UNIVERSITY OF LAGOS · DEPARTMENT OF SYSTEMS ENGINEERING

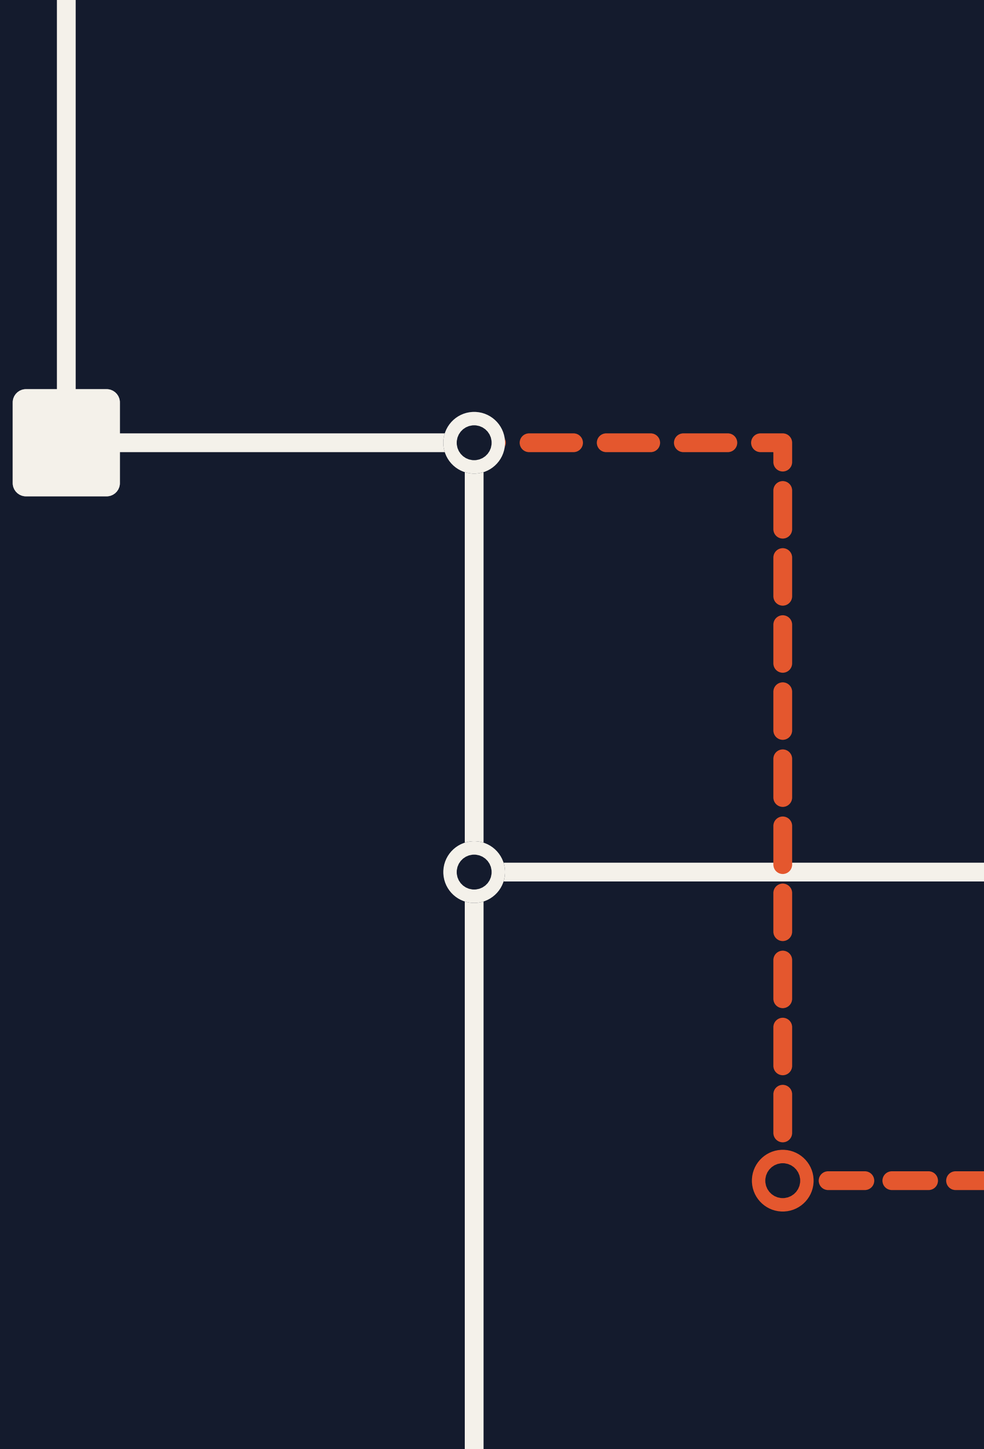
Stochastic–Robust Vehicle Routing Optimization for Last-Mile Delivery

