Supplementary Information for The Value of AI-Driven Predictive Intelligence in Wildfire Relief Logistics: Evidence from the 2026 Ontario Thunder Bay 36 Wildfire

Hooman Razavi[1,2], Mostafa Hajiaghaei-Keshteli[1], Omid Titidezh[3,*], Hossein Bonakdari[2], M. Hamed Mozaffari[4,5], and Neale R. Smith[6]

[1]Tecnologico de Monterrey, School of Engineering and Sciences, Puebla, 72453, Mexico
[2]University of Ottawa, Faculty of Engineering, Ottawa, K1N 6N5, Canada
[3]Shomal University, Department of Civil Engineering, Amol, 46161-558, Iran
[4]National Research Council Canada, Construction Research Center, Ottawa, K1A 0R6, Canada
[5]Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON, M5S 1A1, Canada
[6]Tecnologico de Monterrey, School of Engineering and Sciences, Monterrey, 64849, Mexico
[*]Corresponding author: titidezhomid@gmail.com

Contents. This Supplementary Information accompanies the main article and collects material referenced from it: perimeter-level validation of the hazard surrogate (Supplementary Note A); the solution algorithms, the Monte Carlo validation protocol, objective weighting, and the complexity analysis (Supplementary Notes B–E); and the diagnostic figures, per-seed solver results, and the full sensitivity and reachability tables (Supplementary Note F). Supplementary Figures are numbered S1–S5, Supplementary Tables S1–S6, and Supplementary Algorithms S1–S3.

Supplementary Note A. Perimeter-level validation of the hazard surrogate

The main-text perimeter-validation figure shows the surrogate's fidelity at the *perimeter* level for the two corridor-driving fires, with per-fire IoU and Hausdorff distances averaged over each fire's reporting days. For a single representative snapshot the geometric fit is weaker still (IoU $= 0.000$ and Hausdorff 79.4 km for the 375,922 ha Thunder Bay 36 complex, which no single elliptical lobe can represent; IoU $= 0.009$ and Hausdorff 10.7 km for Savant Lake 89). We report this candidly: the surrogate is not a perimeter-reconstruction tool, and its value for relief routing derives from closure timing and conservatism rather than geometric fit.

The intersection-over-union (IoU) measures the overlap between the simulated and observed burned areas (one denotes perfect overlap, zero a disjoint pair), while the Hausdorff distance measures the largest separation between the two perimeter boundaries; both are computed after projecting the CWFIS perimeters and the double-ellipse footprint into a common UTM-kilometre frame, and both are averaged over every reporting day of each fire. The operative point is that relief routing depends only on *which* road arcs a fire severs and *when*, not on the precise shape of the burn, so this perimeter-level weakness does not undermine the high-precision, well-timed closure prediction reported in the main text (closure ROC-AUC 0.797 over the full horizon).

Supplementary Note B. Solution algorithms

For completeness and reproducibility, Supplementary Algorithms S1 and S2 reproduce the quantum-inspired genetic algorithm and its bidirectional-shuffle decoder as used in this study, following the design of Razavi et al. [20]. They are included for transparency and are not a contribution of the present paper; the only problem-specific adaptation is the objective-driven slot placement that enables deliberate service-timing decisions (main text, Methodology).

In brief, each candidate solution is encoded as a string of qubits whose measurement yields both the subset of communities to attempt and the order in which they are shuffled into daily routes. The qubit amplitudes are updated each generation by a rotation gate that biases sampling toward the best solutions found so far, so the population concentrates on high-quality regions of the search space without an explicit crossover operator. After a chromosome is measured, the bidirectional-shuffle decoder (Supplementary Algorithm S2) assembles feasible vehicle routes and places service times, and a greedy-displacement repair restores capacity, time-window, and admissibility feasibility, dropping must-serve communities only when the hazard layer has rendered them structurally unreachable.

