mathematics

MDPI

Article

When Does Robustness Pay? A Break-Even Analysis for Robust Truck-Capacity Planning in Disruption-Prone Cross-Border Corridors

Amir Karbassi Yazdi [1], Claudia Durán [2,*] and Iván Derpich [3,4,*]

[1] Departamento de Ingeniería Industrial y de Sistemas, Facultad de Ingeniería, Universidad de Tarapacá, Arica 1010069, Chile; akarbassiy@academicos.uta.cl
[2] Departamento de Ingeniería Industrial, Universidad Tecnológica Metropolitana, Santiago 7800002, Chile
[3] Departamento de Ingeniería Industrial, Universidad de Santiago de Chile, Santiago 9160000, Chile
[4] Centre for Innovation in Technology and Design of Materials for the Built Environment CIMAC_ANID, CTI250005, University of Santiago of Chile, Santiago 9170022, Chile
* Correspondence: c.durans@utem.cl (C.D.); ivan.derpich@usach.cl (I.D.); Tel.: +56-974511034 (C.D.); +56-989190614 (I.D.)

Abstract

Carriers operating disruption-prone cross-border corridors must commit capacity before demand is realised, while balancing shortage risk, idle capacity, and environmental costs. This study develops a budgeted robust mixed-integer linear programme for monthly truck-capacity planning with convex piecewise-linear recourse for shortage, idle-capacity, operating, and carbon costs. The robust counterpart is formulated as a compact MILP through linear-programming duality and benchmarked against deterministic and two-stage stochastic models; an ε-constraint extension evaluates the cost–emission trade-off, and a joint extension treats capacity loss as a second adversarial uncertainty source. For the continuous relaxation, we prove that the optimal plan switches from nominal to protected capacity at an exact break-even budget $\Gamma^* = \Sigma_t k_t/\kappa_t$, which depends only on cost and capacity parameters and not on the scale of demand uncertainty. The framework is demonstrated on a designed 144-month dataset that mirrors public statistics for the Cristo Redentor–Los Libertadores corridor; the numerical results are therefore a methodological demonstration rather than an empirical validation. The main capacity adjustment occurs between $\Gamma = 1$ and $\Gamma = 2$, as predicted by $\Gamma^* \approx 1.9$ months. At $\Gamma = 4$, the robust plan increases expected cost under normal operations by 1.8% but reduces scenario-stress expected cost and $CVaR_{95}$ by 6.9% and 7.8%, respectively, capturing 67% of the deterministic-to-stochastic stress-cost reduction with 8.6% fewer trips. This share is stable on independent out-of-sample stress sets but falls to 58% under a shifted, more severe disruption regime, where explicit capacity robustness closes much of the gap. Payload efficiency affects capacity and emissions far more than carbon pricing.

Keywords: robust optimization; budgeted uncertainty; freight capacity planning; cross-border transport; supply chain resilience; road freight decarbonization; stochastic programming

MSC: 65C05; 91B76; 90C46

Check for updates

Academic Editor: Xiang Song

Received: 18 October 2026
Revised: 16 November 2026
Accepted: 23 November 2026
Published: 27 November 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

Road freight plays a central role in overland trade, while reliability and carbon intensity have become important concerns for carriers, shippers, and regulators. Recent reviews

classify road-freight decarbonisation decisions across operational, tactical, and strategic planning levels, with comparatively less attention devoted to tactical decisions in which fleet capacity is committed months in advance [1]. At the same time, freight systems are increasingly exposed to disruptions that can interrupt flows and sharply increase transport costs, motivating a growing literature on resilient freight-network design [2,3]. Evidence from truckload markets further indicates that disruption shocks are absorbed primarily through price adjustments rather than volume adjustments [4]. Consequently, capacity that is not secured in advance may become particularly costly when it is most needed. This creates a fundamental planning problem: how much capacity should be committed before demand is realised when insufficient capacity generates shortage costs, whereas excessive capacity creates idle-capacity costs?

These pressures are particularly relevant in cross-border mountain corridors. The Cristo Redentor–Los Libertadores pass, a major road connection between Chile and Argentina, illustrates this problem. Freight demand is seasonal and uncertain, usable capacity is periodically reduced by weather-related closures, and unserved freight generates heterogeneous delay costs across product groups. Evidence of declining snow-cover extent in the extratropical Chilean Andes further suggests that disruption conditions may not remain stationary over time [5]. Carriers operating in such corridors must therefore determine monthly truck-capacity commitments before demand and disruption conditions are known. Insufficient capacity increases exposure to shortages during peak or disrupted periods, whereas excessive capacity raises fixed and idle-capacity costs and may increase the potential emissions associated with reserved fleet capacity. The resulting decision therefore requires an explicit treatment of uncertainty and of the economic trade-off between resilience and overcapacity.

The operations research literature provides two main approaches to this type of decision: robust optimisation and stochastic programming. Budgeted robust optimisation [6] protects against a limited number of adverse deviations without requiring full probability distributions and has been applied to freight-routing problems under demand uncertainty [7,8]. In these applications, however, the uncertainty budget is generally treated as an exogenous parameter whose effect is examined numerically. Analytical guidance for identifying the protection level at which additional robustness becomes economically justified therefore remains limited. Stochastic programming explicitly represents recourse and has been applied to tactical capacity contracting [9] and intermodal capacity planning [10], but it requires scenario probabilities that may be difficult to estimate for low-frequency disruptions. Comparative studies have consequently examined robust, distributionally robust, and stochastic formulations under alternative uncertainty representations [11–13]. Their relative performance, however, remains highly dependent on the type of uncertainty represented and on the underlying cost structure.

A related trade-off arises between resilience and decarbonisation. Freight-optimisation studies indicate that carbon pricing may have limited effects on transport decisions when service and handling costs dominate [14]. Empirical evidence from a container corridor similarly shows that even a carbon tax of 200 USD/tCO(_2) produces only moderate emission reductions through modal shift [15]. Improvements in vehicle utilisation, by contrast, can generate substantial reductions in both cost and emissions [16,17]. An additional issue is that protective capacity reserved for resilience is not necessarily operated. This distinction motivates separate consideration of latent emissions associated with contracted capacity and realised emissions generated by trips actually performed.

Based on these literature streams, three methodological gaps motivate the present study. First, in budgeted robust freight models, the uncertainty budget is typically treated as an exogenous parameter and evaluated through numerical sensitivity analysis, whereas an-

alytical threshold results remain largely confined to single-period allocation problems [18]. Consequently, an interpretable analytical characterisation of the critical robustness budget for multi-period capacity commitment remains limited. Second, few studies compare deterministic, budgeted robust, and stochastic capacity-planning approaches within a common evaluation framework that considers expected cost, service reliability, tail risk, and both demand and capacity uncertainty. This gap is particularly relevant because differences in the uncertainty represented by each formulation may materially affect their relative performance. Third, the existing literature generally does not distinguish between latent emissions associated with contracted protective capacity and realised emissions generated by trips actually operated. Cross-border mountain corridors provide a relevant setting in which to address these gaps because seasonal demand, disruption-induced capacity losses, heterogeneous shortage costs, and environmental considerations interact within the same tactical capacity decision.

Accordingly, the study addresses three research questions:

RQ1: Under what conditions does budgeted robust capacity planning become economically preferable to nominal planning, and can the corresponding critical uncertainty budget Γ^* be characterised analytically?

RQ2: How do deterministic, budgeted robust, and two-stage stochastic capacity-planning models compare in terms of expected cost, service reliability, and tail risk under demand and disruption uncertainty?

RQ3: How do key operational, economic, and environmental parameters—including shortage cost, payload, disruption severity, carbon pricing, and emission constraints—affect the optimal capacity plan and its cost–service–emission trade-offs?

To address these questions, a budgeted robust mixed-integer linear programme is formulated for monthly truck-capacity planning with convex piecewise-linear recourse representing shortage, idle-capacity, operating, and carbon costs. The resulting min–max problem is reformulated as a compact MILP through linear-programming duality. Particular attention is given to the analytical properties of the robust solution and to the critical robustness budget Γ^*, at which the optimal capacity decision shifts from nominal to protected planning. For the continuous relaxation, this threshold is derived in closed form and proved to be exact (Proposition 1). The robust formulation is benchmarked against deterministic and two-stage stochastic models based on nine demand–capacity–fuel scenarios and is further extended through an variance of epsilon constraint on contracted-fleet emissions and a joint demand–capacity robust formulation in which capacity loss constitutes a second adversarial uncertainty source (M1-DC).

The framework is demonstrated using a designed 144-month dataset representing the Cristo Redentor–Los Libertadores corridor. The dataset reproduces the structure of official freight, customs, fuel, and emission statistics and is mapped to the corresponding public data sources. Because the numerical series are designed rather than observed, the application is interpreted as a methodological demonstration rather than as an empirical validation of corridor-specific performance. The resulting capacity plans are evaluated through Monte Carlo simulation under normal operations, in-sample scenario stress, three independent out-of-sample stress sets, alternative forecast-error specifications, replays of reported closure patterns, and a seven-year rolling-origin evaluation based on the designed series.

The main contributions of this study are threefold.

First, the threshold behaviour of budgeted robust capacity planning is characterised and an exact analytical break-even budget is derived (RQ1). Under convex piecewise-linear recourse, it is proved that the continuous relaxation of the robust counterpart switches from nominal to protected capacity at an exact break-even budget $\Gamma^* = \Sigma_t \, k_t/\kappa_t$, which depends only on the fixed trip cost, usable trip capacity, idle-capacity cost, variable service cost, and

shortage penalty. The threshold admits a newsvendor interpretation and extends to capacity uncertainty. This result provides an interpretable criterion for identifying when protection becomes economically preferable to nominal planning, rather than relying exclusively on a numerical sweep of Γ.

Second, deterministic, budgeted robust, and two-stage stochastic capacity-planning models are evaluated within a common computational framework (RQ2). The comparison quantifies their differences in expected cost, service reliability, tail risk, and required capacity, and identifies the conditions under which demand-oriented robustness is sufficient and those under which explicit modelling of joint demand–capacity disruptions provides additional value. The comparison includes independent out-of-sample stress sets that are not generated from the scenarios used to optimise the stochastic model, and a joint demand–capacity robust extension that removes the asymmetry between the uncertainty represented in the robust and stochastic models.

Third, the sensitivity of the optimal capacity decision to operational, economic, and environmental parameters is systematically analysed (RQ3). In particular, the framework distinguishes between latent emissions associated with contracted protective capacity and realised emissions generated by trips actually operated, allowing the cost–service–emission trade-off to be evaluated without assuming that all reserved capacity is used.

The cross-border mountain-corridor application integrates seasonal demand, disruption-induced capacity losses, product-mix-weighted delay costs, and emissions, providing a computational test bed for the analytical and numerical properties of the framework.

The remainder of the paper is organised as follows. Section 2 reviews the related literature and positions the study. Section 3 presents the dataset, modelling assumptions, notation, uncertainty representation, mathematical formulations, and experimental design. Section 4 reports the computational results, including the robustness frontier, model comparisons, cost–emission trade-offs, sensitivity analyses, and rolling-origin evaluation. Section 5 discusses the theoretical, managerial, and policy implications of the findings, together with their generalisability and implementation pathway. Finally, Section 6 summarises the main conclusions, limitations, and directions for future research. Proofs are given in Appendix C.

