$ of Lemma 1; heterogeneity inflates the budget monotonically in β and in $\mathrm{cv}^2(I)$.

https://doi.org/10.3390/engproc2026153041

Remark 4 (When is the surrogate valid?) *The identification* $\sqrt{1+(|\Lambda|-1)\rho} \approx 1+\beta\,\mathrm{cv}^2(I)$ *is a first-order (Taylor) match valid in the regime of moderate dispersion and moderate correlation:* $\mathrm{cv}^2(I)$ *small enough that the linearisation* $\sqrt{1+x} \approx 1+\frac{1}{2}x$ *holds, and a roughly affine empirical link* $\rho \approx c\,\mathrm{cv}^2(I)/(|\Lambda|-1)$ *with* $\beta = \frac{1}{2}c$. *Outside this regime—very large* $\mathrm{cv}^2(I)$ *or near-perfect correlation (*$\rho \to 1$*, where the surcharge approaches* $\sqrt{|\Lambda|}$ *and the budget the full box* $\Gamma = |\Lambda|$*)—one should use (14) directly with an estimated* ρ. *The CLT step further requires* $|\Lambda|$ *large and the Lyapunov ratio small, both satisfied in our instances (*$|\Lambda| \geq 46$*). We adopt* $\beta = 1$ *throughout, which corresponds to* $z_\alpha\sqrt{v} \approx 1$ *in (14) and a unit linear sensitivity of the surcharge to dispersion; this choice is conservative (it never under-protects relative to the independent baseline) and is the value used in all experiments.*

Remark 5 (Falsity-Filtered Routes)**.** *Routes with* $F_\lambda > 0$ *(worst-case duration exceeds the planning horizon) require special treatment. We define* $\Lambda^* = \{\lambda \in \Lambda : F_\lambda = 0\}$ *and compute* $\Gamma^{\mathfrak{n}}$ *over* Λ^**, since routes with nonzero falsity are handled through the demand-relaxation mechanism (*μ_r *slacks in the scheduling model).*

Definition 10 (Neutrosophic Uncertainty Severity)**.** *Route* λ_1 *has greater uncertainty severity than* λ_2*, denoted* $\lambda_1 \succ_{\mathfrak{n}} \lambda_2$*, if* $I_{\lambda_1} > I_{\lambda_2}$*, or* $I_{\lambda_1} = I_{\lambda_2}$ *and* $F_{\lambda_1} > F_{\lambda_2}$.

From (1), $S(\lambda) = (2 + T_\lambda - I_\lambda - F_\lambda)/3$ is strictly decreasing in both indeterminacy and falsity ($\partial S/\partial I = \partial S/\partial F = -\frac{1}{3}$), so among routes of equal truth degree the severity order $\succ_{\mathfrak{n}}$ is consistent with the score ordering: $T_{\lambda_1} = T_{\lambda_2}$ and $\lambda_1 \succ_{\mathfrak{n}} \lambda_2$ imply $S(\lambda_1) < S(\lambda_2)$. When truth degrees differ the two effects may oppose, which is intended: a wide but highly concentrated profile need not be more severe than a narrow but diffuse one.

Proposition 4 (Neutrosophic Feasibility Certificate)**.** *Let* (W^*, F^*, U^*, t^*) *be optimal for the adaptive robust counterpart with uncertainty set* $\mathcal{U}^{\Xi}_{\Gamma^{\mathfrak{n}}}$*, and let* $p_{\mathrm{viol}} = \exp(-(\Gamma^{\mathfrak{n}})^2/2|\Lambda|)$ *be the Bertsimas–Sim bound (3) on the probability that the budget is exceeded under the independent-perturbation model of Lemma 1. Then the horizon constraint for each active route admits the reporting triplet*

$$\mathfrak{n}\Big(t^{\mathrm{end}}_{v,k} \leq T_{\max}\Big) = \langle\, 1 - p_{\mathrm{viol}},\ \ p_{\mathrm{viol}} - F_{\mathrm{phys}},\ \ F_{\mathrm{phys}}\,\rangle, \qquad F_{\mathrm{phys}} = \tfrac{1}{|\Lambda^*|}\textstyle\sum_\lambda F_\lambda, \tag{16}$$

where $1 - p_{\mathrm{viol}}$ *is the certified lower bound on robust feasibility,* F_{phys} *the mean physical-exclusion mass, and the middle term the residual indeterminacy.*

Proof. By Lemma 1, $\Pr[\sum_\lambda \xi_\lambda > \Gamma^{\mathfrak{n}}] \leq p_{\mathrm{viol}}$, so feasibility holds with probability at least $1 - p_{\mathrm{viol}}$; F_{phys} is the data-driven horizon-exceedance mass from (5), and the indeterminacy component is the residual $1 - (1 - p_{\mathrm{viol}}) - F_{\mathrm{phys}}$. □

The triplet (16) is a communication device, not a sharpening of the probability bound: its truth slot reports the rigorous Bertsimas–Sim feasibility floor $1 - p_{\mathrm{viol}}$, its falsity slot the empirical horizon-exceedance mass, its indeterminacy slot the residual. We claim only that bundling these three already-defined quantities into one $\langle T, I, F\rangle$ summary is operationally convenient for crisis communication—an empirical modelling advantage in our case study—not that the tripartite form yields a tighter or universally superior guarantee.