Algorithm S1 Quantum-inspired genetic algorithm for the wildfire relief routing problem, reproduced from Razavi et al. [20].

initialize quantum population $Q(0)$: $2n$ qubits per individual, $(\alpha, \beta) = (1/\sqrt{2}, 1/\sqrt{2})$
for $g = 1$ **to** G **do**
  for each individual $q \in Q(g-1)$ **do**
    measure q: bit $b_i = 1$ with probability β_i^2, giving selection s and shuffle order o
    $s \leftarrow s \cup P^h$ ▷ must-serve communities always attempted
    routes $\leftarrow$ DECODE(s, o) ▷ Supplementary Algorithm S2
    routes $\leftarrow$ GREEDYDISPLACEMENTREPAIR(routes)
    evaluate fitness by the objective, penalizing any unserved must-serve community
  end for
  update best-so-far solution p^*
  for each individual, each qubit i with observed bit $\neq p^*$'s bit **do**
    rotate (α_i, β_i) by $\pm\Delta\theta$ toward p^*'s bit ▷ $\Delta\theta = 0.05\pi$
  end for
end for
return p^*

Algorithm S2 Bidirectional-shuffle decoding with objective-driven split placement for service timing, reproduced from Razavi et al. [20].

Seed $\leftarrow$ selected communities in ascending index order; $\ell \leftarrow 0$; $r \leftarrow |\text{Seed}| - 1$; Seq $\leftarrow ()$
for each bit b of the shuffle order **do**
  if $b = 0$ **then** append Seed$[\ell]$ to Seq; $\ell \leftarrow \ell + 1$
  else append Seed$[r]$ to Seq; $r \leftarrow r - 1$
  end if
end for
for each community c in Seq **do**
  for every vehicle-day slot (k, t) and both route ends (append, prepend) **do**
    check feasibility: capacity Q, time window $[e_c, l_c]$, duration W, admissibility $\varphi_{ijt} \geq \varphi^{\min}$, risk budget Γ
    compute marginal gain $\Delta = \pi_{ct} - \lambda_2\,\Delta\text{cost} - \lambda_3\,\Delta\text{risk}$
  end for
  place c at the feasible position with maximum Δ; leave c unplaced if none is feasible
end for
return routes

Supplementary Note C. Monte Carlo validation protocol

Supplementary Algorithm S3 summarizes the Monte Carlo protocol. For each replication, one physically consistent fire scenario is drawn (a single weather realization per ignition, not the planning-stage ensemble mean), and the closed or open state of every arc on every day follows from testing its road-path waypoints against that same realized perimeter. This correlated construction ensures that arcs sharing physical road segments fail together, preventing recourse policies from escaping through detours that are in reality the same road; an early independent-per-arc sampler was found to permit exactly such phantom detours and was replaced. Closures are realized per day and held fixed within the day, so within-day order effects arise from the visiting sequence, not from fluctuating closure states. Each plan is executed under three recourse rules: skip-and-continue (primary), in which a vehicle facing a closed arc reroutes along the cheapest currently open path or, failing that, skips the stop and continues; route truncation; and truncation with vehicle stranding, in which a vehicle that cannot return to the depot is lost for the remaining days. The same drawn realizations are applied to the risk-aware ($\lambda_3 > 0$) and risk-blind ($\lambda_3 = 0$) plans, making paired Wilcoxon tests valid. Reported metrics are realized demand coverage (overall and for P^h), reroute and stranding counts, realized cost, and the price of resilience, the relative planned-cost increase of the risk-aware plan. In addition to the sampled scenarios, each plan is replayed deterministically against the recorded July 2026 closure timeline of Highways 599 and 527. Two systematic sensitivity analyses complete the protocol: a schedule-slack sweep over fleet size and vehicle capacity across the realistic range for the event's demand volume, and a planning-start-date sweep along the real fire timeline, which together locate the regime boundary at which risk-aware and risk-blind plans diverge.

Algorithm S3 Correlated Monte Carlo validation of executed relief plans under recourse.

for $r = 1$ **to** R **do**
  draw one weather realization per ignition; grow realized perimeters over days $1,\dots,\bar{T}$
  for each arc (i,j) and day t **do**
    closed$(i,j,t) \leftarrow$ **true** iff any road-path waypoint of (i,j) lies inside the day-t realized perimeter buffer
  end for
  for each plan $\in$ {risk-aware, risk-blind} **do** ▷ same realization for both (paired)
    execute plan under the recourse rule (skip-and-continue primary; truncation and stranding as robustness)
    record demand coverage (overall, P^h), reroutes, strandings, realized cost
  end for
end for
report means, 95% confidence intervals, paired Wilcoxon tests, and price of resilience