2. Literature Review

2.1. Optimisation Under Uncertainty for Freight Capacity Planning

Robust optimization with budgeted (cardinality-constrained) uncertainty sets has become a standard device for protecting freight decisions against demand variability, yet its recent applications remain concentrated at the operational level [6]. Yu et al. [8] derive the robust counterpart of the vehicle routing problem with cross-docking under polyhedral demand uncertainty. Liu et al. [7] extend the approach to electric-vehicle pickup and delivery with time windows, formulating a two-stage adaptive robust model in which routes are fixed before demand is revealed and load quantities adjust afterwards. In both studies, and in the routing literature more generally, the uncertainty budget is treated as an exogenous parameter whose effect is traced empirically across a range of values; the structure of the optimal protection decision is not characterised analytically. Consequently, these applications provide limited analytical guidance for determining the protection level itself.

Tactical capacity commitment has more often been modelled stochastically. Crainic et al. [9] formulate a capacity-contracting problem with three sources of uncertainty: freight demand, spot-market conditions, and contracted-capacity failure. Gbadegoye et al. [10] develop a two-stage stochastic programme with a conditional value-at-risk criterion for road–rail intermodal transport under uncertain demand and train capacity, combining cost, emission, and reliability objectives. These models capture recourse explicitly but require

scenario probabilities, which may be difficult to specify for low-frequency disruptions, and provide limited comparison with distribution-free robust alternatives.

A third line of work compares modelling paradigms directly. Xia et al. [11] report improved out-of-sample performance for a data-driven distributionally robust model relative to stochastic and robust formulations. Jin et al. [13] reach a similar conclusion for humanitarian logistics with uncertain third-party supply capacity, where a scenario-wise distributionally robust model dominates stochastic programming. Hosseini Shekarabi et al. [12] benchmark an extended robust method against two-stage stochastic programming and the Bertsimas–Sim approach for a perishable-product network. Closest in spirit to the analytical results of the present paper, Wang et al. [18] derive closed-form solutions for a distributionally robust multi-product newsvendor with a capacity constraint and identify capacity thresholds below which products receive no allocation.

2.2. Decarbonisation of Road Freight and the Effectiveness of Carbon Policy

The decarbonisation of road freight has been surveyed comprehensively by Koç et al. [1], who classify planning models into operational, tactical, and strategic levels. Within the optimisation literature, carbon policy is most often evaluated through its effect on route and mode choice. Wu et al. [14] compare a carbon tax and an emissions trading scheme for oversize and heavyweight cargo and find that carbon price fluctuations have only a small effect on routing decisions. An empirical corridor study points in the same direction: for the Yiwu–Ningbo container corridor, Zhu et al. [15] estimate that a carbon tax of 200 USD per tonne of CO2 reduces corridor emissions by only 14.2% through modal shift. Taken together, these studies suggest that carbon prices at realistic levels exert limited leverage on freight decisions in which service and handling costs dominate.

By contrast, a growing body of evidence identifies vehicle utilisation as a potentially important source of emission reductions. Solakivi and Ojala [16] estimate that high-capacity trucks offer an emission-reduction potential of 42% in mass-based and 38% in volume-based transport relative to standard semitrailers. At a national scale, Xu et al. [17] combine proprietary truck-hailing data with a fleet model and show that improvements in load factor and reductions in empty running could cut China's road-freight emissions substantially, and that these operational gains are complementary to technological change. However, this literature generally evaluates emissions from planned or realised flows and does not explicitly distinguish emissions associated with reserved capacity from those generated by trips actually operated. This distinction is particularly relevant in tactical resilience planning, where protective capacity may be contracted without necessarily being dispatched.

2.3. Resilience to Disruption in Freight Corridors

Research on freight resilience increasingly integrates disruption risk into network design. Müllerklein and Fontaine [2] formulate a two-stage stochastic mixed-integer programme combining resilience strategies across strategic, tactical and operational levels under disruption uncertainty. Wei et al. [3] address vulnerability assessment and reduction for intermodal freight transportation networks. Both studies focus primarily on network-level resilience rather than on the amount of capacity that should be committed to a single disruption-prone corridor.

Recent empirical evidence also clarifies how freight markets absorb shocks. Using monthly state-level data for the United States, Hsu et al. [4] find that climate shocks raise truckload spot prices by 1.9% per standard deviation with minimal effect on volumes, suggesting that adjustment occurs mainly through prices rather than volumes. For mountain corridors, disruption conditions may also be non-stationary. Cordero et al. [5] document a decline in snow cover extent in the extratropical Chilean Andes of about 19% per decade

over 2001–2022. Such evidence reinforces the need to incorporate capacity risk explicitly into tactical planning.

2.4. Comparison of Related Studies

Table 1 positions the most closely related studies across six dimensions relevant to the proposed framework: tactical capacity, emissions, disruption, out-of-sample evaluation, cross-border application, and analytical threshold characterisation. The comparison shows that these dimensions have largely been studied separately. In particular, none of the studies summarised in Table 1 jointly addresses tactical capacity commitment, disruption risk, emissions, out-of-sample evaluation, and an analytical characterisation of the robustness threshold.

Table 1. Positioning of the proposed framework relative to related studies.

Ref.	Setting	Method and Uncertainty	Tact. Cap.	Emiss.	Disr.	OOS	X-Border	Threshold
[8]	Vehicle routing with cross-docking	Budgeted RO; demand	–	–	–	–	–	–
[7]	Electric-vehicle pickup and delivery	Two-stage adaptive RO; demand	–	–	–	–	–	–
[11]	Modular bus capacity	Data-driven DRO vs. SP/RO; demand	✓	–	–	✓	–	–
[18]	Multi-product newsvendor	Moment-based DRO; demand	✓	–	–	–	–	(✓)
[9]	Shipper capacity contracting	Stochastic bin packing; loads, spot market, capacity loss	✓	–	✓	–	–	–
[10]	Road–rail intermodal freight	Two-stage SP with CVaR; demand and capacity	✓	✓	(✓)	–	–	–
[12]	Perishable supply network	Extended RO vs. SP/BS; multiple uncertainties	✓	✓	✓	–	–	–
[13]	Humanitarian logistics	Scenario-wise DRO vs. SP; demand and supply capacity	✓	–	✓	✓	–	–
[15]	Container corridor modal shift	Emission accounting and discrete choice	–	✓	–	–	–	–
[14]	Oversize cargo routing	GA with carbon pricing	–	✓	–	–	–	–
[3]	Intermodal freight network	Vulnerability assessment and reduction; disruption	–	–	✓	–	–	–
[2]	Resilient transport network	Two-stage SP with Benders decomposition; disruption and cost	✓	–	✓	–	–	–
This study	Cross-border mountain corridor	Γ-robust MILP, SP, ε-constraint; budgeted demand and disruption scenarios	✓	✓	✓	✓ MC + rolling	✓	✓ Γ*

Notes: ✓ = addressed; (✓) = partially addressed; – = not addressed. Tact. cap. = tactical capacity; Emiss. = emissions; Disr. = disruption; RO = robust optimisation; DRO = distributionally robust optimisation; SP = stochastic programming; GA = genetic algorithm; BS = Bertsimas–Sim; CVaR = conditional value-at-risk; MILP = mixed-integer linear programme; OOS = out-of-sample; MC = Monte Carlo; Γ* = critical robustness threshold. For this study, we distinguish latent contracted-fleet emissions from realised operating emissions.

3. Methodology

3.1. Dataset Description

This study uses a designed monthly panel representing import road freight through the Cristo Redentor–Los Libertadores pass, a major cross-border road corridor between Argentina and Chile. The panel covers January 2014 to December 2025, yielding 144 observations; each observation represents one month for a fixed corridor and import direction.

The panel is organised in a research workbook with nine linked sheets (Table 2). The numerical series are designed to be economically plausible and reproducible. Provenance fields map each variable to the corresponding official repository, allowing subsequent replacement of the designed values with observed data. These include the Observatorio Logístico de Chile for freight flows and vehicles, the Servicio Nacional de Aduanas for customs declarations, the Comisión Nacional de Energía for diesel prices, HuellaChile for emission factors, and MOP for road distances.

Table 2. Structure of the research workbook.

Sheet	Content	Size	Role in the Methodology
Monthly_Panel	Freight, trucks, payload, diesel, disruption, product shares, derived cost, and CO2e	144 × 33	Forecasting, uncertainty calibration, rolling-origin evaluation
Model_Parameters	Cost, fuel, emission, capacity, and service parameters	15 parameters	Model coefficients and calibration
Demand_Scenarios	Demand, capacity, and diesel multipliers with probabilities	9 scenarios	Stochastic benchmark and stress evaluation
Gamma_Design	Γ levels and fixed experimental settings	8 levels	Robustness frontier
Product_Mix	Shares and payload, delay-cost, and emission factors by HS group	6 groups	Mix-weighted coefficients
Sensitivity_Design	Factors, levels, and hypotheses	10 tests	Sensitivity and ablation
Source_Registry	Official repositories and provenance fields	7 sources	Replication and empirical replacement
Data_Dictionary/README	Definitions, units, provenance, and replication instructions	2 sheets	Documentation

Note: Data_Dictionary and README are reported jointly as documentation sheets.

The public repositories are used as provenance and external plausibility references rather than as direct observations for fitting the designed demand series. Publicly reported corridor closures are additionally used in Section 4.10 for external comparison and replay experiments. Consequently, the demand seasonality and numerical calibration reported here should not be interpreted as empirically estimated corridor parameters; full calibration with observed customs and freight-flow data remains a priority for future research.

The designed panel reproduces the structure of the official statistics (trend, seasonality, closure months, fuel prices, and emission factors), but its values are not observations. Accordingly, all numerical results in this paper, including the rolling-origin evaluation in Section 4.8, constitute a methodological demonstration of the framework rather than an empirical validation for the corridor. Two properties limit the dependence of the conclusions on the designed values. First, up to integer rounding, the model is positively homogeneous of degree one in nominal demand: scaling all $\underline{d_t}$ by a common factor scales trips, costs, and emissions by the same factor and leaves fill rates, shortage probabilities, percentage

changes, and Γ* unchanged. Second, Γ* depends only on cost and capacity parameters (Proposition 1). Section 4.10 compares the disruption component of the panel with publicly reported closure records and shows that the designed panel understates the severity of recent closures.

In the designed panel, monthly freight averages 137,091 t, with a standard deviation of 13,961 t (Table 3). Annual volume grows from 1.51 Mt in 2014 to 1.77 Mt in 2025, corresponding to an average annual growth rate of 1.49%.

Table 3. Descriptive statistics of the designed monthly calibration panel (n = 144).

Variable	Unit	Mean	SD	Min	Max
Freight demand	t/month	137,091	13,961	107,356	172,113
Freight vehicles	trucks/month	5577	576	4436	7319
Average load	t/truck	24.99	0.93	22.24	27.07
Diesel price	USD/L	0.846	0.070	0.668	1.007
Disruption index	0–1	0.129	0.090	0.050	0.386
Closure days	days/month	0.60	0.85	0	3
Rolling 12-month demand CV	–	0.147	0.037	0.050	0.196
Utilisation of nominal capacity	–	0.946	0.038	0.839	1.058
Estimated CO2e	t/month	2215	229	1762	2907

Note: SD = standard deviation; CV = coefficient of variation.

- Demand exhibits marked seasonality: the monthly index peaks in March (1.108) and May (1.095) and bottoms out in July (0.922) and September (0.928).
- Border closures occur in 58 of the 144 months and last at most 3 days.
- Nominal-capacity utilisation averages 0.946 and exceeds unity in some months, motivating the explicit utilisation constraint introduced below.