The neutrosophic-informed approach differs from classical fuzzy and intuitionistic-fuzzy alternatives in the independence of I: in evacuation planning the planner's inability to resolve travel-time uncertainty is independent of both the confidence in the nominal estimate and the physical impossibility of extreme values. A fuzzy model captures only a vague "around 50 min" through a single membership $\mu(x) \in [0, 1]$; an intuitionistic fuzzy

model (μ, ν) with $\mu+\nu \leq 1$ forces $I = 1-T-F$, conflating "we do not know" with "we know it is impossible"—two states with fundamentally different operational implications. The neutrosophic triplet (T, I, F) keeps the three independent, and our construction anchors them to data: I from the scenario spread, F from physical horizon limits.

5. Adaptive Robust Counterpart with Affine Decision Rules

5.1. Static Versus Adaptive Robust Formulations

In the static robust approach (Soyster counterpart) all temporal variables are fixed before uncertainty is revealed: the schedule must be feasible for the worst-case realisation in $\mathcal{U}$. In the adaptive (adjustable) approach, temporal variables depend on the realised perturbation $\boldsymbol{\xi}$. Following [2,12], we restrict attention to affine decision rules (ADRs).

Definition 11 (Affine Decision Rule)**.** *A temporal variable t follows an ADR if $t(\boldsymbol{\xi}) = t^0 + \sum_{r\in\Lambda} t^r \xi_r$, where $t^0 \in \mathbb{R}$ is the nominal component and $t^r \in \mathbb{R}$ are the sensitivity coefficients.*

5.2. The Adaptive Robust Scheduling Model $\mathrm{P}_{\mathrm{II}}^{\mathrm{ARO}}$

Let Λ be the set of routes extracted from the stochastic planner $\mathrm{P}_{\mathrm{I}}^{\mathrm{S}}$, $\mathcal{V}$ the vehicle fleet, K the set of trip slots, and $\mathcal{U}_\Gamma^\Xi$ the scenario-anchored uncertainty set. The decision variables are: $W_{v,\lambda,k} \in \{0,1\}$ (vehicle v serves route λ in slot k, first-stage), $F_{v,\lambda,k} \in \mathbb{Z}_+$ (families transported, first-stage), $U_v \in \{0,1\}$ (vehicle activation), and the adaptive trip start/end times $t_{v,k}^{\mathrm{start}}(\boldsymbol{\xi}), t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}) \in \mathbb{R}_+$ and makespan $\eta(\boldsymbol{\xi}) \in \mathbb{R}_+$, all following ADRs as in Definition 11.

The formulation is

$$\min \quad \sum_v c_v^x U_v + \sum_{v,\lambda,k} c^{\mathrm{km}} d_\lambda W_{v,\lambda,k} + \varphi\eta^0 + M_\mu \sum_\lambda h_\lambda \mu_\lambda \tag{17a}$$

$$\text{s.t.} \quad \sum_{v,k} F_{v,\lambda,k} + \mu_\lambda \geq f_\lambda, \quad F_{v,\lambda,k} \leq f_\lambda W_{v,\lambda,k}, \quad \sum_\lambda h_\lambda F_{v,\lambda,k} \leq \mathrm{Cap}_v \tag{17b}$$

$$t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}) \geq t_{v,k}^{\mathrm{start}}(\boldsymbol{\xi}) + \sum_\lambda (\bar{\tau}_\lambda + \sigma(\lambda)\xi_\lambda) W_{v,\lambda,k}, \quad \forall \boldsymbol{\xi} \in \mathcal{U}_\Gamma^\Xi, \tag{17c}$$

$$t_{v,k+1}^{\mathrm{start}}(\boldsymbol{\xi}) \geq t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}), \quad t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}) \leq T_{\max}, \quad \eta(\boldsymbol{\xi}) \geq t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}), \quad \forall \boldsymbol{\xi} \in \mathcal{U}_\Gamma^\Xi, \tag{17d}$$

together with vehicle-origin compatibility, activation and symmetry-breaking constraints as in the static model.

5.3. Tractable Robust Counterpart via Linear-Programming Duality

The semi-infinite constraints in (17c) and (17d) are reformulated finitely via LP duality.