Under Demand and Travel-Time Uncertainty

Ajayi Daniel Olawale

Supervisor: Dr. John Ogbemhe

FEIC Conference



A delivery plan assumes tomorrow is known.

It isn't.

43–95%

of days, the cost-optimal plan overloads at least one vehicle — five benchmark instances, 2,000 simulated days each

97–100% load on its fullest vehicle: minimising distance packs vans full



No capacitated last-mile model treats stochastic demand and robust travel time together — or tells an operator when protection is worth its cost.

AIM

Develop and evaluate an integrated stochastic–robust routing model.

Two established answers, used apart

PLANS FOR THE AVERAGE DAY — GUARDS AGAINST THE WORST DAY

Stochastic programming

Optimises the average over many scenarios and prices failure — but says nothing about the bad day.

Robust optimization

Survives every outcome inside a set — and pays for that protection whether or not the bad day is likely.

RESEARCH QUESTIONS

RQ1 Limits of existing methods

RQ2 How to improve the baseline models

RQ3 How effective the integrated model is

RQ4 Performance as uncertainty varies

RQ5 Cost, reliability and computing trade-offs

Four models, one routing skeleton, one scorecard

INPUT

1

Benchmarks

CVRPLIB P- and E-instances with published optimal solutions, up to 22 customers

UNCERTAINTY

2

Augmentation

Normal demand, LogNormal travel time, cv = 0.25, κ = 2, T_max = 1.1 × longest route

MODELS

3

Four models

One shared routing skeleton: differences come only from the uncertainty treatment

- **DET** deterministic
- **SP** stochastic
- **RO** robust
- **SR** integrated

SOLVE

4

Exact MILP

Python, PuLP and HiGHS; robust counterparts by LP duality; 60 in-sample scenarios

TEST

5

Out of sample

2,000 fresh scenarios and one cost basis: routing + vehicles + θ · E[overflow]

REPORT

6

Metrics

Capacity and time service levels, Price of Robustness, θ*, worst overflow

VALIDATED

Reproduces published CVRPLIB optima exactly on six instances, with zero gap · an independent checker re-verifies every plan · 16 automated tests

Each uncertainty, treated by its nature

DEMAND



Varies both ways

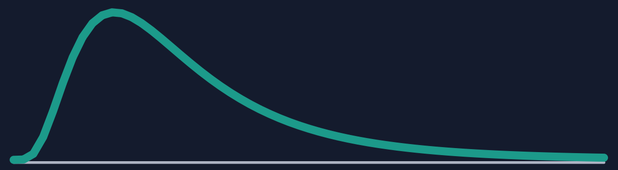
Two-stage recourse: routes are fixed today; in each scenario, load above a van's capacity is an overflow priced at θ. A robust capacity budget Γd adds protection.