Nine scenarios describe joint shocks to demand, usable capacity, and diesel price (Table 4). These scenarios are designed stress states rather than empirically estimated events; the associated probabilities are calibration weights used consistently in M2 and in the scenario-stress evaluation. They range from weak demand to a joint extreme in which demand rises 28% while capacity falls 32%. Their probability-weighted means are 1.067 for the demand multiplier and 0.948 for the capacity multiplier. The designed product mix (Table 5) is represented through three share-weighted coefficients that adjust payload, shortage cost, and emission intensity.

Table 4. Joint demand-disruption-fuel scenarios used in M2 and stress evaluation.

ID	Scenario	Prob.	Demand Mult.	Capacity Mult.	Diesel Mult.	Closure Days
S01	Low demand	0.10	0.85	1.00	0.95	0
S02	Normal	0.30	1.00	1.00	1.00	0
S03	High demand	0.20	1.12	1.00	1.03	0
S04	Peak demand	0.10	1.22	0.98	1.05	0
S05	Moderate closure	0.10	1.05	0.86	1.03	2
S06	Severe closure	0.06	1.08	0.70	1.05	5
S07	Fuel-price shock	0.06	1.03	0.98	1.25	1
S08	Joint extreme	0.04	1.28	0.68	1.30	6
S09	Recovery surge	0.04	1.32	0.92	1.08	1

Note: Scenario probabilities sum to 1.00; Prob. = probability; mult. = multiplier.

Table 5. Product groups and mix-adjustment factors.

Group	Share	Payload Factor	Delay-Cost Index	Emission Adjustment	Perishability
Industrial inputs	0.34	1.05	1.10	1.00	Low
Machinery & equipment	0.16	0.82	1.40	1.03	Low
Food & agricultural	0.22	0.92	1.60	1.02	High
Consumer goods	0.18	0.88	1.30	1.01	Medium
Chemicals & pharma	0.06	0.78	1.80	1.05	Medium
Other	0.04	1.00	1.00	1.00	Low
Weighted factor	1.00	φPF = 0.9358	φDCI = 1.3320	φEIA = 1.0140	–

Note: ϕ_{PF}, ϕ_{DCI}, and ϕ_{EIA} denote the share-weighted payload, delay-cost, and emission-intensity factors, respectively. Perishability is descriptive only and does not enter the optimisation models directly; its effect is captured through the delay-cost index, which scales the shortage penalty via φ_{DCI} in Equation (11).

3.2. Research Process

The methodological workflow comprises six sequential stages (Figure 1): data preparation, calibration, uncertainty representation, optimisation, simulation-based and rolling-origin evaluation, and synthesis of results.

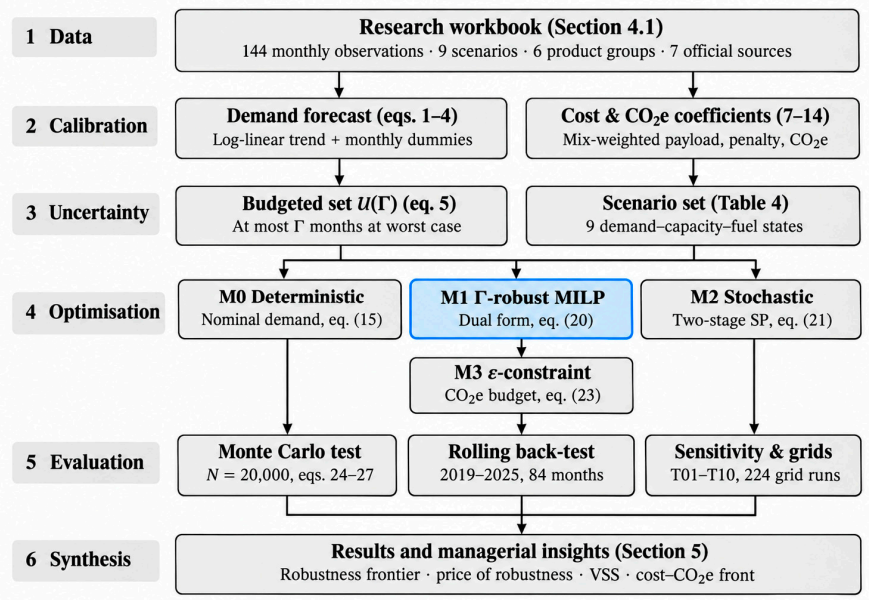


Figure 1. Research workflow. Arrows indicate the sequence of methodological stages; the shaded box highlights the Γ-robust MILP as the core methodological contribution.

Data (Section 3.1): the monthly panel, scenario set, and product mix define the calibration inputs.

Calibration (Sections 3.5 and 3.6): the calibration panel is transformed into demand forecasts and cost, capacity, and emission coefficients.

Uncertainty (Section 3.5): demand risk is represented by the budgeted uncertainty set $\mathcal{U}(\Gamma)$ used by the robust model and as the discrete scenarios used by the stochastic model.

Optimisation (Sections 3.7 and 3.8): M0 provides the deterministic benchmark, M1 the Γ-robust formulation, M2 the stochastic benchmark, and M3 the emission-constrained extension.

Evaluation (Section 3.9): each capacity plan is fixed and evaluated through Monte Carlo simulation, rolling-origin evaluation, one-factor sensitivity analysis, and two-factor interaction experiments.

Synthesis (Section 4): the results are summarised through the robustness frontier, price of robustness, VSS, and cost–emission frontier.

For each rolling-origin year, calibration, optimisation, and evaluation are repeated using only information available before the corresponding planning period.

3.3. Modelling Assumptions

The formulation relies on twelve assumptions concerning planning scope, demand and capacity, cost and emissions, and uncertainty representation. Sensitivity tests relax selected assumptions where indicated.

Scope and timing

A1—Planning horizon. A single corridor and import direction are considered over T = 12 monthly periods.

A2—Here-and-now capacity. Truck capacity is contracted before demand is observed, as an integer number of truck trips per month. Each contracted trip incurs the fixed dispatch cost c_{fix} whether or not it is used.

Capacity and demand

A3—Usable capacity. A trip carries at most $Q \cdot \varphi_{PF} \cdot u_{max}$ tonnes. Monthly capacity is reduced by the expected disruption loss $(1 - \ell \cdot DI_t)$, where DI_t is the historical mean disruption index for calendar month t. This assumption is relaxed in T07–T08.

A4—Uncertain demand. Demand is the only adversarial uncertain parameter in the robust model. The uncertainty set considers upward deviations from nominal demand $\bar{d}_t$ of at most $\hat{d}_t$. Lower realised demand is not included in the robust uncertainty set; its economic effect is captured ex post through the idle-capacity penalty. Capacity disruption enters the baseline formulation through its expected loss (A3), thereby isolating the effect of demand protection in the threshold analysis. The joint extension M1-DC (Section 3.8) relaxes this assumption by treating capacity loss as a second adversarial uncertainty source.

A5—Budgeted uncertainty. At most $\Gamma \in [0, T]$ months reach their worst-case demand simultaneously [6]. Its effect is examined in T01.

A6—No backlog. Unserved freight in month t is penalised and not carried forward to month t + 1.

Costs and emissions

A7—Linear trip costs. Variable trip cost is linear in distance: a non-fuel cost per km plus fuel at the monthly diesel price. A 10% empty-repositioning leg is included.

A8—Fuel-based emissions. Emissions are proportional to diesel burned (2.68 kg CO_2e/L, scaled by φ_{EIA}) and are priced at p_{CO2}. This assumption is varied in T04 and T09.

A9—Optimal recourse. Once demand is observed, the served quantity is $\min(dt, C_t)$, and only the trips required to serve this quantity are operated.

A10—Convex recourse. The effective shortage penalty exceeds the variable cost of serving one tonne ($b = 126.5 > a_t \approx 20$ USD/t), which makes the monthly recourse function convex and piecewise linear in demand and supports the exact robust reformulation for integer Γ developed in Section 3.7.

Uncertainty distributions

A11—Forecast errors. Log-demand errors are assumed to be normally distributed, independent across months, and homoscedastic; consequently, demand follows a lognormal specification. Robustness to this assumption is examined in Section 4.9 using autocorrelated, heteroscedastic, heavy-tailed, and structural-break error processes, together with alternative deviation calibrations.

A12—Scenario probabilities. Scenario probabilities are treated as given calibration weights in M2; in the stress evaluation, scenario states are independently sampled by month according to these weights. To avoid evaluating M2 exclusively under the same scenario

structure used for optimisation, Section 3.9 defines independent out-of-sample stress sets that do not rely directly on these calibration weights.

3.4. Notation

Tables 6–8 summarise the sets, parameters, and decision variables used throughout the formulations. A bar denotes a nominal value, except for $\underline{E}(x)$ which denotes contracted-fleet emissions; a hat denotes a maximum deviation. The superscript $k \in \{0, 1\}$ identifies the nominal (0) and worst-case (1) demand states.

Table 6. Sets and indices.

Symbol	Definition	Size
$\mathcal{T} = \{1, ..., T\}$, t	Planning months	T = 12
$\mathcal{S} = \{1, ..., S\}$, s	Demand–capacity–fuel scenarios	S = 9
$\mathcal{G}$, g	Product groups	6
$\mathcal{H} = \{1, ..., H\}$, h	Historical months used for estimation	H = 144
$\mathcal{K} = \{0, 1\}$, k	Demand state: nominal (0), worst case (1)	2
$\mathcal{U}(\Gamma)$	Budgeted demand-uncertainty set	–

Table 7. Model parameters and baseline values.

Symbol	Definition [Baseline]	Unit
	Demand and uncertainty	
$\bar{d}_t$	Nominal (forecast) demand in month t	t
σ_t	Standard deviation of demand in month t	t
$\hat{d}_t$	Maximum demand deviation, $\hat{d}_t = z_{SL}\ \delta\ \sigma_t$	t
z_{SL}	Standard normal quantile of the service target [1.96]	–
δ	Demand-deviation scale [1]	–
Γ	Uncertainty budget [4]	months
	Capacity and operations	
Δ	Corridor distance [380]	km
Q	Nominal truck payload [26]	t/truck
u_{max}	Maximum utilisation [0.90]	–
$\varphi_{PF}, \varphi_{DCI}, \varphi_{EIA}$	Mix-weighted payload, delay-cost, and emission factors [0.9358, 1.3320, 1.0140]	–
DI_t	Mean disruption index for calendar month t	–
ℓ	Capacity lost per unit of disruption index [0.22]	–
r	Minimum capacity reserve [0.05]	–
	Cost and emissions	
fc_L, fc_E	Fuel use on loaded/empty legs [0.36, 0.30]	L/km
r_E	Empty-repositioning share [0.10]	–
P_t	Diesel price in month t	USD/L
c_{nf}	Non-fuel operating cost [0.78]	USD/km
c_{fix}	Fixed cost per contracted trip [145]	USD/trip
p_{sh}	Shortage penalty before mix weighting [95]	USD/t
h	Idle-capacity penalty [12]	USD/t
e_f	Diesel emission factor [2.68]	kg CO2e/L
p_{CO_2e}	Carbon price [55]	USD/tCO2e
p_s	Probability of scenario s	–
DM_s, CM_s, PM_s	Demand, capacity, and diesel multipliers of scenario s	–
ε	Fleet CO2e budget in the ε-constraint model	tCO2e

Note: Values in brackets denote the baseline calibration. Derived parameters F, e, q, π_t, a_t, and b are defined in Section 3.6. Parameters of the joint demand–capacity model M1-DC are defined in Section 3.8.