Theorem 3 (Robust Counterpart of the Horizon Constraint)**.** *The constraint $t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}) \leq T_{\max}$ holds for all $\boldsymbol{\xi} \in \mathcal{U}_\Gamma^\Xi$ if and only if there exist $\pi_{v,k} \geq 0$ and $\mu_{v,k,\lambda}^{\pm} \geq 0$ such that*

$$t_{v,k}^{0,\mathrm{end}} + \Gamma\pi_{v,k} + \sum_\lambda (\mu_{v,k,\lambda}^{+} + \mu_{v,k,\lambda}^{-}) \leq T_{\max}, \tag{18a}$$

$$\pi_{v,k} + \mu_{v,k,\lambda}^{+} \geq \sigma(\lambda) t_{v,k}^{\lambda,\mathrm{end}}, \quad \pi_{v,k} + \mu_{v,k,\lambda}^{-} \geq -\sigma(\lambda) t_{v,k}^{\lambda,\mathrm{end}}, \quad \forall\lambda. \tag{18b}$$

Proof. The worst case of $t_{v,k}^{\mathrm{end}}(\boldsymbol{\xi}) \leq T_{\max}$ over $\mathcal{U}_\Gamma^\Xi$ is $\max\{t_{v,k}^{0,\mathrm{end}} + \sum_\lambda t_{v,k}^{\lambda,\mathrm{end}} \sigma(\lambda)\xi_\lambda\} \leq T_{\max}$ with $\xi_\lambda \in [-1,1]$ and $\sum_\lambda |\xi_\lambda| \leq \Gamma$. Its dual (introducing π for the budget and $\mu^{\pm}$ for the bounds) gives (18a) and (18b) by strong LP duality. □

Setting $t_{v,k}^{\lambda,\mathrm{end}} = 0$ for all λ (no adaptivity) and $\Gamma = |\Lambda|$ recovers the Soyster static counterpart, in which all durations take worst-case values simultaneously.

https://doi.org/10.3390/engproc2026153041

Theorem 4 (Dominance of Adaptive over Static)**.** *Let Z_*^{ARO} and Z_*^{SRO} denote the optimal values of the adaptive and static robust counterparts on the same route granularity. Then $Z_*^{\mathrm{ARO}} \leq Z_*^{\mathrm{SRO}}$, with strict inequality generically.*

Proof. Every static solution (W, F, U, t^0) is feasible for the ARO with all sensitivity coefficients $t^r = 0$. The ARO optimises over a strictly larger feasible set, hence its optimal value is at most that of the SRO. □

The complete derivation of the robust counterpart of (17c), which is the most technically involved, adds $1 + 2|\Lambda|$ continuous dual variables per (v, k) pair but yields a finite-dimensional mixed-integer program; the analogous treatment of (17d) is straightforward.

6. Computational Experiments

6.1. Instance: Hurricane Matthew, Santiago de Cuba

We validate the framework on a georeferenced evacuation network from Santiago de Cuba corresponding to Hurricane Matthew (October 2016, Category 4, $V_{\max} = 230$ km/h). The network has $|\mathcal{A}| = 8$ coastal origin zones, $|\mathcal{R}| = 5$ intermediate transit nodes, $|\mathcal{F}| = 6$ shelters with combined capacity $\sum_f \pi_f = 7300$ persons and $|\mathcal{E}| = 57$ directed arcs. The fleet $\mathcal{V}$ comprises 20 vehicles (4 large buses 45p, 6 medium buses 30p, 5 trucks 20p, 5 minibuses 12p), total per-trip capacity 520, with $K = 10$ maximum trips per vehicle. Three demand scenarios capture trajectory uncertainty: pessimistic ($P^{\omega_1} = 8196$, $p = 0.25$), nominal ($P^{\omega_2} = 6830$, $p = 0.50$) and optimistic ($P^{\omega_3} = 5464$, $p = 0.25$), with $\rho = 0.20$ deviation factor. The stochastic planner $\mathrm{P}_{\mathrm{I}}^{\mathrm{S}}$ identifies 46 distinct consolidated routes spanning all eight origins, filling every shelter to 100% of capacity. The decomposition for the adaptive model yields 145 route segments, reflecting the fine-grained structure required by the affine decision rules. The underlying network representation follows standard network-flow modelling principles [15].

6.2. Anchored Versus Generic Uncertainty Sets

We compare three constructions: Generic Box (Soyster)—$\sigma_{\mathrm{box}}(\lambda) = \rho\bar{\tau}_\lambda$ with $\rho = 0.25$, $\Gamma = |\Lambda|$; Generic Budget (B–S)—same σ, $\Gamma = \sqrt{|\Lambda|} \approx 6.78$; Scenario-Anchored Neutrosophic (SAN)—$\sigma(\lambda)$ from (9) with $\kappa = 0.75$ and Γ^{n} from (15) with $\beta = 1$. Table 2 reports, for each construction, the geometry of the resulting uncertainty set and the robust makespan and cost it induces on the 46 consolidated routes.

Table 2. Uncertainty-set comparison: geometry and resulting robust makespan (46 consolidated routes).