Supplementary Note D. Objective weighting and normalization

The three objective terms, priority-weighted coverage, transportation cost, and arc risk, are measured in different units, so the weights $\lambda_1, \lambda_2, \lambda_3$ are set in two steps rather than chosen ad hoc. First, each term is *normalized* to a comparable $[0,1]$ scale by dividing it by its nominal range on the instance (maximum attainable prize; cost of the longest admissible plan; maximum per-plan risk exposure $|K|\,\bar{T}\,\Gamma$), so the raw weights carry no hidden unit effects. Second, the normalized trade-off is explored rather than fixed: the coverage-versus-cost frontier is traced by ε-constraint optimization on the exactly solvable instances (maximizing coverage subject to a swept cost bound, all other constraints intact), which recovers the true Pareto front including any non-convex regions a pure weighted sum would miss; and the risk weight λ_3 is swept from the risk-blind optimum to the point where the risk budget Γ becomes binding. Reported results therefore do not hinge on a single choice of λ_3: the regime analysis reports outcomes across the whole λ_3 sweep, accompanied by a weight-sensitivity analysis. We adopt the ε-constraint front as the primary trade-off instrument because it is exact and interpretable, and use the normalized weighted sum only as the scalarization the metaheuristics optimize at scale.

Supplementary Note E. Complexity and scalability

The model generalizes the CVRPTW and is therefore NP-hard, so exact solution is inherently limited in scale. With $n = |P|$ communities, $m = |K|$ vehicles, and $\bar{T}$ planning days, the dominant decision variables are the routing variables x_{ijkt}, numbering $O(n^2 m\bar{T})$, together with $O(nm\bar{T})$ assignment variables y_{ikt} and $O(nm\bar{T})$ continuous time variables. The binding constraint families scale as $O(n^2 m\bar{T})$ (time propagation and admissibility) and $O(nm\bar{T})$ (flow, capacity, time windows, risk budget), so total model size grows quadratically in the number of communities and linearly in fleet size and horizon. For the real instance ($n = 12$, $m = 4$, $\bar{T} = 7$) this yields on the order of 4,000 binary routing variables; the exact solver reaches proven optimality up to roughly $n = 20$ within the time budget, beyond which the branch-and-bound tree grows too large. The quantum-inspired metaheuristic sidesteps this growth: its per-generation cost is dominated by decoding, which is $O(n^2 m\bar{T})$ per individual, so the whole search scales polynomially and handles n up to at least 60 within minutes. Two practical implications follow. First, for a single event of realistic size, tens of affected communities and a fleet of a handful of vehicles, the framework returns high-quality solutions well within a daily planning cycle. Second, the quadratic growth in n means province- or multi-event-scale problems (hundreds of communities) would require decomposition, for example by geographic cluster or by fire; the arc-availability interface supports this naturally, because each cluster's hazard field is computed independently.

Supplementary Note F. Supplementary results and figures

This note collects the diagnostic figures, per-seed solver results, and the full sensitivity and reachability tables referenced from the main-text Results; the headline figures and tables remain in the main text.

Table S1 reports the per-size solver comparison in full. On the ten-community instances the exact solver (CBC) proves optimality on eight of ten seeds, and the QGA reaches within about 2% of that optimum in roughly one second; at $n = 15$ and $n = 20$ CBC no longer proves optimality within its 150-second per-seed budget, yet the QGA matches or exceeds CBC's time-limited incumbent while remaining roughly twice as good as the GA, PSO, and ALNS baselines. Figures S1 and S2 then give the full convergence and machine-learning diagnostics behind the main-text summaries.