ONE EXACT MILP

```
min  routing + vehicles
     + θ · E[overflow]
```

- 3-index routing with MTZ subtour elimination
- Robust capacity with budget Γd
- Robust duration ≤ T_max with budget Γt
- Overflow per vehicle, per scenario

TRAVEL TIME



Only runs late

Budgeted robust set: up to Γt roads on a route take their worst-case delay at once, and the route must still finish within T_max. LP duality keeps it linear.



Γ_d

How many customers on a route may order more than forecast at once?



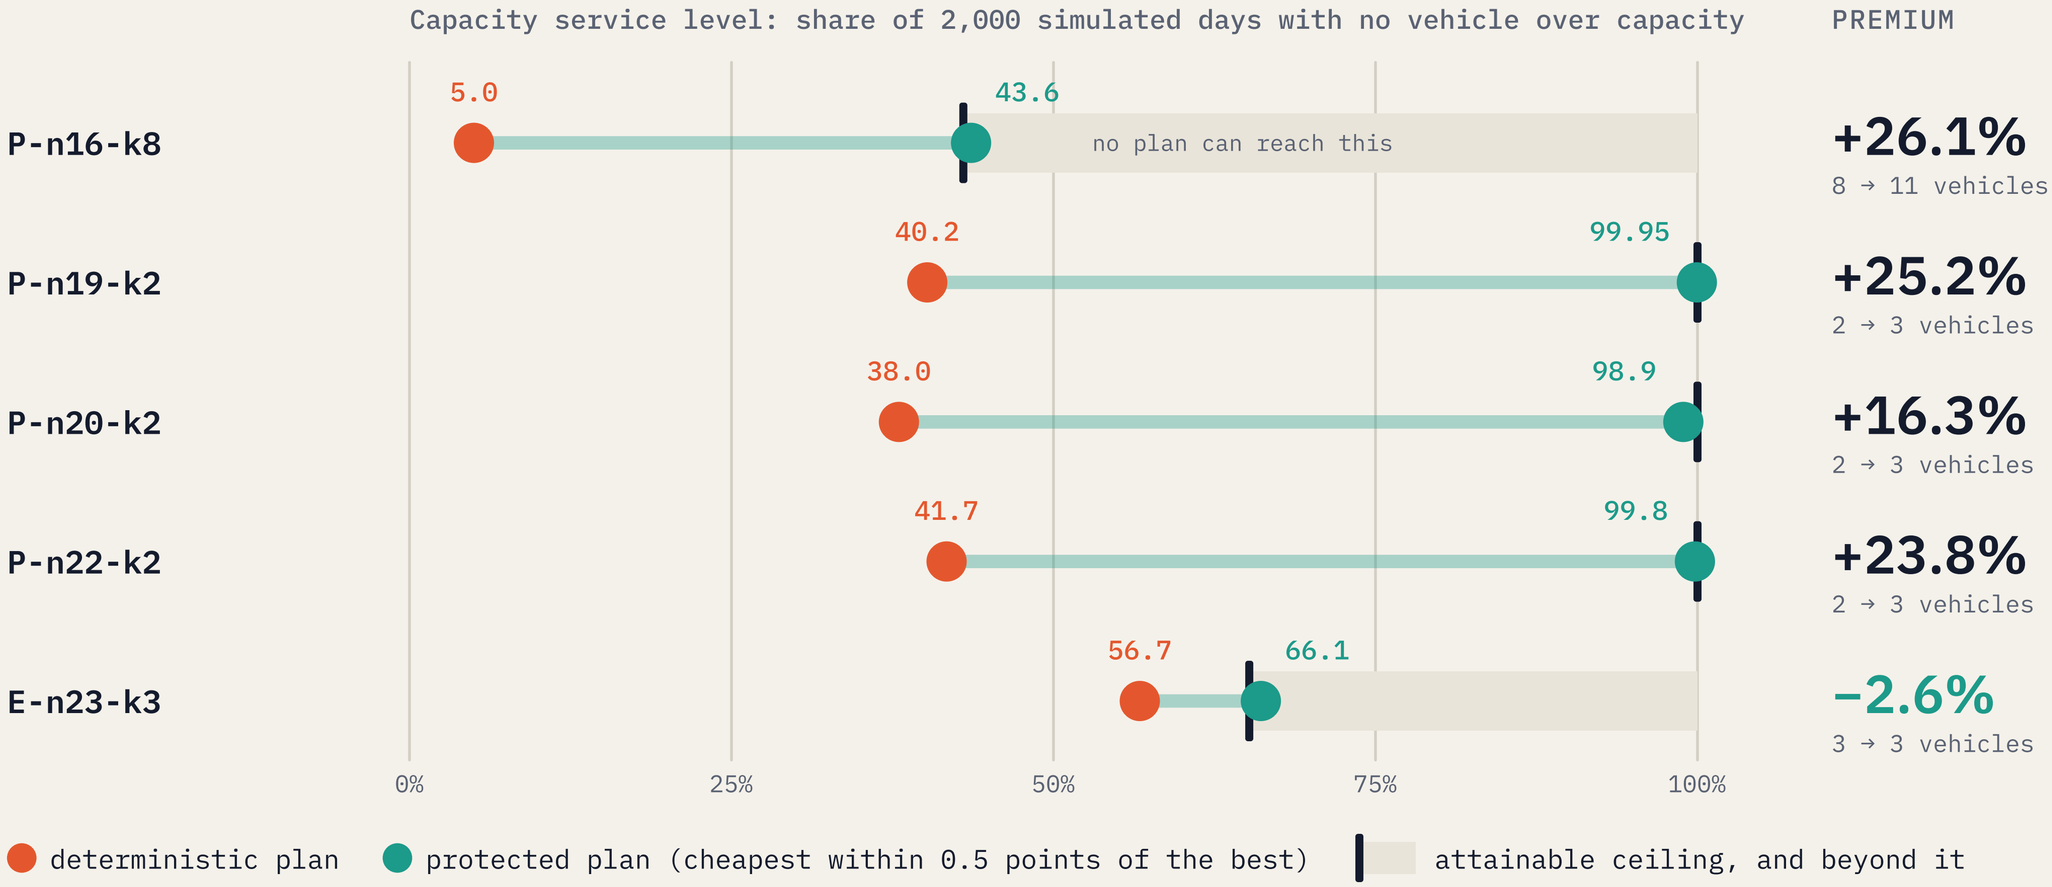
Γ_t

How many roads on a route may be congested at once?

SWITCHES RECOVER THE OTHER MODELS

Γ = 0, no recourse	→	DET
recourse only	→	SP
budgets only	→	RO
budgets + recourse	→	SR

Protection lifts reliability to what is attainable — for 16–26% more



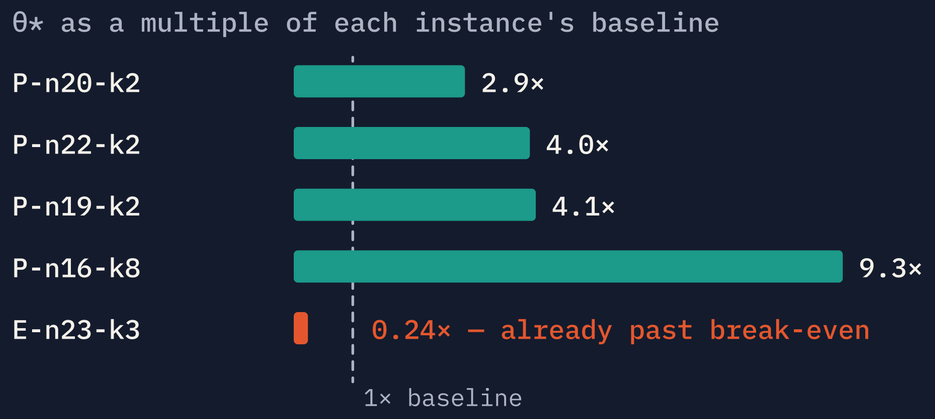
A single order larger than a van defeats every plan, which caps what is attainable (P-n16-k8: capacity 35, largest mean order 31). Protection reaches that ceiling on four of the five instances.