Method	$\bar{\sigma}_{\mathrm{rel}}$	$\sigma_{\max}$	Γ	η^{rob} (min)	Cost ($)	$\Delta\eta$
Generic Box	25.0%	20.3	46	352	4288	—
Generic Budget	25.0%	20.3	6.78	318	3890	−9.7%
SAN ($\kappa = 0.75$)	15.0%	12.2	7.42	305	3740	−13.4%
SAN ($\kappa = 0.50$)	10.0%	8.1	7.12	295	3614	−16.2%
SAN ($\kappa = 1.00$)	20.0%	16.2	7.68	322	3938	−8.5%
Nominal ($\rho = 0$)	0%	0.0	0	282	3430	−19.9%

The SAN set with $\kappa = 0.75$ reduces makespan by 13.4% relative to the generic box in the static context (46 consolidated routes), maintaining feasibility for all realisations in the set. The improvement stems from two sources: the anchored deviations $\sigma(\lambda)$ are route-specific and typically smaller than the uniform 25% of the generic approach, and the neutrosophic budget $\Gamma^{\mathrm{n}} = 7.42$ is calibrated to the actual indeterminacy heterogeneity. When applied to the full adaptive formulation—operating on 145 route segments—the cal-

ibration yields $\Gamma^{n} = 13.53$ (since $\sqrt{145} \approx 12.04$ with variance correction), with $\bar{\sigma}_{rel} = 15\%$ at $\kappa = 0.75$.

6.3. Neutrosophic Indeterminacy Profile of the Route Network

Table 3 presents the neutrosophic characterisation of a representative set of routes. The uniform $I_\lambda = 0.15$ (since $\kappa = 0.75$ and the maximum demand deviation is 20%, yielding a 15% distance deviation) reflects the structural correlation in this instance: all eight origin zones share the same hurricane-track geometry and the coupling parameter κ applies homogeneously to every arc. The truth degree $T_\lambda = 0.50$ follows from the symmetric three-point scenario ensemble, whose probability-weighted mean absolute deviation is exactly half of the worst-case spread; it should be read as "on a probability-weighted basis, the realised travel time sits halfway between the nominal value and its worst case", not as a complement of I_λ. In networks with heterogeneous weather exposure, I_λ would vary across routes, making the neutrosophic calibration (15) bind (see Section 6.6). All routes have $F_\lambda = 0$, confirming physical feasibility within the planning horizon $T_{max} = 480$ min; the uniform score $S(\lambda) = 0.783$ provides a single-figure operational summary. Heterogeneous instances would also exhibit a non-trivial ordering under $\succ_n$ (Definition 10) that could be exploited for priority-based fleet allocation.

Table 3. Neutrosophic characterisation $\langle T_\lambda, I_\lambda, F_\lambda \rangle$ for selected representative routes, $\kappa = 0.75$.

Route	Origin	$\bar{\tau}_\lambda$	$\sigma(\lambda)$	T_λ	I_λ	F_λ	$S(\lambda)$
λ_1	Cayo Granma	52.4	7.86	0.500	0.150	0.000	0.783
λ_{23}	El Caney	63.7	9.56	0.500	0.150	0.000	0.783
λ_{38}	San Pedrito	40.2	6.03	0.500	0.150	0.000	0.783

6.4. Per-Origin Decomposition and Scenario Evaluation

The Hierarchical Origin-Based Decomposition (DJO) applies identically to the static and adaptive formulations, since the block-diagonal structure (vehicle–origin uniqueness) is preserved under affine decision rules. Across the eight origins (demands 396–1704 persons, 2–3 vehicles, 10–40 route segments each), the largest subproblem is El Caney (1704 persons, 40 segments, 720 binaries). The DJO reduces total binaries from 24,795 (monolithic, $19 \times 145 \times 9$) to 3069 (sum of subproblems), an 87.6% dimensional reduction.

A key advantage of the ARO is the ability to adapt temporal variables to realised perturbations. Table 4 evaluates the ARO solution under each of the three P_I^S scenarios, showing how affine rules adjust the schedule while the first-stage assignment W^* remains fixed. Under the optimistic scenario ($\zeta = 0.85$), distances contract by 15% and the makespan drops to 470 ut—a 14.5% reduction from the 550 ut worst-case guarantee; the pessimistic scenario ($\zeta = 1.15$) inflates distances by 15% to 545 ut, close to the formal worst-case bound ($\Delta = -1$ ut). This near-saturation confirms that the scenario-anchored set tightly envelopes actual realisations, avoiding the excessive conservatism of generic constructions. The expected makespan $\mathbb{E}[\eta] = 506.8$ ut provides a probabilistically weighted operational estimate.

Table 4. Evaluation of the ARO schedule across P_1^S scenarios. Affine rules adapt times while first-stage W^* remains fixed; 19 vehicles active, 145 route segments.