Table 8. Decision variables.

Symbol	Domain	Definition	Model
x_t	$\mathbb{Z}+$	Contracted truck trips in month t (first-stage decision)	M0–M3
C_t	$\mathbb{R}+$	Usable capacity, Ct = πt xt (auxiliary)	M0–M3
g_t^k	$\mathbb{R}$	Epigraph of recourse cost in month t and demand state k	M0, M1, M3
λ	$\mathbb{R}+$	Dual variable of the uncertainty-budget constraint	M1, M3
μ_t	$\mathbb{R}+$	Dual variable of the upper bound $z_t \leq 1$ in Equation (18)	M1, M3
y_{ts}	$\mathbb{R}+$	Tonnes served in month t, scenario s	M2
u_{ts}	$\mathbb{R}+$	Tonnes unserved (shortage) in month t, scenario s	M2
o_{ts}	$\mathbb{R}+$	Idle capacity in month t, scenario s	M2
g_t^{jk}	$\mathbb{R}$	Epigraph of recourse cost in month t, demand state j, and capacity state k	M1-DC
$\lambda_d, \lambda_c, \nu_t$	$\mathbb{R}+$	Dual variables of the demand budget, capacity budget, and monthly state constraints	M1-DC

3.5. Demand Forecasting and the Uncertainty Set

Nominal demand and its uncertainty are estimated in three stages: demand forecasting, deviation calibration, and construction of the budgeted uncertainty set.

Step 1—Forecast model. Log demand is regressed on a linear trend and eleven monthly dummies (January is the reference month):

$$lnD_h = \beta_0 + \beta_1 h + \sum_{m=2}^{12} \gamma_m 1[month(h) = m] + e_h, \qquad e_h \sim N(0, s^2), \quad h \in H \tag{1}$$

In Equation (1), β_1 captures the temporal trend, whereas γ_m represents month-specific seasonal effects relative to January. The residual variance, adjusted for the 13 estimated coefficients, is used to quantify forecast uncertainty:

$$s^2 = \frac{1}{H - 13} \sum_{h \in H} \hat{e}_h^2 \tag{2}$$

Step 2—Nominal demand and maximum deviation. Because demand is assumed to be lognormal, the point forecast includes the variance correction $s^2/2$:

$$\overline{d}_t = exp\left(\hat{\beta}_0 + \hat{\beta}_1(H + t) + \hat{\gamma}_{m(t)} + \frac{s^2}{2}\right), \qquad \sigma_t = \overline{d}_t \sqrt{e^{s^2} - 1} \tag{3}$$

The expressions in Equation (3) follow from the moments of a lognormal random variable. If ln D ~ N(m, s^2), then E[D] = exp(m + s^2/2) and SD(D) = E[D]√(exp(s^2) − 1). In Equation (3), $\overline{d}_t$ denotes expected demand and σ_t its standard deviation. The maximum demand deviation is then defined as:

$$\hat{d}_t = z_{SL}\, \delta\, \sigma_t \tag{4}$$

s = 0.0707, z_{SL} = 1.96, δ = 1.

Step 3—Budgeted uncertainty set. Following Bertsimas and Sim [6], the uncertainty set allows deviations in all months while limiting their aggregate magnitude through Γ:

$$U(\Gamma) = \left\{ d_t \in R^T : d_t = \overline{d}_t + \hat{d}_t z_t, \;\; 0 \leq z_t \leq 1, \;\; \sum_{t \in T} z_t \leq \Gamma \right\} \tag{5}$$

Γ = 0 recovers the deterministic case, whereas Γ = T = 12 yields full protection, as in Soyster [19]. Intermediate values control the degree of conservatism.

Calendar-specific inputs are obtained from the corresponding months of the designed calibration panel:

$$DI_t = \frac{1}{|H_t|} \sum_{h \in H_t} DI_h, \qquad P_t = P_{H-12+t} \tag{6}$$

In Equation (6), $\mathcal{H}_t$ denotes the historical observations associated with calendar month t. Diesel prices are generated using a seasonal naive specification based on the corresponding month of the final calibration year.

3.6. Capacity, Cost and Emission Structure

The operational coefficients entering the optimisation models are derived from the parameters in Table 7.

Fuel and emissions per trip. A trip consumes fuel on the loaded leg and on a fraction of the empty-repositioning leg:

$$F = \Delta\,(fc_L + r_E fc_E) = 380 \times (0.36 + 0.10 \times 0.30) = 148.2\ L/trip \tag{7}$$

Emissions per trip are obtained from fuel consumption as follows:

$$e = \frac{F \cdot ef \cdot \varphi_{EIA}}{1000} = 0.403\ tCO_2e/trip \tag{8}$$

Capacity per trip. Effective load q combines nominal payload, product-mix adjustment, and maximum utilisation. Usable capacity π_t additionally accounts for disruption-related capacity loss:

$$q = Q\,\varphi_{PF}\,u_{max} = 21.90\ t/trip, \qquad \pi_t = q\,(1 - l\ DI_t), \qquad C_t = \pi_t\,x_t \tag{9}$$

Cost per tonne served. The variable trip cost combines distance-related operating cost and fuel expenditure. The carbon charge is then included in the per-tonne service cost:

$$v_t = \Delta\,c_{nf} + F\,P_t, \qquad a_t = \frac{v_t + p_{CO_2}\,e}{q} \approx 20\ USD/t \tag{10}$$

Shortage penalty. The shortage penalty is adjusted by the product-mix delay-cost index:

$$b = p_{sh}\,\varphi_{DCI} = 95 \times 1.332 = 126.5\ USD/t \tag{11}$$

Monthly recourse cost. For realised demand d and capacity C, the monthly recourse cost is defined as

$$f_t(d, C) = a_t min(d, C) + b\,(d - C)^+ + h\,(C - d)^+ \tag{12}$$

where the three terms represent, respectively, operating, shortage, and idle-capacity costs. Since $a_t < b + h$, which holds for the baseline calibration ($20 < 138.5$), Equation (12) can be equivalently written as the maximum of two affine functions:

$$f_t(d, C) = max\{a_t d + h(C - d),\ \ a_t C + b(d - C)\} \tag{13}$$

Equation (13) shows that $f_t(d,C)$ is convex and piecewise linear.

Annual cost and emissions. For a capacity plan x and demand vector d, the annual cost, realised emissions, and contracted-fleet emissions are defined as

$$Z(x,d) = \sum\nolimits_{t\in T}\left[c_{fix}\,x_t + f_t(d_t, \pi_t x_t)\right], \qquad E(x,d) = \sum\nolimits_{t\in T}\frac{e\,min(d_t, \pi_t x_t)}{q}, \quad \overline{E}(x) = e\sum\nolimits_{t\in T} x_t \tag{14}$$

Here, E(x,d) measures emissions from trips actually operated, whereas $\overline{E}(x)$ represents the emissions associated with fully utilised contracted capacity. The latter is the quantity constrained in the ε-constraint formulation.

3.7. Deterministic and Γ-Robust Capacity Models

Both formulations determine the integer number of contracted trips x_t; they differ in the treatment of demand uncertainty.

Model M0—deterministic benchmark. M0 is formulated using nominal demand $\overline{d}_t$. The recourse function in Equation (13) is represented through the epigraph variable g_t^0, which satisfies $g_t^0 = f_t\,(\overline{d}_t, C_t)$ at optimality:

$$\min_{x, g^0} \sum_{t\in T}\left(c_{fix}\,x_t + g_t^0\right) \tag{15a}$$

$$s.t. \quad g_t^0 \geq a_t\,\bar{d}_t + h\left(\pi_t x_t - \bar{d}_t\right) \qquad \forall t \in T \tag{15b}$$

$$g_t^0 \geq a_t \pi_t x_t + b\left(\bar{d}_t - \pi_t x_t\right) \qquad \forall t \in T \tag{15c}$$

$$\pi_t x_t \geq (1+r)\,d_t, \quad x_t \in Z_+ \qquad \forall t \in T \tag{15d}$$

Constraints (15a–d) define the deterministic formulation. Equation (15a) minimises fixed contracting and monthly recourse costs; Equation (15b,c) represent the excess-capacity and shortage regimes, respectively; and Equation (15d) enforces the minimum 5% capacity reserve and integrality of x_t.

Model M1—Γ-robust model. The robust formulation minimises the worst-case annual cost over the budgeted uncertainty set $U(\Gamma)$:

$$\min_{x \in Z_+^T} \sum_{t \in T} c_{fix}\,x_t + \max_{d_t \in U(\Gamma)} \sum_{t \in T} f_t(d_t, \pi_t x_t) \quad s.t.\ (15d) \tag{16}$$

Problem (16) is a min–max structure. To obtain a tractable MILP formulation, the inner maximisation is reformulated through epigraph representation and linear-programming duality.

Step 1—Extreme demand states. Let g_t^0 and g_t^1 denote the recourse costs associated with nominal demand ($z_t = 0$) and worst-case demand ($z_t = 1$), respectively. Define

$$D_{t^0} = \bar{d}_t,\ D_{t^1} = \bar{d}_t + \hat{d}_t.$$

Because $f_t(d,C)$ is convex in d,

$$f_t\left(\bar{d}_t + \hat{d}_t z_t, C_t\right) \leq (1 - z_t)\,g_t^0 + z_t\,g_t^1 = g_t^0 + \left(g_t^1 - g_t^0\right) z_t, \quad z_t \in [0,1] \tag{17}$$

For $z_t \in \{0, 1\}$, Equation (17) holds with equality.

Step 2—Adversarial problem. Substituting Equation (17) into the inner maximisation of Equation (16) yields

$$\max_{z} \sum_{t \in T} \left(g_t^1 - g_t^0\right) z_t \quad s.t. \quad \sum_{t \in T} z_t \leq \Gamma : \textit{dual variable } [\lambda], z_t \leq 1 \ : \ \textit{dual variable } [\mu_t], z_t \geq 0 \tag{18}$$

Step 3—Dual reformulation. Since Problem (18) is feasible and bounded, strong duality applies. Introducing $\lambda \geq 0$ for the uncertainty-budget constraint and $\mu_t \geq 0$ for the upper bound on z_t, the dual problem is

$$\min_{\lambda,\mu} \Gamma\lambda + \sum_{t \in T} \mu_t \quad s.t. \quad \lambda + \mu_t \geq g_t^1 - g_t^0, \quad \lambda \geq 0, \quad \mu_t \geq 0 \qquad \forall t \tag{19}$$

Here, λ represents the marginal cost of increasing the uncertainty budget, whereas μ_t captures the residual worst-case contribution associated with month t. Combining the outer minimisation in Equation (16) with the dual problem in Equation (19) gives the compact robust counterpart.