Table S1. Solver comparison on synthetic instances (10 seeds each, equal budget; CBC under a 150 s per-seed limit); gap is measured against CBC's best bound.

n	Solver	Objective	CPU (s)	Gap (%)	Opt. (/10)
10	Exact (CBC)	−8.84	57.0	42.3	8
10	QGA	−9.04	1.3	45.8	–
10	GA	−19.69	2.9	312.4	–
10	PSO	−19.28	2.8	301.7	–
10	ALNS	−19.28	3.7	286.1	–
15	Exact (CBC)	−7.00	150.0	307.7	0
15	QGA	−7.04	2.8	303.4	–
20	Exact (CBC)	−11.41	147.2	189.2	0
20	QGA	−7.55	6.0	159.1	–

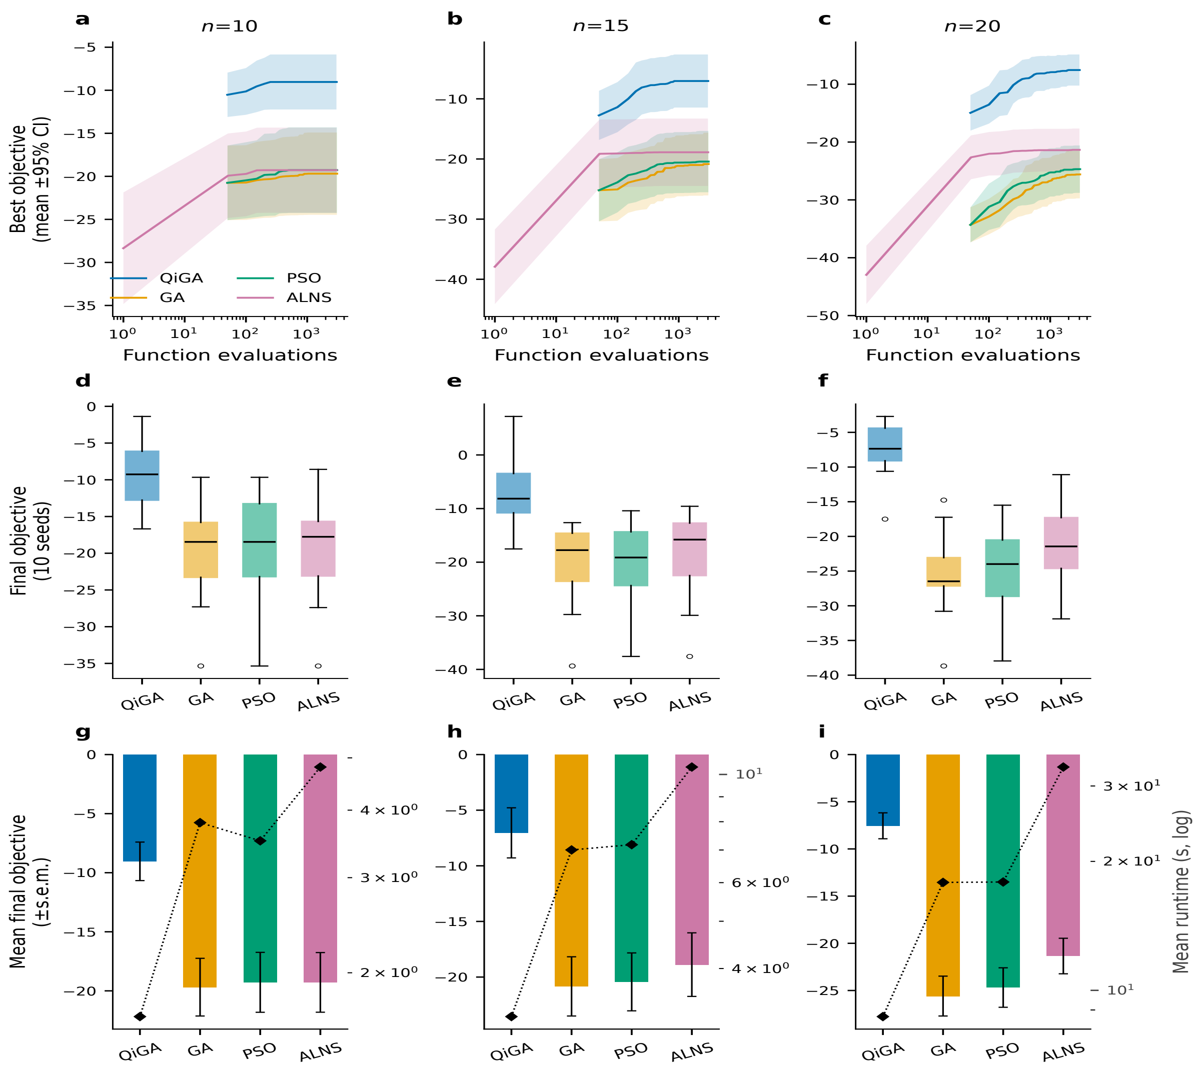

Figure S1. Convergence and solution-quality diagnostics for the metaheuristic comparison across instance sizes $n = 10, 15, 20$ (ten random seeds each, equal function-evaluation budget). (a–c) Best-objective convergence versus function evaluations, with 95% confidence bands, for the QGA, GA, PSO, and ALNS; the QGA (blue) converges fastest and to the best value. (d–f) Distributions of the final objective over the ten seeds. (g–i) Mean final objective (± s.e.m., left axis) and mean runtime (right axis). The objective is reported as a negative maximization value, so lower is better.