Metric	Optimistic $\zeta = 0.85$	Nominal $\zeta = 1.00$	Pessimistic $\zeta = 1.15$	Worst-Case ($\mathcal{U}_\Gamma^\Xi$)
Makespan (ut)	470	506	545	550
Δ Makespan	−76	−40	−1	0 (ref.)
Total distance (km)	1271.6	1496.0	1720.4	—
Variable cost (USD)	381.48	448.80	516.12	—
Total cost (USD)	5381.48	5808.80	6266.12	10,616.88
Persons evacuated	7300	7300	7300	7300

6.5. Sensitivity to Budget and Coupling

Table 5 reports the dual sensitivity of the ARO. The price of robustness is modest across the entire Γ range: from 546 ut (nominal) to 550 ut at $\Gamma^n = 13.53$, a 0.7% increase, since the budget limits simultaneous deviations to ∼13 of 145 routes at $\bar{\sigma}_{\text{rel}} = 15\%$. The value $\Gamma^n = 13.53$ sits just above $\sqrt{|\Lambda|} = 12.04$, the extra 1.5 units protecting against indeterminacy heterogeneity. Moving from $\sqrt{|\Lambda|}$ to Γ^n tightens the Bertsimas–Sim violation bound (3) from $\exp(-12.04^2/290) = 60.7\%$ to $\exp(-13.53^2/290) = 53.2\%$—equivalently, it raises the certified robust-feasibility floor from 39.3% to 46.8%—at a 0.2 ut cost. We stress that 53.2% is an upper bound on the probability of budget exceedance, not a success probability; it is deliberately conservative, since the operative guarantee here is the deterministic scenario containment of Theorem 2 (Table 4: the pessimistic realisation falls 1 ut inside the worst case). On the coupling side, makespan moves only from 546 to 552 ut while cost rises from USD 5800 to USD 12,480 as κ goes from 0 to 1: in the ARO, cost is far more sensitive to uncertainty parametrisation than time. We adopt $\kappa = 0.75$ as the operational default, reflecting that hurricane road degradation is substantial but sub-proportional to demand (a 20% demand rise maps to ∼ 15% travel-time increase).

Table 5. Joint sensitivity of the ARO: budget Γ (left, SAN, $\kappa = 0.75$, $|\Lambda| = 145$) and coupling κ (right). The "viol. bound" column is the Bertsimas–Sim upper bound on the probability of budget exceedance (3), i.e., a lower bound on robust feasibility of 1 minus the listed value.

Γ	η^{rob}	$\Delta\eta/\eta^{\text{nom}}$	Viol. Bound	κ	$\bar{\sigma}_{\text{rel}}$	Γ^n	Cost ($)
0.00	546	0.0%	100.0%	0.00	0.0%	—	5800
4.00	547	0.2%	94.6%	0.25	5.0%	12.40	6890
8.00	548	0.4%	80.2%	0.50	10.0%	12.85	8240
12.04	549	0.5%	60.7%	0.75	15.0%	13.53	10,617
13.53	550	0.7%	53.2%	1.00	20.0%	13.98	12,480
20.0	551	0.9%	25.2%	—	—	—	—
145	∼552	1.1%	≈0%	—	—	—	—

6.6. Analytical Sensitivity and a Heterogeneous Instance

The Hurricane Matthew network is a single, near-homogeneous instance. To probe the framework's behaviour across the dimensions the calibration is meant to capture—network size, indeterminacy dispersion, scenario count and coupling—we exploit the fact that the uncertainty-set construction is available in closed form: Γ^n, the neutrosophic profile and the violation bound depend on the data only through $|\Lambda|$, $\bar{I}$ and $\text{cv}^2(I)$, and can, therefore, be evaluated analytically without re-solving the mixed-integer program. (The makespan and cost figures, by contrast, are specific to the solved Matthew instance and are not extrapolated here.)

https://doi.org/10.3390/engproc2026153041

Scaling of the budget. By Definition 9, $\Gamma^{n} = \sqrt{|\Lambda|}\,(1 + \beta\,\mathrm{cv}^2(I))$ with $\beta = 1$. Table 6 reports Γ^{n} over a grid of network sizes and dispersion levels. Two regularities stand out: Γ^{n} grows as $\sqrt{|\Lambda|}$ in size (sub-linear, so the budget covers a shrinking fraction of routes as the network grows), and linearly in $\mathrm{cv}^2(I)$ (so dispersion, not mean indeterminacy, drives the heterogeneity surcharge). The Matthew instance sits at $|\Lambda| = 145$, $\mathrm{cv}^2(I) \approx 0.124$, giving $\Gamma^{n} = 13.53$.

Table 6. Analytical scaling of $\Gamma^{n} = \sqrt{|\Lambda|}(1 + \mathrm{cv}^2(I))$ across network size $|\Lambda|$ and indeterminacy dispersion $\mathrm{cv}^2(I)$ ($\beta = 1$). Values are capped at $|\Lambda|$ (the full-box budget).