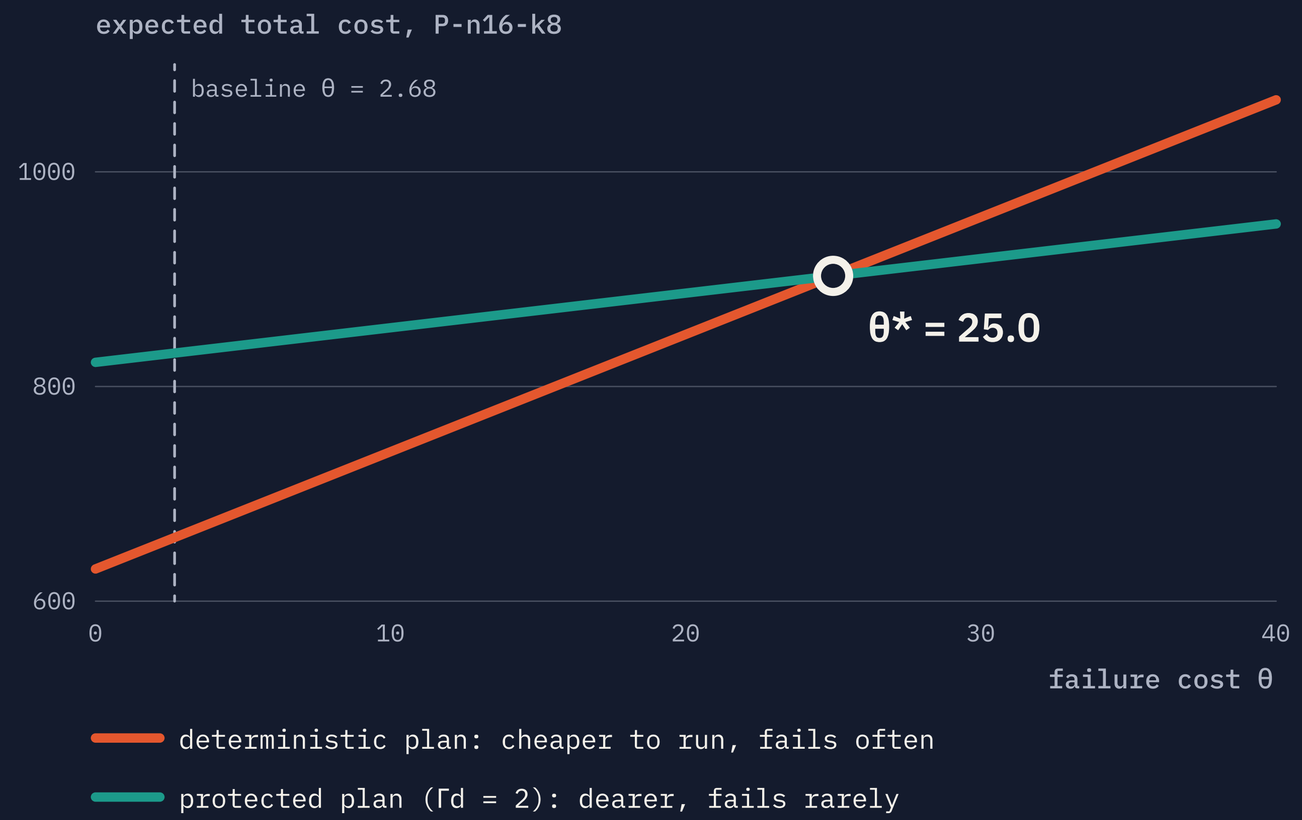
When does protection pay for itself?