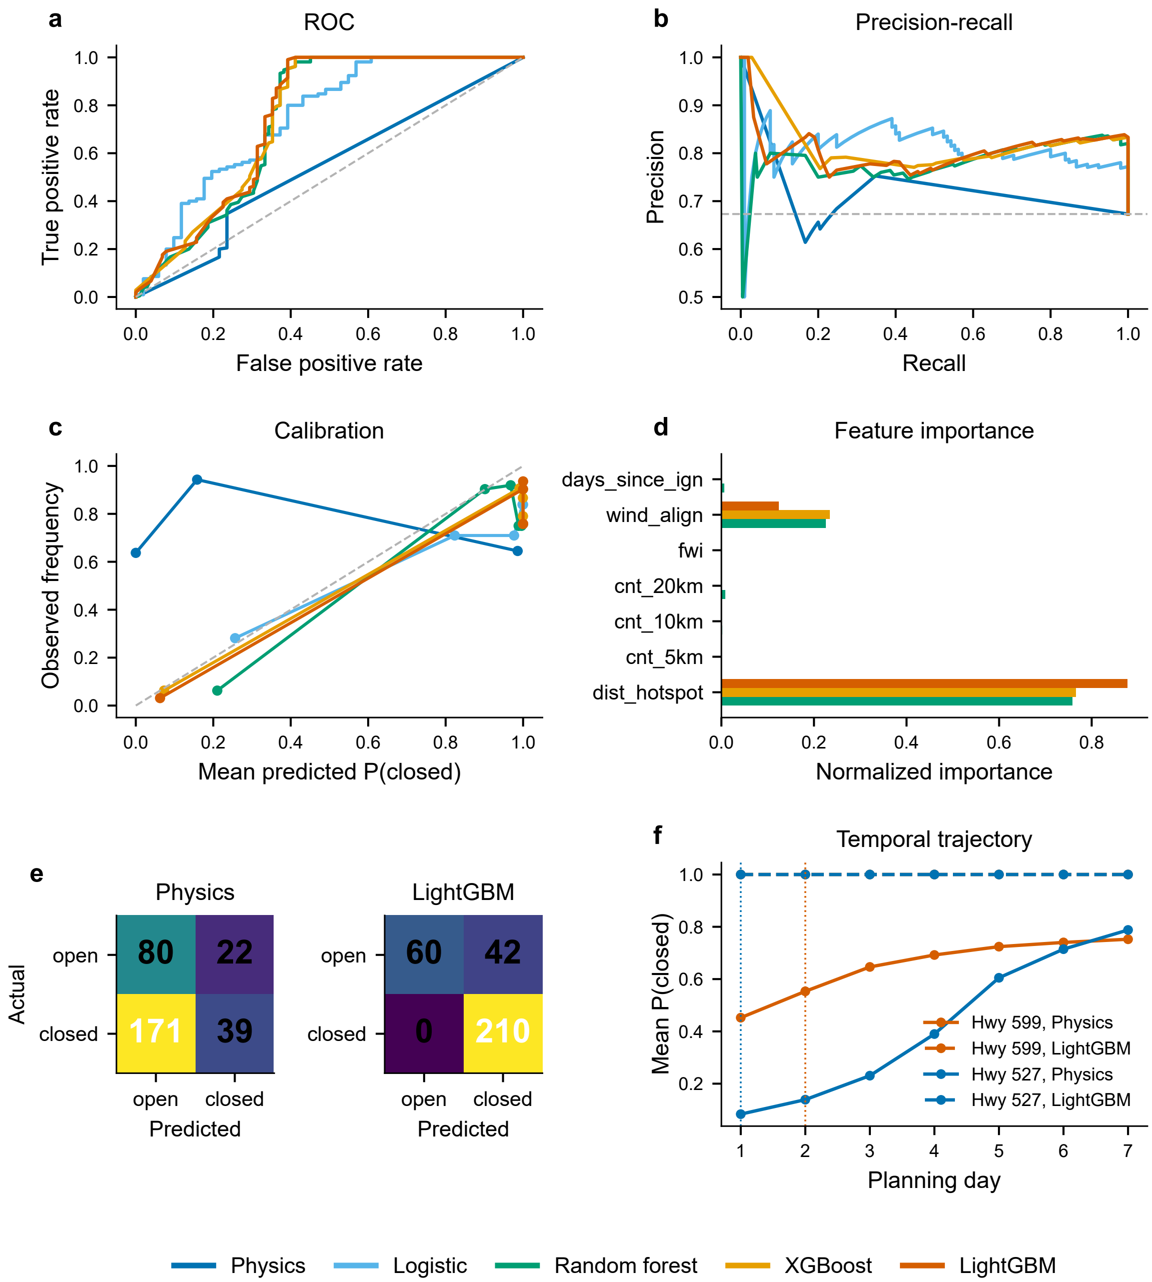

Figure S2. Machine-learning hazard-classifier diagnostics on the matched two-day closure-onset window ($n = 312$ arc-days), for the logistic, random-forest, XGBoost, and LightGBM classifiers alongside the physics ensemble. (a) ROC curves, (b) precision–recall curves, and (c) reliability (calibration) curves. (d) Normalized feature importance for the learned models. (e) Confusion matrices at the operating threshold, contrasting the conservative physics ensemble with the high-recall LightGBM classifier. (f) Predicted closure-probability trajectories along the Highway 599 and 527 corridors over the horizon, with the observed closure days marked. The learned models are far more accurate classifiers yet forbid many more arcs, which is why they induce lower realized coverage.

The next three tables summarize the resilience of the executed plans under recourse. Table S2 reports realized demand coverage, cost, and the price of resilience under the three recourse rules over 500 correlated Monte Carlo realizations of the real instance; the risk-aware and risk-blind plans are statistically indistinguishable under