$\lvert\Lambda\rvert \backslash \mathrm{cv}^2(I)$	0.00 (Homog.)	0.124 (Matthew)	0.50	1.00
46	6.78	7.62	10.17	13.56
100	10.00	11.24	15.00	20.00
145	12.04	13.53	18.06	24.07
300	17.32	19.47	25.98	34.64

A heterogeneous illustrative instance. Consider a network of $|\Lambda| = 145$ segments split into two exposure classes: a fraction $a = 0.40$ of exposed coastal segments with $I = 0.30$ and the remaining 0.60 of sheltered inland segments with $I = 0.05$. The aggregate indeterminacy is $\bar{I} = 0.40(0.30) + 0.60(0.05) = 0.15$—identical to the Matthew mean—yet the dispersion is now $\mathrm{Var}(I) = 0.40(0.15)^2 + 0.60(0.10)^2 = 0.015$, so $\mathrm{cv}^2(I) = 0.015/0.15^2 = 0.667$ and $\Gamma^{n} = 12.04(1 + 0.667) = 20.07$. The same mean indeterminacy, thus, yields a budget 48% larger than the near-homogeneous Matthew case, and the Bertsimas–Sim violation bound falls from 53.2% to $\exp(-20.07^2/290) = 24.9\%$: the calibration automatically concentrates protection where dispersion—hence latent correlation—is higher. The two classes are separated under the severity preorder ($\lambda_{\text{exposed}} \succ_{n} \lambda_{\text{sheltered}}$ since $0.30 > 0.05$), with scores $S_{\text{exposed}} = (2 + T - 0.30 - F)/3$ and $S_{\text{sheltered}} = (2 + T - 0.05 - F)/3$; this ordering is precisely the lever a priority-based fleet allocation would use, and it is invisible to a mean-only or box description. This illustrates analytically why the neutrosophic calibration is consequential exactly when the network departs from homogeneity—the regime the single Matthew instance cannot exhibit.

Scenario count and coupling. The deviation $\sigma(\lambda) = \bar{\tau}_{\lambda}^{\mathrm{tp}} \max_{\omega} |\zeta^{\omega} - 1|$ depends on the ensemble only through the extreme distance factor; adding interior scenarios refines the probability concentration (hence T) but leaves σ—and, therefore, I and the set geometry—unchanged unless they extend the range. The coupling κ scales the range linearly, $\max_{\omega} |\zeta^{\omega} - 1| = \kappa \max_{\omega} |\delta^{\omega} - 1|$, as already traced in the right half of Table 5. Together these closed-form dependencies give a complete, instance-independent picture of how the constructed set responds to the three data dimensions the reviewer highlights, complementing the solved Matthew case.

The complete hierarchical chain $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}} \to \mathrm{P}_{\mathrm{II}}^{\mathrm{D}}(\mathrm{DJO}) \to \mathrm{P}_{\mathrm{I}}^{\mathrm{S}} \to \mathrm{P}_{\mathrm{II}}^{\mathrm{ARO}}$ executes in under four minutes on the IAN-style hardware (Intel® i7, 32 GB RAM, CBC solver): $\mathrm{P}_{\mathrm{I}}^{\mathrm{D}}$ takes 2.1 s, DJO Phase 1 (GAP) 0.03 s, $\mathrm{P}_{\mathrm{I}}^{\mathrm{S}}$ 8.4 s and $\mathrm{P}_{\mathrm{II}}^{\mathrm{ARO}}$ 214 s, for a total of 224.5 s. This is well within the 24–48 h operational window preceding landfall, confirming tractability even with affine decision rules and neutrosophic calibration.

6.7. Discussion: Practical Implications for Emergency Management

The computational results yield several insights for emergency planners, including one counterintuitive finding that warrants careful interpretation.

The SRO–ARO cost inversion, resolved by a same-granularity ablation. The static robust solution (Soyster) yields a lower total cost (USD 4288) than the adaptive solution (USD 10,617)—a 147.6% premium for the ARO that apparently contradicts Theorem 4.

The contradiction is only apparent because the two figures are computed at different route granularities, whereas Theorem 4 compares the two formulations at the same granularity. To isolate the value of adaptivity cleanly, Table 7 decomposes the gap into two orthogonal effects. The granularity (decomposition) effect is the surplus variable transportation cost incurred by splitting 46 consolidated routes into 145 independently-costed segments, USD 4817 − USD 920 = USD 3897; it has nothing to do with robustness and would be present even for two deterministic models. The residual robustness/adaptivity effect is what remains once granularity is held fixed. At the common 46-route granularity, Theorem 4 guarantees $Z_*^{\mathrm{ARO}} \leq Z_*^{\mathrm{SRO}}$, since every static solution is ARO-feasible with zero sensitivity coefficients; adaptivity can only reduce cost when the comparison is apples-to-apples. The reported inversion is, therefore, entirely a modelling-granularity artefact, not evidence against the dominance theorem.

Table 7. Same-granularity ablation: attribution of the SRO–ARO total-cost gap. The granularity effect (route decomposition) is orthogonal to robustness; at fixed granularity Theorem 4 gives $Z_*^{\mathrm{ARO}} \leq Z_*^{\mathrm{SRO}}$.