θ* = 25

On P-n16-k8, a failed delivery must cost 9.3 times the baseline before the protected plan becomes the cheaper one.

Exact in closed form: with routes fixed, total cost is a straight line in θ. Unchanged across a 50-fold sweep of the failure cost.





Turn the two dials together

Buying capacity protection without duration protection makes the plan late more often than buying no protection at all.

COST ABOVE THE UNPROTECTED PLAN · (Γd, Γt)

capacity alone	(1, 0)	+139.0
duration alone	(0, 2)	+0.0
both together	(1, 2)	+185.5
interaction		**+46.5**

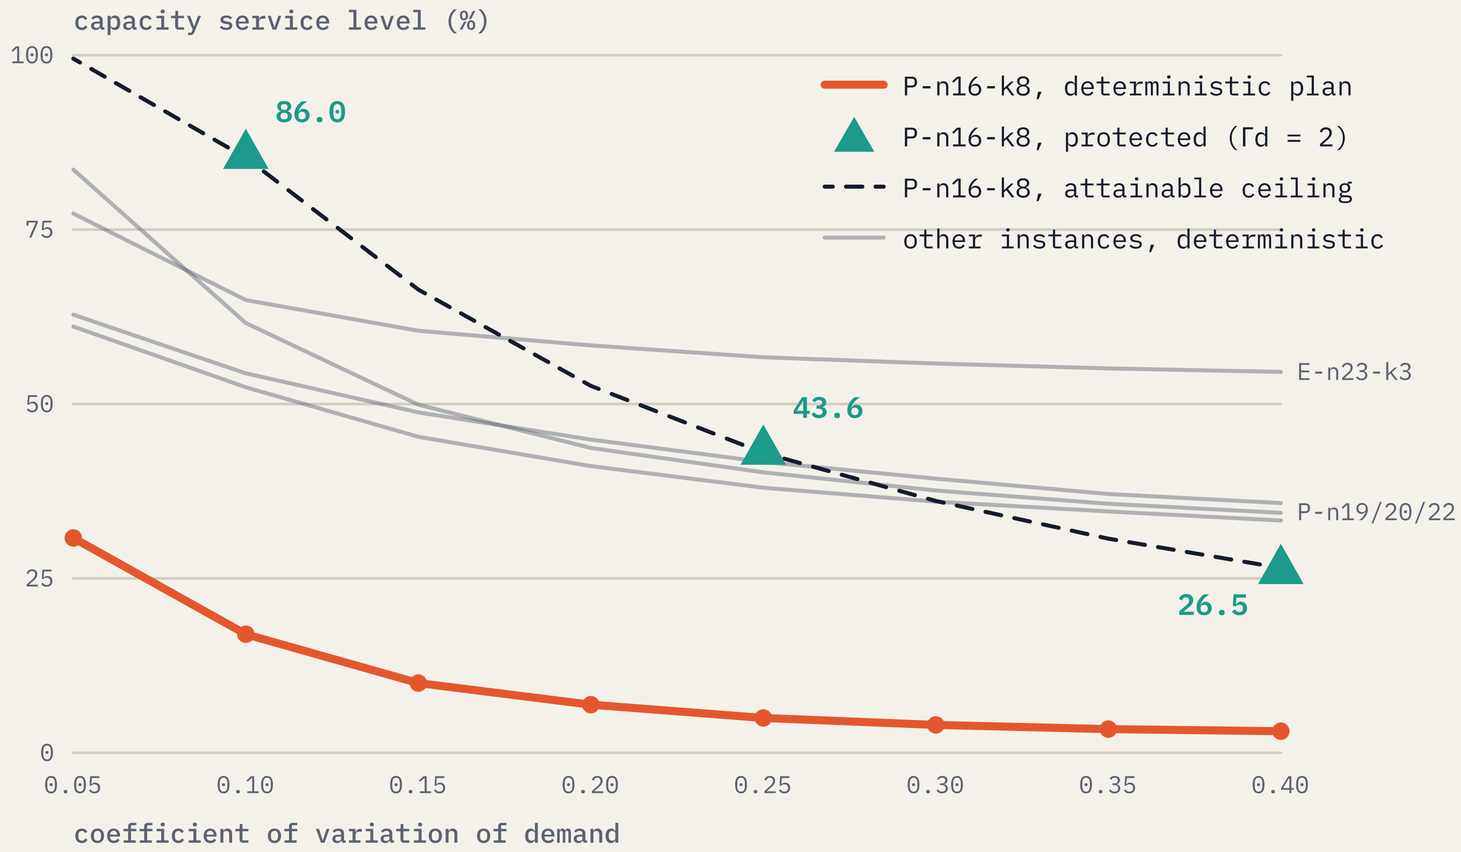
The budgets also saturate at route length: on P-n16-k8 nothing changes beyond Γd = 3, tested up to 15.