Robust counterpart (M1).

$$min \sum_{t \in T} c_{fix}\,x_t + \sum_{t \in T} g_t^0 + \Gamma\lambda + \sum_{t \in T} \mu_t \tag{20a}$$

$$g_t^k \geq a_t D_t^k + h\left(\pi_t x_t - D_t^k\right) \qquad \forall t \in T,\ k \in \{0,1\} \tag{20b}$$

$$g_t^k \geq a_t \pi_t x_t + b\left(D_t^k - \pi_t x_t\right) \qquad \forall t \in T,\ k \in \{0,1\} \tag{20c}$$

$$\mu_t \geq g_t^1 - g_t^0 - \lambda,\ \mu_t \geq 0, \qquad \lambda \geq 0, \quad g_t^k \in R, \quad \forall t \in T \tag{20d}$$

$$\pi_t x_t \geq (1+r)\,\bar{d}_t, \qquad x_t \in Z_+, \quad \forall t \in T \tag{20e}$$

Equation (20b,c) represent the two affine components of the recourse function at nominal and worst-case demand. Equation (20d) incorporates the dual representation of the budgeted uncertainty set, while Equation (20e) retains the reserve and domain constraints.

Properties. For integer Γ, the vertices of $U(\Gamma)$ satisfy $z_t \in \{0,1\}$; consequently, Equation (17) is exact at the relevant extreme points and the robust counterpart in Equation (20) recovers the exact worst-case value of Equation (16). For fractional Γ, Equation (17) provides an upper-bounding interpolation; exactness is not guaranteed because an extreme point may contain one fractional deviation component. When $\Gamma = 0$, M1 reduces to the deterministic model M0, whereas $\Gamma = T = 12$ corresponds to full Soyster-type protection. For $T = 12$, the formulation contains 12 integer variables, 37 continuous variables, and 72 constraints and is solved using branch-and-bound with a relative optimality gap of 10^{-6}.

Break-even robustness budget. The threshold behaviour reported in Section 4.2 can be characterised exactly for the continuous relaxation of M1, where $x_t \in \mathbb{R}_+$. With approximately 8000 trips per month, integrality changes each monthly decision by less than one trip. Let $C_t = \pi_t\, x_t$, $L_t = (1 + r)\, \bar{d}_t$ and $U_t = \bar{d}_t + \hat{d}_t$, and define for each month

$$k_t = \frac{c_{fix}}{\pi_t} + h, \qquad \kappa_t = b + h - a_t, \qquad \Gamma^* = \sum_{t \in T} \frac{k_t}{\kappa_t} \tag{21}$$

Raising capacity by one tonne in every month costs k_t per month in the nominal state, namely the fixed trip cost per tonne plus the idle-capacity cost. In each of the Γ months that reach worst-case demand, the same tonne avoids the shortage penalty b net of the variable service cost a_t and does not incur the idle cost h. For identical months, the worst-case objective therefore changes by $Tk - \Gamma\kappa$, and protection Γ becomes economically preferable when $T > \Gamma^* = Tk/\kappa$. The following result makes this argument exact and extends it to heterogeneous months.

Proposition 1 (Exact break-even budget). *Consider the continuous relaxation of M1 and suppose that (i) $0 < k_t < \kappa_t$ for all t; (ii) $r\underline{d_t} < \hat{d}_t$ for all t, so that the reserve does not already cover worst-case demand; and (iii) the plateau condition $\lambda_{lo} = max_t\ (a_t - h)\hat{d}_t < \lambda_{hi} = min_t\ (b\hat{d}_t - \kappa_t\, r\underline{d_t})$ holds. Let $C_t(\lambda) = min\{U_t, max\{L_t, (A_t - \lambda)/\kappa_t\}\}$ with $A_t = bU_t - (a_t - h)\underline{d_t}$. Then: (a) if $\Gamma < \Gamma^*$, every optimal plan satisfies $C_t \le C_t(\lambda_{hi})$ for all t (low-protection regime); (b) if $\Gamma > \Gamma^*$, every optimal plan satisfies $C_t \ge C_t(\lambda_{lo})$ for all t (high-protection regime); (c) if all months are identical, then $C_t(\lambda_{hi}) = L_t$ and $C_t(\lambda_{lo}) = U_t$, so the optimal plan is the nominal-reserve plan for $\Gamma < \Gamma^*$ and the fully protected plan for $\Gamma > \Gamma^*$, and every common capacity in $[L, U]$ is optimal at $\Gamma = \Gamma^*$.*

The proof in Appendix C dualises the adversarial problem, solves the monthly subproblems in closed form for a fixed dual price λ, and shows that the derivative of the resulting one-dimensional convex problem equals $\Gamma - \Gamma^*$ throughout the interval $(\lambda_{lo}, \lambda_{hi})$. Three implications follow. First, Γ^* is exact rather than approximate: it is the budget at which the continuous relaxation is indifferent between low and high protection. Second, Γ^* does not depend on $\underline{d_t}$, $\hat{d}_t$ δ, z_{SL} or r, provided that conditions (ii) and (iii) hold; the scale of demand uncertainty and the reserve determine the size of the protection layer, $\Sigma_t\ (U_t - L_t)$, but not the location of the threshold. At baseline, condition (iii) holds whenever $\delta \ge 0.37$ and $r \le 0.13$. Third, Γ^*/T is the average critical shortage probability k_t/κ_t of a risk-neutral newsvendor with capacity cost k_t and shortage cost $b - a_t + h$ (Appendix C).

Hence, Γ^* is the expected number of shortage months per year that a risk-neutral planner would accept, and robust protection becomes worthwhile when the planner insures against more adverse months than this number. With heterogeneous months, monthly capacities move only within the bounds $C_t(\lambda_{hi})$ and $C_t(\lambda_{lo})$, which explains the small changes observed for $\Gamma < \Gamma^*$ and $\Gamma > \Gamma^*$ in Table 10. A simplified first-order version, which neglects $a_t - h$ relative to b and the expected disruption loss in π_t, is discussed in Section 5.1 (Equation (31)).

3.8. Stochastic Benchmark and ε-Constraint Model

Model M2—two-stage stochastic programme. Model M2 provides the stochastic benchmark. In contrast to the worst-case protection adopted in M1, M2 minimises expected annual cost over the nine scenarios reported in Table 4. The first-stage decision determines the number of contracted trips x_t before uncertainty is realised. In the second stage, for each scenario s, the model determines the served freight y_{ts}, shortage u_{ts}, and idle capacity o_{ts}.

$$min \sum_{t \in T} c_{fix}\, x_t + \sum_{s \in S} p_s \sum_{t \in T} (a_{ts}\, y_{ts} + b\, u_{ts} + h\, o_{ts}) \tag{22a}$$

Subject to

$$y_{ts} + u_{ts} = DM_s\, \underline{d_t} \quad \forall t, s \tag{22b}$$

$$y_{ts} + o_{ts} = CM_s\, \pi_t\, x_t \quad \forall t, s \tag{22c}$$

$$\pi_t x_t \geq (1 + r)\ \underline{d_t}, \quad x_t \in Z_+, \quad y_{ts}, u_{ts}, o_{ts} \geq 0, \qquad a_{ts} = \frac{\Delta c_{nf} + F\, PM_s\, P_t + p_{CO_2} e}{q} \tag{22d}$$

Equation (22b,c) enforce, respectively, scenario-specific demand and capacity balances, whereas Equation (22d) imposes the reserve, integrality, and non-negativity conditions. The scenario multiplier PM_s affects only the fuel-cost component of a_{ts}. The deterministic equivalent contains 12 integer variables and 324 continuous variables.

Value of the stochastic solution (VSS). The value of the stochastic solution measures the additional expected cost incurred when the stochastic programme is replaced by an expected-value solution. Let RP denote the optimal value of Model M2 and let x^{EV} be the solution of the corresponding expected-value model obtained using the scenario-weighted mean demand, capacity, and fuel-price multipliers. The corresponding expected result EEV is obtained by evaluating x^{EV} in the stochastic model. Thus,

$$VSS = EEV - RP = \sum\nolimits_{s \in S} p_s\, Z\left(x^{EV}, \xi_s\right) - RP\ \geq 0 \tag{23}$$

where ξ_s denotes the joint demand, capacity, and fuel-price realisation associated with scenario s. Because RP and EEV are both computed on the nine scenarios used to construct M2, this VSS is an in-sample quantity. Section 4.4 therefore also reports an out-of-sample counterpart, the difference between the expected costs of x^{EV} and of the M2 solution on stress sets that are not generated from Table 4.

Model M3—ε-constraint for cost versus emissions. To characterise the trade-off between worst-case cost and emissions, Model M3 extends the robust formulation by imposing an upper bound on contracted-fleet emissions $\overline{E}(x)$, as defined in Equation (14):

$$M3(\varepsilon): \quad min\ (20a) \quad s.t.\ (20b) - (20e), \qquad e \sum_{t \in T} x_t \leq \varepsilon \tag{24}$$

The emission budget ε is defined as a fraction of the baseline contracted-fleet emissions $\overline{E}(x^*)$ obtained for $\Gamma = 4$, and is varied from 100% to 90% in 1% increments. Solving Model M3 for each value of ε generates the efficient frontier between worst-case cost and contracted-fleet emissions. The ε-constraint approach is preferred to a weighted-sum formulation because it preserves the emission limit in physical units and can recover non-supported efficient solutions. Infeasibility for sufficiently small values of ε identifies the point at which the minimum reserve requirement in Equation (20e) can no longer be satisfied.

Model M1-DC—joint demand–capacity robust model. In M1, capacity disruption enters only through the expected loss factor $(1 - \ell\, DI_t)$, whereas M2 represents joint demand–capacity shocks explicitly. To remove this asymmetry, M1-DC treats capacity loss as a second adversarial uncertainty source. Realised capacity is $\gamma_t\ \pi_t\ x_t$, and demand and the capacity factor vary in the joint budgeted set

$$U(\Gamma_d, \Gamma_c) = \left\{(d, \gamma) : d_t = d_t + \hat{d}_t z_t,\ \gamma_t = 1 - \hat{\rho}\, w_t,\ z, w \in [0,1]^T,\ \sum_t z_t \le \Gamma_d,\ \sum_t w_t \le \Gamma_c\right\} \tag{25a}$$

where $\hat{\rho}$ is the maximum relative capacity loss in an adverse month and Γ_c is the number of months in which it can occur. Because f_t is jointly convex in (d, C) (Equation (13)) and C is affine in w_t, the worst case is attained at binary (z, w). Each month, therefore, has four extreme states (j, k) $\in \{0, 1\}^2$, with demand D_t^j and capacity $C_t^k = (1 - \mathrm{k}\,\hat{\rho})\pi_t$ x_t. Dualising the linear relaxation of the adversarial problem over these states yields the compact counterpart

$$min \sum_t c_{fix} x_t + \sum_t g_t^{00} + \Gamma_d \lambda_d + \Gamma_c \lambda_c + \sum_t \nu_t \tag{25b}$$

$$g_t^{jk} \ge a_t D_t^j + h\left(C_t^k - D_t^j\right), \quad g_t^{jk} \ge a_t C_t^k + b\left(D_t^j - C_t^k\right), \quad C_t^k = (1 - k\hat{\rho})\,\pi_t x_t \quad \forall t,\ j,k \in \{0,1\} \tag{25c}$$

$$\nu_t + \lambda_d \ge g_t^{10} - g_t^{00}, \quad \nu_t + \lambda_c \ge g_t^{01} - g_t^{00}, \quad \nu_t + \lambda_d + \lambda_c \ge g_t^{11} - g_t^{00}, \quad \nu_t, \lambda_d, \lambda_c \ge 0, \quad (20e) \tag{25d}$$

Formulation (25) is an upper bound on the exact worst case. In all instances reported in Section 4.4, the exact worst case, computed by dynamic programming over binary (z, w), coincides with the optimal value of (25), so the bound is tight. M1-DC reduces to M1 when $\Gamma_c = 0$. Following the logic of Proposition 1, protection against capacity loss has its own break-even budget.