Model	Granularity	Var. Transport (USD)	Total (USD)	Adaptivity
SRO (Soyster)	46 routes	920	4288	no
ARO (affine)	145 segments	4817	10,617	yes
Granularity effect	46 → 145	+3897	—	—
At fixed granularity	46 routes	—	$Z_*^{\mathrm{ARO}} \leq Z_*^{\mathrm{SRO}}$	Theorem 4

Structural versus monetary value of adaptivity. The ARO's 0.7% adaptation margin (4 ut out of 550) appears negligible monetarily, but its significance is qualitative: the affine sensitivity coefficients (average 2.25, maximum 4.20) are explicit recourse instructions—if perturbation ξ_λ materialises on segment λ, shift the next trip's departure by t^λ units—an operational transparency absent from the static solution, which commits irrevocably to worst-case durations.

When is adaptivity worth the cost? The framework's value lies not in unconditionally preferring ARO over SRO but in providing a principled mechanism—the (T, I, F) characterisation and the derived Γ^{n}—to quantify the price of adaptivity. For the homogeneous Matthew instance ($I_\lambda = 0.15$) the premium is high and the flexibility gain modest; under heterogeneous exposure (exposed coastal routes $I_\lambda \sim 0.30$ versus sheltered inland $I_\lambda \sim 0.05$) the calibration yields route-specific budgets that concentrate adaptive protection where needed, making the ARO advantage more pronounced (Section 6.6).

Communicability. A potential modelling convenience—not a universal advantage—is that the planner can report a single tripartite summary rather than an isolated makespan: here, "confidence $T = 0.50$ that the route runs near nominal, indeterminacy $I = 0.15$ absorbed by the adaptive rules, falsity $F = 0$, against a violation bound of 53.2%". We found this $\langle T, I, F \rangle$ format easier to convey to operational staff than a lone probability bound, and present it as an empirical observation about crisis communication, not a claim of inherent superiority over scalar reporting.

Table 8 positions the framework against the most relevant recent contributions along five axes. To our knowledge, none of these works combines all five—stochastic demand, adaptive temporal robustness, data-driven uncertainty-set construction, origin-based decomposition and neutrosophic calibration—in a single hierarchical pipeline; we state this as a gap in the surveyed literature rather than as an absolute priority claim, since the five ingredients individually are well established. The closest works isolate subsets of these axes: Goerigk et al. [16] address temporal robustness with box sets but without stochastic

demand, adaptivity or calibration; Bayram and Yaman [9] model stochastic demand with dynamic resource allocation but use no temporal uncertainty set; Eriskin and Karatas [10] and Li et al. [17] build budgeted robust sets but with a fixed, non-anchored geometry. Our specific increment is the construction of the set—anchoring its size to the scenario ensemble and its budget to a neutrosophic indeterminacy profile—rather than the use of any single ingredient in isolation.

Table 8. Comparison with recent evacuation optimisation frameworks. SP: stochastic programming; RO: robust optimisation; DRA: dynamic resource allocation.

Work	Demand	Temporal	Adaptive	$\mathcal{U}$-Design	Decomp.	Neut.
Bayram–Yaman [18]	SP	—	—	—	Benders	—
Goerigk et al. [16]	—	RO	—	Box	—	—
Sun et al. [7]	Scen.	—	—	Scen.	—	—
Eriskin–Karatas [10]	—	RO	—	Budget	—	—
Bayram–Yaman [9]	SP	—	—	—	DRA	—
Zhang et al. [17]	RO	—	—	Budget	—	—
This work	**SP**	**ARO**	**Affine**	**Anchored**	**DJO**	**Yes**

7. Conclusions

We introduced a neutrosophic-informed framework for constructing uncertainty sets in adaptive robust pre-disaster evacuation planning. Scenario anchoring derives temporal deviations $\sigma(\lambda)$ from the P_I^S ensemble, producing a set tied to the operational environment rather than generic worst cases; in the static context it cuts makespan conservatism by 13.4% relative to Soyster boxes while preserving feasibility (Theorem 2). Neutrosophic calibration maps the (T, I, F) profile onto the set geometry: for Matthew (145 segments), $\Gamma^n = 13.53$ (Definition 9) replaces the ad-hoc $\sqrt{|\Lambda|}$ heuristic with a data-driven surrogate, tightening the Bertsimas–Sim violation bound from 60.7% to 53.2% (feasibility floor 46.8%) at a 0.7% premium, while deterministic containment provides the operative guarantee. Affine decision rules deliver scenario-dependent makespans of 470–545 ut under fixed first-stage assignments. The framework quantifies the price of adaptivity once the SRO–ARO granularity artefact is isolated, and the full pipeline runs in under four minutes—well within emergency decision windows.

Several extensions remain open. Interval-valued neutrosophic sets [19] would let the (T, I, F) degrees themselves be uncertain, modelling imprecise uncertainty levels; ARO formulations on the same consolidated route structure as the SRO would isolate the pure benefit of affine adaptivity from decomposition effects; distributionally robust ambiguity sets [11] informed by the neutrosophic profile offer a middle ground between stochastic and robust models; equity-weighted objectives using neutrosophic priority measures would strengthen the humanitarian dimension; and real-time reoptimisation, with $\mathcal{U}$ reshaped as hurricane-track forecasts update, would turn the offline pipeline into an adaptive decision-support tool.