every rule. Table S3 isolates the effect of planning timing by sliding the planning-horizon start across the real evacuation week, and Table S4 replays the plan deterministically against the recorded Highway 527 and 599 closure dates as an operational stress test.

Table S2. Resilience comparison, real 12-community instance, QGA plans, 500 correlated Monte Carlo realizations (mean ± 95% CI); risk-aware and risk-blind plans are identical.

Recourse rule	Plan	% served	Realized cost ($)	Price of resilience
Skip-and-continue	Risk-aware	75.6 ± 0.0	7081.3	+0.0%
	Risk-blind	75.6 ± 0.0	7081.3	–
Truncation	Risk-aware	75.6 ± 0.0	7081.3	+0.0%
	Risk-blind	75.6 ± 0.0	7081.3	–
Truncation + stranding	Risk-aware	75.6 ± 0.0	7081.3	+0.0%
	Risk-blind	75.6 ± 0.0	7081.3	–

Table S3. Temporal schedule-slack sweep: planning-horizon start (offset from 2026-07-13) versus whether risk-aware and risk-blind plans diverge. No divergence occurs across the range.

Offset	Day-1 date	φ (depot→Armstrong)	Identical?
0	2026-07-13	1.000	Yes
1	2026-07-14	1.000	Yes
2	2026-07-15	1.000	Yes
3	2026-07-16	1.000	Yes
4	2026-07-17	0.983	Yes
5	2026-07-18	0.850	Yes
6	2026-07-19	0.617	Yes