Proposition 2 (Capacity-loss threshold). *In M1-DC with identical months and $\Gamma_c \le \Gamma_d \le T$, consider capacities between D^1 and $D^1/(1-\hat{\rho})$. On this range, the worst-case objective is affine in C with slope $Tk - \Gamma_c[(1-\hat{\rho})(b-a) + h]$. Additional protection against capacity loss is therefore optimal if and only if $\Gamma_c > \Gamma_c^*$, where*

$$\Gamma_c^* = \sum_{t \in T} \frac{k_t}{(1-\hat{\rho})(b - a_t) + h} \tag{26}$$

The proof is given in Appendix C. Baseline values are $\hat{\rho} = 0.14$, corresponding to the moderate-closure scenario S05, and $\hat{\rho} = 0.30$, corresponding to the severe-closure scenario S06, with $\Gamma_d = 4$ and $\Gamma_c \in \{0, 1, 2, 3\}$.

3.9. Evaluation and Experimental Design

Each optimised capacity plan x is held fixed during evaluation. Monte Carlo replications introduce new forecast-error realisations, while the stress regime additionally samples scenario multipliers from the nine-state scenario set defined in Table 4. The recourse function in Equation (12) is applied ex post, with served freight given by $min(d_t, C_t)$.

Simulation regimes. For replication n = 1, ..., N, with N = 20,000 and random seed 2026, month t is subject to a lognormal forecast error $\omega_t^{(n)} \sim \mathcal{N}(0, 1)$. Under the stress regime, a scenario state $s_t^{(n)}$ is additionally sampled according to the calibration probabilities p_s.

$$d_t^{(n)} = \bar{d}_t\, exp\left(s\,\omega_t^{(n)} - \frac{s^2}{2}\right) \cdot DM_{s_t^{(n)}}, \qquad C_t^{(n)} = CM_{s_t^{(n)}}\,\pi_t x_t, \qquad P_t^{(n)} = PM_{s_t^{(n)}}\,P_t \tag{27}$$

Two evaluation regimes are considered. Under normal operations, all scenario multipliers are set to unity and only forecast error is retained. Under scenario stress, Equation (27) incorporates demand, capacity, and diesel-price shocks in addition to forecast uncertainty.

Independent out-of-sample stress sets. Because the scenario-stress regime samples the same nine states and probabilities that are used to construct M2, it gives the stochastic plan an informational advantage. Three additional stress sets, each with N = 20,000 replications

and a separate random seed, are therefore generated without using the M2 scenario tree directly. OOS-1 (perturbed scenarios) rescales the deviation of every multiplier from unity by an independent U(0.5, 1.5) factor and redraws the scenario probabilities from a Dirichlet distribution with concentration parameters $20p_s$ in each replication. OOS-2 (continuous distribution) replaces the discrete tree with continuous laws matched only to the first two moments of the scenario mixture: a lognormal demand shock (mean 1.067, standard deviation 0.116), a monthly closure indicator with probability 0.24, a beta-distributed capacity loss with conditional mean 0.20 and standard deviation 0.09, a lognormal diesel shock (mean 1.042, standard deviation 0.082), and a Gaussian copula with correlation 0.5 between the demand shock and closure severity. OOS-3 (shifted regime) is a hold-out law that deliberately differs from the calibration: the closure probability doubles to 0.48, mean closure severity rises by 25% to 0.25, and closures persist from one month to the next (Markov persistence 0.5).

Alternative forecast-error specifications. To examine Assumption A11, both regimes are re-simulated with (i) AR(1) log-errors with coefficient 0.5 and unchanged marginal variance; (ii) heteroscedastic errors whose standard deviation is 30% higher in June–September than in the other months, with the same average variance; (iii) Student-t errors with four degrees of freedom scaled to unit variance; (iv) a structural break that, with probability 0.25 per year, shifts log demand permanently by $N(0.03, 0.06^2)$ from a random month onwards; and (v) all four departures combined. The robust plan is also re-optimised with alternative deviation calibrations: $z_{SL} = 1.645$ and 2.326, month-specific standard deviations, and the 97.5% quantile of the scaled t(4) distribution.

Observed-closure replays. Two closure patterns reported for the corridor (Section 4.10) are replayed on top of the normal regime: a 2026-type event, in which a continuous closure of about 34 days from mid-July to mid-August leaves 14 of 31 operating days in each of July and August (capacity multiplier 0.45), and a 2023-type season with 58 closure days distributed, by assumption, as 12, 18, 16, and 12 days in June–September.

Performance metrics. Let $Z^{(n)}$ be the annual cost associated with replication n, and let $Z_{(1)} \le \cdots \le Z_{(N)}$ denote the ordered annual costs. Expected cost, value-at-risk, and conditional value-at-risk are defined as:

$$E[Z] = \frac{1}{N}\sum_{n=1}^{N} Z^{(n)}, \qquad VaR_{95} = Z_{(\lceil 0.95N \rceil)}, \qquad CVaR_{95} = \frac{1}{N - \lceil 0.95N \rceil + 1}\sum_{n \ge \lceil 0.95N \rceil} Z_{(n)} \tag{28}$$

Service performance and realised emissions are measured by

$$Fill = \frac{\sum_{n,t} min\left(d_t^{(n)}, C_t^{(n)}\right)}{\sum_{n,t} d_t^{(n)}}, \qquad P(short) = \frac{1}{NT}\sum_{n,t} 1\left[d_t^{(n)} > C_t^{(n)}\right], \qquad E[E(x,d)] = \frac{1}{N}\sum_{n} E\left(x, d^{(n)}\right) \tag{29}$$

The resulting indicators are expected annual cost $E[Z]$, tail risk $CVaR_{95}$, fill rate, monthly shortage probability $P(short)$, and expected realised emissions $E[E(x,d)]$.

Price of robustness. The price of robustness at protection level Γ is defined as the incremental expected cost under normal operating conditions relative to the deterministic solution:

$$PoR(\Gamma) = E_{normal}[Z(x_\Gamma)] - E_{normal}[Z(x_0)] \tag{30}$$

Rolling-origin evaluation. For each planning year, Y = 2019, ..., 2025, Equations (1)–(6) are re-estimated using information available up to December of $Y - 1$. Models M0, M1 ($\Gamma = 4$), and M2 are then re-solved and evaluated over year Y. This procedure yields seven rolling-origin planning periods, corresponding to 84 monthly observations per model.

Experimental design. Table 9 summarises the computational experiments. One-factor sensitivity tests vary individual parameters around the baseline ($\Gamma = 4$), whereas

two-factor grids examine interactions between Γ and selected economic, operational, and environmental parameters.

Table 9. Computational experimental design.

Experiment	Factors and Levels	Models	Runs
E1 Robustness frontier	$\Gamma = \{0, 0.5, 1, 2, 4, 6, 8, 12\}$	M1	8
E2 Model comparison	Baseline	M0, M1, M2, EV	4
E3 Cost–emission frontier	ε = 100–90% of $\bar{E}(x^*)$ (1% steps)	M3	11
E4 Sensitivity/ablation	T01–T10; 2–8 levels/test	M1	45
E5 Rolling-origin evaluation	Y = 2019–2025	M0, M1, M2	21
E6 Two-factor interactions	Γ × {penalty, carbon price, deviation scale, disruption loss, CO_2e budget}	M1, M3	224
E7 Out-of-sample stress	OOS-1, OOS-2, OOS-3	M0, M1, M1-DC, M2, EV	24
E8 Joint demand–capacity robustness	$\hat{\rho} \in \{0.14, 0.30\}$; $\Gamma_c \in \{0, 1, 2, 3\}$	M1-DC	8
E9 Forecast-error specification	6 error processes × 2 regimes; 5 deviation calibrations	M0, M1, M2	41
E10 Reserve and emission policy	$r \in \{0, 0.025, 0.05, 0.075, 0.10\}$; 3 instruments	M0, M1, M2, M3	23
E11 Observed-closure replays	2026-type event; 2023-type season	M0, M1, M1-DC, M2, EV	12
E12 Threshold verification	Relaxed Γ-sweep in 0.05 steps; 14 parameter cases	M1, M1-DC (relaxed)	1357

Note: EV = expected-value solution.

Implementation. The baseline models are implemented in MATLAB version R2026b and solved with a relative optimality gap of 10^{-6}. All models have also been re-implemented independently in Python 3 with the open-source HiGHS solver (via SciPy). This implementation reproduces the baseline results in Tables 10 and 11 to within Monte Carlo error (differences below 0.1% in expected cost and below 0.6% in $CVaR_{95}$, comparable to the variation across random seeds) and is used for experiments E7–E12. Monte Carlo simulation uses N = 20,000 replications and fixed seeds to ensure reproducibility; both implementations are provided as Supplementary Materials. The baseline results of the robustness frontier and the comparison of the three capacity plans are reported in Tables 10 and 11, respectively.

Table 10. Robust plans by protection level Γ (20,000 Monte Carlo replications per evaluation regime).

Γ	Trips	Worst-Case Obj. (MUSD)	Normal Operations			Scenario Stress				PoR (MUSD)
			E[cost] (MUSD)	Fill (%)	P(short) (%)	E[cost] (MUSD)	CVaR95 (MUSD)	Fill (%)	P(short) (%)	
0	89,527	51.76	54.00	98.95	23.4	76.00	98.00	89.51	60.0	0.00
0.5	89,627	52.71	53.98	98.98	23.0	75.90	97.87	89.57	59.7	−0.02
1	89,777	53.60	53.96	99.02	22.4	75.76	97.68	89.66	59.3	−0.04
2	96,935	55.16	54.94	99.90	3.2	70.50	90.04	93.19	41.1	0.95
4	97,030	55.57	54.98	99.90	3.1	70.45	89.95	93.23	40.8	0.98
6	97,037	55.95	54.98	99.90	3.1	70.45	89.95	93.23	40.8	0.98
8	97,075	56.32	54.99	99.90	3.1	70.43	89.91	93.25	40.7	0.99
12	97,086	57.00	54.99	99.90	3.1	70.42	89.91	93.25	40.7	1.00

Note: PoR = price of robustness, measured relative to Γ = 0 under normal operations; the small negative values at Γ = 0.5 and 1 are within Monte Carlo error.

Table 11. Monte Carlo performance of the three capacity plans (20,000 replications per evaluation regime).

Plan	Trips	Normal Operations				Scenario Stress				CO_2e Under Stress (kt)
		E[cost] (MUSD)	$CVaR_{95}$ (MUSD)	Fill (%)	P(short) (%)	E[cost] (MUSD)	$CVaR_{95}$ (MUSD)	Fill (%)	P(short) (%)	
Deterministic	89,527	54.00	58.72	98.95	23.4	76.00	98.00	89.51	60.0	31.86
Robust ($\Gamma = 4$)	97,030	54.98	56.95	99.90	3.1	70.45	89.95	93.23	40.8	33.18
Stochastic	106,140	58.41	59.14	99.998	0.1	67.77	83.94	95.97	23.9	34.15

Note: E[cost] = expected annual cost; CVaR95 = 95% conditional value-at-risk; P(short) = monthly shortage probability; CO2e = carbon dioxide equivalent. The last column reports expected realised emissions under scenario stress; normal-operation P(short) for the stochastic plan is from the re-implementation described in Section 3.9.