Author Contributions: Conceptualization, Y.F.-F. and S.M.A.-A.; methodology, formal analysis, investigation and original draft preparation, Y.F.-F.; review and editing, S.M.A.-A., R.M.-P. and G.B.-A.; supervision, S.M.A.-A. and G.B.-A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The software implementation supporting the case study—Lab Virtual runner, MGEA scenario generator and the complete $P_I^D \rightarrow P_{II}^D \rightarrow P_I^S \rightarrow P_{II}^R \rightarrow P_{II}^{ARO}$ chain—is openly available in the Zenodo software archive [20], https://doi.org/10.5281/zenodo.19139457.

Acknowledgments: This work is part of the doctoral thesis of Y.F.-F. at the Faculty of Mathematics and Computation of the Universidad de La Habana. The authors acknowledge the Research Department and Computer Science Program at the Universidad Politécnica Estatal del Carchi (Ecuador) and the Center for Optimization and Applications at the Universidad de Cienfuegos (Cuba).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Kall, P.; Wallace, S.W. *Stochastic Programming*; Wiley: Chichester, UK, 1994.
2. Ben-Tal, A.; El Ghaoui, L.; Nemirovski, A. *Robust Optimization*; Princeton University Press: Princeton, NJ, USA, 2009.
3. Kamyabniya, A.; Sauré, A.; Salman, F.S.; Bénichou, N.; Patrick, J. Optimization models for disaster response operations: A literature review. *OR Spectr.* **2024**, *46*, 737–783. [CrossRef]
4. Soyster, A.L. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* **1973**, *21*, 1154–1157. [CrossRef]
5. Ben-Tal, A.; Nemirovski, A. Robust optimization—Methodology and applications. *Math. Program.* **2002**, *92*, 453–480. [CrossRef]
6. Bertsimas, D.; Sim, M. The price of robustness. *Oper. Res.* **2004**, *52*, 35–53. [CrossRef]
7. Sun, H.; Li, J.; Wang, T.; Xue, Y. A novel scenario-based robust bi-objective optimization model for humanitarian logistics network under risk of disruptions. *Transp. Res. E* **2022**, *157*, 102578. [CrossRef]
8. Siddig, M.; Song, Y. Multistage stochastic programming with a random number of stages: Applications in hurricane disaster relief logistics planning. *Eur. J. Oper. Res.* **2025**, *321*, 925–941. [CrossRef]
9. Bayram, V.; Yaman, H. A joint demand and supply management approach to large scale urban evacuation planning: Evacuate or shelter-in-place, staging and dynamic resource allocation. *Eur. J. Oper. Res.* **2024**, *313*, 171–191. [CrossRef]
10. Eriskin, L.; Karatas, M. Applying robust optimization to the shelter location-allocation problem: A case study for Istanbul. *Ann. Oper. Res.* **2024**, *339*, 1589–1635. [CrossRef]
11. Bertsimas, D.; Sim, M.; Zhang, M. Adaptive distributionally robust optimization. *Manag. Sci.* **2019**, *65*, 604–618. [CrossRef]
12. Lorca, A.; Sun, X.A. Multistage adaptive robust optimization for the unit commitment problem. *Oper. Res.* **2016**, *64*, 32–51. [CrossRef]
13. Smarandache, F. *A Unifying Field in Logics: Neutrosophic Logic. Neutrosophy, Neutrosophic Set, Neutrosophic Probability and Statistics*; American Research Press: Rehoboth, NM, USA, 2005.
14. Ahmad, F. Interactive neutrosophic optimization technique for multiobjective programming problems: An application to pharmaceutical supply chain management. *Ann. Oper. Res.* **2022**, *311*, 551–585. [CrossRef]
15. Ahuja, R.K.; Magnanti, T.L.; Orlin, J.B. *Network Flows: Theory, Algorithms, and Applications*; Prentice Hall: Upper Saddle River, NJ, USA, 1993.
16. Goerigk, M.; Deghdak, K.; T'Kindt, V. A two-stage robustness approach to evacuation planning with buses. *Transp. Res. B* **2015**, *78*, 66–82. [CrossRef]
17. Li, H.; Yu, D.; Zhang, Y.; Yuan, Y. A two-stage robust optimization model for emergency service facilities location-allocation problem under demand uncertainty and sustainable development. *Sci. Rep.* **2025**, *15*, 2895. [CrossRef] [PubMed]
18. Bayram, V.; Yaman, H. Shelter location and evacuation route assignment under uncertainty: A Benders decomposition approach. *Transp. Sci.* **2018**, *52*, 416–436. [CrossRef]
19. Smarandache, F. *Introduction to Neutrosophic Statistics*; Sitech and Education Publisher: Craiova, Romania, 2014.
20. Fernández-Fernández, Y. *Metamodelo de Optimización Robusta para Planificación de Evacuaciones ante Huracanes*, version 1.0.0; GitHub: Geneva, Switzerland; Zenodo: Geneva, Switzerland, 2026. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.