Time service level (%), integrated model, P-n16-k8

Γt ↓ Γd →	Γd = 0	Γd = 1	Γd = 2
Γt = 0	99.5 no protection	74.5 worse than none	99.3
Γt = 1	99.3	90.3	99.0
Γt = 2	99.0	99.0	99.3
Γt = 3	99.0	98.5	99.0

98.5 or more · 90 to 98.5 · below 90

More uncertainty: protection costs more, but pays off sooner



Fragility does not need high uncertainty: even at demand cv = 0.05, the deterministic plan fails on 16–69% of days.

PREMIUM · P-n16-k8, Γd = 2

+20.7 → +35.5%

rises from demand cv 0.10 to 0.40, one more vehicle each step

BREAK-EVEN FAILURE COST θ*

60.6 → 21.7

falls: the more uncertain demand is, the sooner protection pays

TRAVEL TIME HAS A CEILING TOO

At travel-time cv = 0.40, Γt = 2 is infeasible for any fleet: 7 of 15 customers cannot be served within T_max, even on a route of their own.

Key Performance Indicators (KPI)

COMMON EX-POST SCORECARD

Fixed routes

set by model

Same scenarios

independent holdout

Shared costs

routing + fleet + overflow

Evaluation KPIs

- Capacity, duration and joint service levels with 95% Wilson intervals
- Mean overflow and duration excess; fleet and peak-route utilisation
- Mean cost plus 95th-percentile VaR and upper-tail CVaR

Reproducibility exports

- Paired bootstrap intervals for model cost differences
- Scenario-level cost, feasibility, overflow and route-use exports
- Solver status, runtime, incumbent, bound and available gap metadata

Overflow is a penalised excess-load proxy. It does not simulate a physical depot return or a failed delivery.

The reporting pipeline runs across instances, but the reduced solves are inconclusive

40
model-instance rows

13
feasible incumbents

27
no usable solution

0
proven optimal

Coverage

Four instances × DET, SP, RO and SR

Run limits

20 training / 100 evaluation scenarios / 10 s per solve

Other checks

Demand and time CV sweeps at 0.10, 0.25, 0.40 on P-n16-k8

Second training seed 1201 on P-n16-k8 only

These checks verify expanded calculations and exports. They do not establish new benchmark-wide model-effectiveness claims.

The integrated model is useful, but its advantage depends on the regime

01 **The deterministic cost optimum is fragile**

43–95% overload frequency across the main benchmark runs.

02 **Protection has a price that depends on θ**

Break-even θ* is the clearer threshold; θ is a calibrated proxy.

03 **The budgets interact**

Capacity protection can make routes late unless duration is protected too.

WHAT THE EVIDENCE DOES NOT SHOW

- No field validation on a live delivery network
- Independent demand and arc-time assumptions
- No customer time windows
- Exact tier is limited to small instances ($n \leq 22$)
- Multi-seed inference across all instances remains open

Thank You

Daniel Ajayi
Stochastic–Robust Vehicle Routing