4. Results

4.1. Baseline Calibration and Evaluation Setting

The baseline calibration uses the designed 144-month series covering January 2014–December 2025 and determines monthly truck capacity for 2026. The mean forecast demand is 151,179 t/month. The log-linear trend-and-seasonality model yields a residual standard deviation of 0.0707 in logarithms, corresponding to a monthly forecast error of approximately 7.1%.

The baseline parameter values are reported in Table 7, while the computational experiments are summarised in Table 9. The product mix is incorporated through three share-weighted factors: payload ($\phi_{PF} = 0.9358$), delay cost ($\phi_{DCI} = 1.3320$), and emission intensity ($\phi_{EIA} = 1.0140$). These adjustments result in an effective payload of 21.90 t/trip, an effective shortage penalty of 126.5 USD/t, and emissions of 0.403 tCO_2e/trip.

Results are reported for the deterministic, Γ-robust, and two-stage stochastic plans under the normal and scenario-stress evaluation regimes defined in Section 3.9. Monetary results are expressed in million US dollars (MUSD).

4.2. Robustness Frontier and the Price of Robustness

The robustness frontier exhibits a pronounced threshold effect. The main adjustment occurs between $\Gamma = 1$ and $\Gamma = 2$, where contracted capacity increases from 89,777 to 96,935 trips (+8.0%). Under scenario stress, expected annual cost decreases from 75.76 to 70.50 MUSD (−6.9%), and $CVaR_{95}$ decreases from 97.68 to 90.04 MUSD (−7.8%). Under normal operations, the monthly shortage probability simultaneously decreases from 22.4% to 3.2%.

Beyond $\Gamma = 2$additional protection produces only marginal changes in the capacity plan. Increasing Γ from 2 to 12 adds 151 trips (+0.16%), while the worst-case objective increases from 55.16 to 57.00 MUSD. Monte Carlo performance remains essentially unchanged over this range (Table 10 and Figure 2).

The threshold is predicted by Proposition 1. With the baseline coefficients, the average marginal protection cost and benefit are $\underline{k} = 18.82$ USD/t and $\underline{k} = 117.75$ USD/t, so Equation (21) gives $\Gamma^* = 1.92$ months. A sweep of the continuous relaxation in steps of 0.05 locates the jump between $\Gamma = 1.90$ and $\Gamma = 1.95$, where total capacity increases from 90,225 to 96,985 trips; for integer Γ, the transition therefore occurs between $\Gamma = 1$ and $\Gamma = 2$, as observed in Table 10. The plateau condition of Proposition 1 holds comfortably ($\lambda_{lo} = 0.20$ MUSD $< \lambda_{hi} = 1.61$ MUSD). The simplified screening version in Equation (31), which neglects the variable service cost and the expected disruption loss, gives $\Gamma^* \approx 1.77$ and slightly understates the exact threshold.

The price of robustness is comparatively small. Under normal operations, protection at $\Gamma = 4$ increases expected annual cost by 0.98 MUSD (+1.8%). Under scenario stress, the same plan reduces expected cost by 5.55 MUSD and $CVaR_{95}$ by 8.05 MUSD, corresponding

to a stress-cost reduction approximately 5.7 times the normal-operation price of robustness. The monthly shortage probability under scenario stress falls from 60.0% at $\Gamma = 0$ to 41.1% at $\Gamma = 2$, and remains close to 41% thereafter. The price of robustness similarly stabilises at approximately 0.95–1.00 MUSD beyond $\Gamma = 2$.

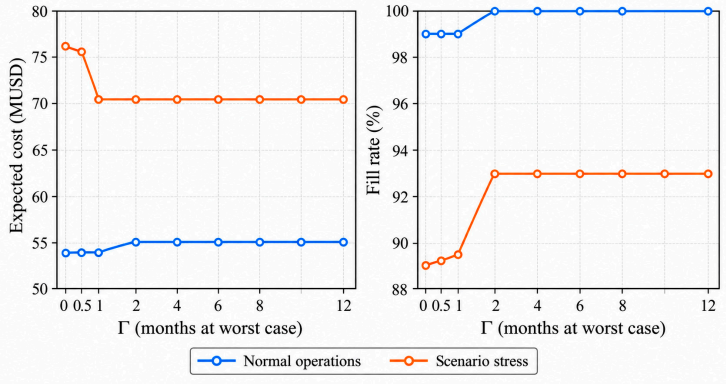


Figure 2. Expected cost (**left**) and fill rate (**right**) as a function of Γ, under normal operations and scenario stress.

4.3. Demand Forecast and Monthly Capacity Plans

The 2026 forecast preserves the seasonal pattern of the calibration series, with higher demand during March–May and lower demand during July–September. Forecast demand ranges from approximately 139 kt in July to 168 kt in March, corresponding to a peak-to-trough variation of approximately 21%. The nominal forecast, its 95% prediction interval, and the corresponding worst-case demand path are reported in Figure A1.

The three models generate distinct monthly capacity profiles (Figure 3).

The deterministic plan contracts 89,527 trips and closely follows nominal demand with the minimum reserve. At $\Gamma = 4$, the robust plan increases capacity to 97,030 trips (+8.4%), with the additional capacity concentrated in higher-demand months. The stochastic plan contracts 106,140 trips (+18.6% relative to the deterministic plan) and remains above the other two profiles throughout the year because it also accounts for joint demand–capacity scenarios.

This difference reflects the distinct uncertainty structures represented by M1 and M2.

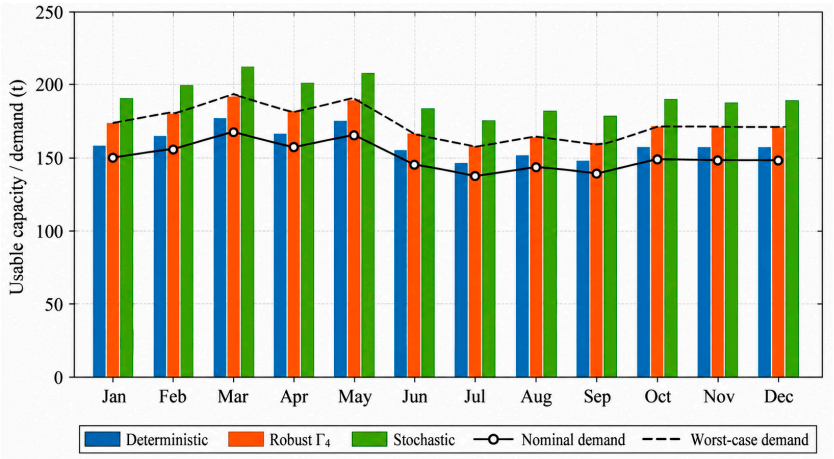


Figure 3. Monthly usable capacity under the deterministic, robust ($\Gamma = 4$), and stochastic plans relative to nominal and worst-case demand, 2026.

4.4. Deterministic, Robust, and Stochastic Plans Compared, Including Out-of-Sample Stress Tests

The ranking of the three plans differs across evaluation regimes (Table 11).

Under normal operations, the deterministic plan has the lowest expected cost (54.00 MUSD), whereas the robust plan has the lowest $CVaR_{95}$ (56.95 MUSD) at an additional expected cost of 0.98 MUSD. The stochastic plan has the highest expected cost, at 58.41 MUSD. Under scenario stress, the stochastic plan achieves the lowest expected cost (67.77 MUSD), the lowest $CVaR_{95}$ (83.94 MUSD), and the highest fill rate (95.97%). The robust plan occupies an intermediate position, while the deterministic plan has the highest expected cost (76.00 MUSD) and lowest fill rate (89.51%).

Relative to the deterministic plan under scenario stress, the robust plan captures 5.55 MUSD of the 8.23 MUSD deterministic-to-stochastic reduction in expected cost (67%). It requires 9110 fewer contracted trips than the stochastic plan, produces approximately 1 kt less CO_2e in expected realised emissions, and costs 3.43 MUSD less under normal operations.

Within the nine-scenario stochastic formulation, $RP = 66.45$ MUSD and $EEV = 67.86$ MUSD, yielding $VSS = 1.41$ MUSD (2.08%).

This VSS is computed on the same nine scenarios used to construct M2 and is therefore an in-sample measure. Its small magnitude does not appear to be an artefact of in-sample evaluation. Rather, it reflects that the expected-value plan already contains a substantial buffer (100,634 trips, because the reserve is applied to the expected-value demand and capacity). On the independent stress sets in Table 12, the expected-cost difference between the expected-value and stochastic plans is 1.12–1.27 MUSD under OOS-1 and OOS-2 and increases to 2.79 MUSD under the shifted regime OOS-3.

Table 12. Expected annual cost and $CVaR_{95}$ (MUSD) under in-sample and independent out-of-sample stress sets (20,000 replications each).

Plan	Trips	In-Sample Stress		OOS-1 Perturbed		OOS-2 Continuous		OOS-3 Shifted	
		E[cost]	$CVaR_{95}$	E[cost]	$CVaR_{95}$	E[cost]	$CVaR_{95}$	E[cost]	$CVaR_{95}$
Deterministic	89,519	75.94	97.46	76.06	106.38	75.50	95.97	87.38	117.61
Robust ($\Gamma = 4$)	97,017	70.41	89.40	70.71	98.31	69.44	87.54	80.09	108.14
M1-DC ($\hat{\rho} = 0.14, \Gamma_c = 2$)	98,106	69.90	88.47	70.23	97.37	68.88	86.59	79.28	106.89
M1-DC ($\hat{\rho} = 0.30, \Gamma_c = 2$)	99,467	69.41	87.46	69.75	96.31	68.33	85.58	78.39	105.40
Expected value (EV)	100,634	68.90	86.57	69.28	95.41	67.78	84.65	77.57	104.19
Stochastic	106,131	67.74	83.43	68.16	92.00	66.51	81.56	74.78	99.00
Robust share of Det.→Sto. reduction (%)		67		68		67		58	
Out-of-sample VSS: EV − Sto. (MUSD)		1.16		1.12		1.27		2.79	

Note: OOS-1 rescales scenario severities and redraws probabilities; OOS-2 uses continuous distributions matched to the first two moments of the scenario mixture; OOS-3 doubles the closure probability, raises mean severity by 25%, and adds month-to-month persistence (Section 3.9). In-sample values are re-simulated with an independent random stream and differ from Table 11 by less than the Monte Carlo variation across seeds. Robust share = (Det. − Robust)/(Det. − Sto.) in expected cost.

Asymmetry between M1 and M2. The comparison above is not between equally informed models. M1 hedges adversarially against demand deviations only, and capacity disruption enters through its expected value, whereas M2 explicitly represents joint demand–capacity–fuel shocks, including capacity losses of up to 32%. Part of the stress-cost gap between the two plans therefore reflects the uncertainty represented rather than the robust–stochastic paradigm itself. Two experiments address this asymmetry: independent out-of-sample stress sets (Table 12) and the joint demand–capacity robust model M1-DC (Table 13).

Table 13. Joint demand–capacity robust plans (M1-DC, $\Gamma_d = 4$); costs in MUSD.