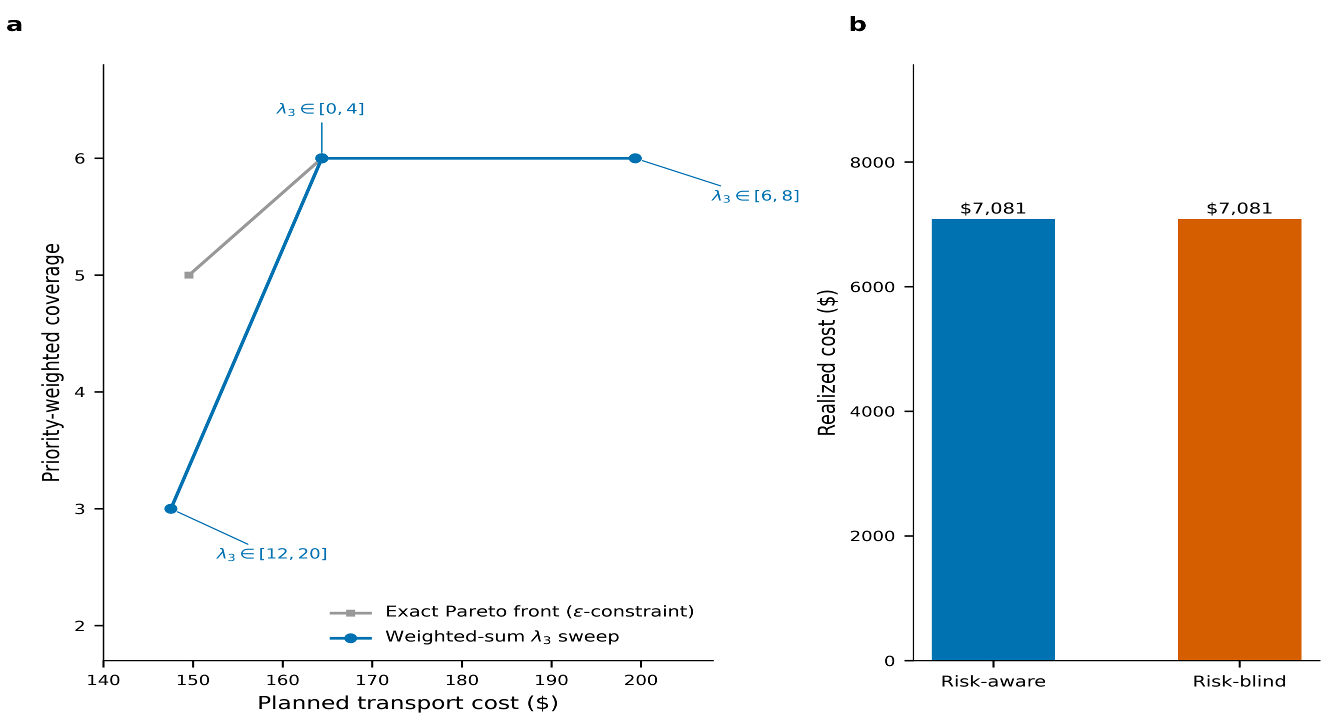

Figure S3. Cost-of-robustness analysis. (a) Coverage-cost trade-off on the toy instance. Increasing λ_3 replaces a risky shortcut with a safe detour, producing a 21.2% price of resilience. (b) Realized costs of risk-aware and risk-blind plans. In the real 2026 instance, both plans coincide, reflecting the absence of meaningful route alternatives.

Table S4. Actual-2026 replay: the framework's plan replayed deterministically against the real Hwy 527 (2026-07-13) and Hwy 599 (2026-07-14) closures, an operational stress test.

Recourse rule	% served	Realized cost ($)
Skip-and-continue	57.3	5164.9
Truncation	57.3	5164.9
Truncation + stranding	57.3	5164.9

A positive price of resilience appears only when the network offers a genuine choice between a risky and a safe route. Figure S3 contrasts the real instance, which supplies no such alternative, with a constructed toy instance in which sweeping the risk weight λ_3 does trade cost for robustness.

Because per-community relief quantities are estimated rather than measured, Table S5 checks that the coincidence of the two objectives is not an artefact of the assumed demand level. Scaling total demand by half and by one-and-a-half leaves the plans identical and the two single-corridor communities unserved at every scale, confirming that the finding rests on schedule slack and network topology, not on the absolute pallet counts.

Table S5. Demand-sensitivity sweep: total demand scaled $\times 0.5/1.0/1.5$. Plans stay identical and Mishkeegogamang and Pickle Lake remain unserved, confirming a topology property; totals reflect per-community rounding.

Scale	Total demand	% served (aware)	% served (blind)	Plans identical?
$\times 0.5$	45	77.8	77.8	Yes
$\times 1.0$	82	75.6	75.6	Yes
$\times 1.5$	122	75.4	75.4	Yes

To show that the regime boundary generalizes beyond this single event, Figure S4 places all 184 synthetic instances in the three-dimensional space of fleet capacity, closure day, and schedule-slack margin, where a single plane cleanly separates the coincident from the diverging instances.

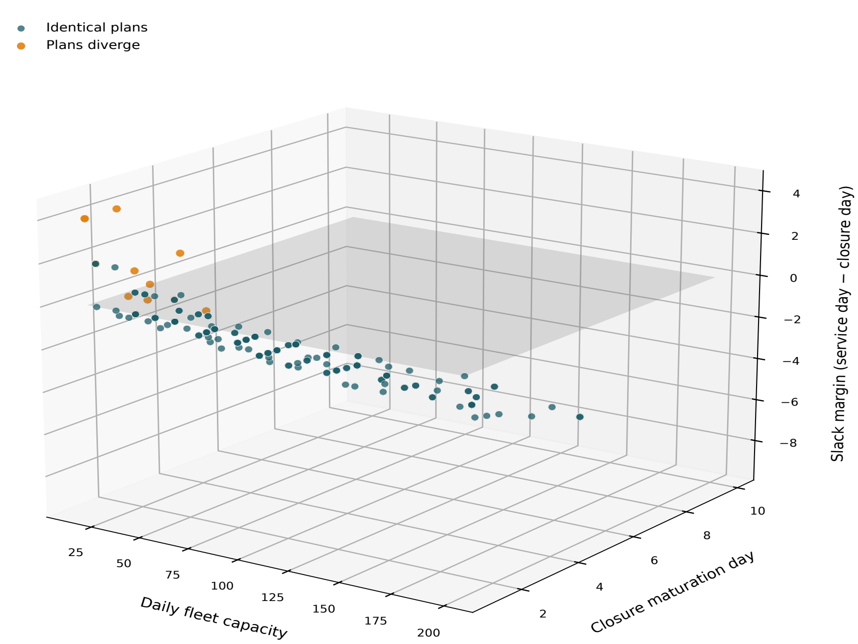

Figure S4. 3D view of the regime boundary across 184 synthetic instances over daily fleet capacity, closure-maturation day, and schedule-slack margin. Points represent instances where risk-aware and risk-blind plans match (teal) or diverge (orange), establishing the zero-margin plane as the decision boundary for anticipatory routing.

Finally, Figure S5 shows the executed routes on the real network and Table S6 traces per-community admissibility day by day, making the discrete reachability cliffs of the single-corridor communities explicit.

Figure S5. Executed relief plan for the real twelve-community instance under realistic parameters, showing depot routes from Thunder Bay. Communities are marked as served, unattempted, or unreachable. Mishkeegogamang and Pickle Lake are unattempted due to Highway 599 closure from day 3, resulting in identical risk-aware and risk-blind routes.

Table S6. Per-community network admissibility over planning days 1–7: whether any depot-to-community path uses only $\varphi_{ijt} \geq \varphi_{\min} = 0.35$ arcs. ✓ = reachable.

Community	Pr.	D1	D2	D3	D4	D5	D6	D7
Armstrong	3	✓	✓	✓	✓	✓	–	–
Whitesand FN	3	✓	✓	✓	✓	✓	–	–
Collins FN	3	✓	✓	✓	✓	–	–	–
Gull Bay FN	3	✓	✓	✓	✓	✓	✓	✓
Lac des Mille Lacs FN	3	✓	✓	✓	✓	✓	✓	✓
Lac La Croix FN	3	✓	✓	✓	✓	✓	✓	✓
Mishkeegogamang FN	3	✓	✓	–	–	–	–	–
Pickle Lake	2	✓	✓	–	–	–	–	–
Atikokan	2	✓	✓	✓	✓	✓	✓	✓
Upsala	2	✓	✓	✓	✓	✓	✓	✓
Savant Lake	2	✓	✓	✓	✓	✓	✓	✓
Sioux Lookout	1	✓	✓	✓	✓	✓	✓	✓