$\hat{\rho}$	Γ_c	Trips	Worst-Case Obj. (MUSD)	Normal Operations	In-Sample Stress		OOS-3 Shifted	
				E[cost]	E[cost]	Fill (%)	E[cost]	Fill (%)
0.14	0	97,017	55.59	55.00	70.41	93.25	80.09	88.20
0.14	1	97,423	58.27	55.14	70.22	93.41	79.78	88.40
0.14	2	98,106	60.72	55.38	69.90	93.67	79.28	88.74
0.14	3	110,962	61.17	60.37	67.57	96.89	73.19	93.50
0.30	0	97,017	55.59	55.00	70.41	93.25	80.09	88.20
0.30	1	97,866	61.40	55.30	70.04	93.56	79.48	88.61
0.30	2	99,467	66.76	55.88	69.41	94.12	78.39	89.35
0.30	3	132,772	69.91	69.10	71.23	99.01	72.20	97.46
Stochastic (M2)		106,131	–	58.44	67.74	95.99	74.78	92.01

Note: $\Gamma_c = 0$ coincides with M1 ($\Gamma = 4$). The worst-case objective of the compact counterpart (25) equals the exact worst case obtained by dynamic programming in every case. Values are from the Python re-implementation (Section 3.9).

Out-of-sample stress evaluation. Table 12 evaluates all plans on the in-sample stress regime and on the three independent stress sets. The ranking is unchanged in every set: the stochastic plan has the lowest expected cost and $CVaR_{95}$, followed by the expected-value plan, the two M1-DC plans, the robust plan, and the deterministic plan. The share of the deterministic-to-stochastic expected-cost reduction captured by the robust plan is 67% in-sample, 68% under OOS-1 and 67% under OOS-2, so this conclusion is not an artefact of evaluating M2 on its own scenarios. Under the shifted regime OOS-3, however, the share falls to 58%, because the robust plan does not hedge against the more frequent and more persistent capacity losses. The M1-DC plan with $\hat{\rho} = 0.30$ and $\Gamma_c = 2$ raises this share to 71% while contracting 6.3% fewer trips than the stochastic plan. Under normal operations, the ordering is reversed, with expected costs of 55.00 (robust), 55.38 and 55.88 (M1-DC), 56.29 (expected value) and 58.44 MUSD (stochastic).

Joint demand–capacity robustness. Table 13 reports M1-DC for $\hat{\rho} \in \{0.14, 0.30\}$ and $\Gamma_c \in \{0, 1, 2, 3\}$ with $\Gamma_d = 4$. For $\Gamma_c \leq 2$, capacity robustness adds a moderate layer of 406–2450 trips to the robust plan, raises the expected normal-operation cost by 0.14–0.88 MUSD, and lowers the expected stress cost by 0.2–1.0 MUSD in-sample and by 0.3–1.7 MUSD under OOS-3. Consistent with Proposition 2, a second threshold appears at $\Gamma_c^* = 2.19$ for $\hat{\rho} = 0.14$ and $\Gamma_c^* = 2.63$ for $\hat{\rho} = 0.30$; the relaxation sweep locates the jumps in $[2.15, 2.20]$ and $[2.60, 2.65]$, respectively. Beyond this threshold, every month is protected against a simultaneous demand surge and capacity losses. With $\hat{\rho} = 0.14$ and $\Gamma_c = 3$, the plan contracts 110,962 trips (4.6% more than the stochastic plan) and attains slightly lower expected stress cost than the stochastic plan both in-sample (67.57 vs. 67.74 MUSD) and under OOS-3 (73.19 vs. 74.78 MUSD), at a normal-operation premium of 1.93 MUSD. With $\hat{\rho} = 0.30$, full protection requires 132,772 trips and is clearly over-conservative under normal operations (69.10 MUSD). Hence, once capacity loss is represented in the uncertainty set, a distribution-free robust plan can match the stress performance of the stochastic plan; the remaining choice concerns the normal-operation cost that the planner is willing to pay for protection without probabilistic information.

For comparison, Figure 4 shows the annual-cost distributions of the deterministic, robust, and stochastic plans under scenario stress.

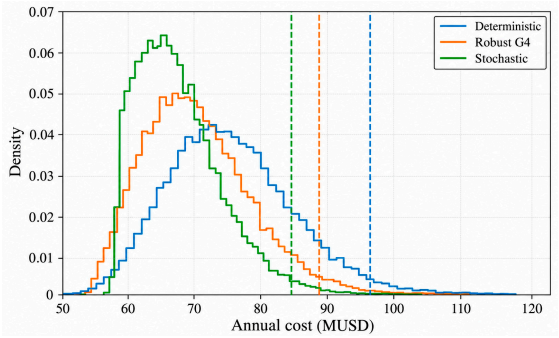


Figure 4. Annual-cost distributions under scenario stress for the deterministic, robust, and stochastic plans. Dashed vertical lines indicate $CVaR_{95}$.

4.5. Cost–Emissions Trade-Off

The contracted-fleet CO_2e budget can be reduced to 93% of its baseline value before the model becomes infeasible. Baseline contracted-fleet emissions are 39.08 kt CO_2e for the $\Gamma = 4$ solution.

Budgets from 100% to 93% of baseline remain feasible; at 92% and below, the 5% minimum reserve constraint cannot be satisfied. Across the feasible range, the fill rate decreases from 99.90% to 99.13%, while the monthly shortage probability increases from 3.1% to 20.3% (Table 14).

Table 14. ε-constraint solutions for the robust model ($\Gamma = 4$) under normal operations.

CO_2e Budget (% of Baseline)	Budget (kt)	Worst-Case Cost (MUSD)	E[cost] (MUSD)	Fill (%)	P(short) (%)
100	39.08	55.57	54.97	99.90	3.1
97	37.91	56.81	54.18	99.73	7.6
94	36.73	58.08	53.85	99.34	16.2
93	36.34	58.50	53.91	99.13	20.3
≤92	≤35.95	infeasible	–	–	–

The worst-case cost increases approximately linearly by 0.42 MUSD for each one-percentage-point reduction in the emission budget. This corresponds to an implied robust abatement cost of approximately 1070 USD/tCO_2e, calculated as $(58.50 - 55.57)$ MUSD/$(39.08 - 36.34)$ kt, approximately 19 times the baseline carbon price of 55 USD/tCO_2e.

Expected cost under normal operations initially decreases, reaching a minimum of 53.85 MUSD at the 94% budget, as moderate tightening removes protective capacity that is rarely used under normal conditions. Below this point, increasing shortage costs offset the capacity savings (Figure 5). Overall, tighter emission budgets progressively reduce service performance while producing only limited changes in expected normal-operation cost.

Sensitivity to the minimum reserve. Because infeasibility of the ε-constraint is caused by the reserve requirement, Table 15 varies r from 0 to 0.10. The reserve affects neither the robust plan at $\Gamma = 4$ (97,017 trips in all cases, because the protected capacity U_t exceeds the reserve level L_t) nor the location of the threshold: $\Gamma^* = 1.92$ for every r, as implied by Proposition 1, and the integer transition always occurs between $\Gamma = 1$ and $\Gamma = 2$. The reserve instead determines the deterministic plan and the room for decarbonisation under a contracted-fleet cap. The maximum feasible reduction in contracted-fleet emissions falls from 12.1% at r = 0 to 7.7% at r = 0.05 and 3.3% at r = 0.10, whereas the marginal worst-case cost of tightening the cap remains 0.42 MUSD per percentage point. The reserve is therefore the binding bottleneck for decarbonisation through contracted-capacity caps, but not a driver of the robustness threshold. Without a reserve, the deterministic plan has a monthly

shortage probability close to 50% under normal operations, and the price of robustness becomes negative (−1.10 MUSD).

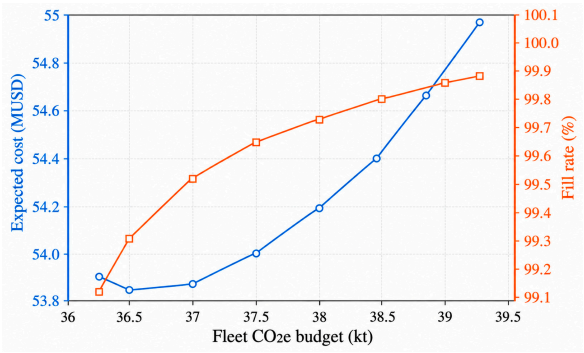


Figure 5. ε-constraint frontier for expected cost and fill rate versus the contracted-fleet CO_2e budget ($\Gamma = 4$).

Table 15. Sensitivity of capacity, robustness threshold, and emission-cap feasibility to the minimum reserve r.

r	Trips		Γ^* (Equation (21))	Integer Transition	PoR (MUSD)	Stress E[cost] (MUSD)		Max. CO_2e Reduction (%)	Slope (MUSD/pp)
	Det.	Robust ($\Gamma = 4$)				Det.	Robust		
0.00	85,257	97,017	1.92	$\Gamma = 1 \rightarrow 2$	−1.10	80.76	70.41	12.1	0.42
0.025	87,389	97,017	1.92	$\Gamma = 1 \rightarrow 2$	0.27	78.18	70.41	9.9	0.42
0.05	89,519	97,017	1.92	$\Gamma = 1 \rightarrow 2$	0.97	75.94	70.41	7.7	0.42
0.075	91,652	97,017	1.92	$\Gamma = 1 \rightarrow 2$	1.11	74.00	70.41	5.5	0.42
0.10	93,785	97,017	1.92	$\Gamma = 1 \rightarrow 2$	0.86	72.37	70.41	3.3	0.42

Note: Max. CO_2e reduction is the largest feasible cut in contracted-fleet emissions relative to the robust plan of each row; slope is the increase in worst-case cost per percentage point of reduction. PoR is measured under normal operations relative to the deterministic plan with the same r. Values from the Python re-implementation; the r = 0.05 row corresponds to Tables 10 and 14.

Choice of emission-policy instrument. The contracted-fleet cap in M3 constrains latent emissions $\underline{E}(x)$, which include protective capacity that is rarely operated. Table 16 compares three instruments applied to the robust plan. A contracted-fleet cap at 95% of baseline reduces contracted emissions by 1.95 kt but realised emissions by only 0.13 kt (0.4%), while the monthly shortage probability rises from 3.1% to 12.8%; it therefore penalises resilience buffers rather than emissions. A cap on realised emissions would not penalise unused buffers: all three plans emit 33.0–33.4 kt under normal operations, well below the contracted-fleet budget at which M3 becomes infeasible. However, because realised emissions are proportional to the freight served for a given vehicle technology, a binding realised cap acts as a volume cap: a 5% reduction would require leaving about 91 kt of demand unserved, at an implied abatement cost of about 5750 USD/tCO_2e. An emission-intensity standard (gCO_2e per tonne-km) is neutral with respect to both capacity buffers and volume: it is identical across the deterministic, robust, and stochastic plans (48.4 gCO_2e/t-km at baseline) and can be met only by efficiency measures. Raising payload from 26 to 27.4 t reduces intensity by 5.1%, realised emissions by 5.1%, and expected normal-operation cost by 2.64 MUSD without affecting the fill rate. Intensity-based instruments therefore align decarbonisation with the efficiency lever identified in Section 4.6 without discouraging resilience.

Table 16. Comparison of emission-policy instruments applied to the robust plan (normal operations).

Instrument (Robust Plan, $\Gamma = 4$)	Contracted CO_2e (kt)	Realised CO_2e (kt)	Fill (%)	P(short) (%)	E[cost] (MUSD)	Abatement Cost (USD/tCO_2e)
No emission instrument	39.07	33.34	99.90	3.1	55.00	–
Contracted-fleet cap, 95% (M3)	37.12	33.20	99.50	12.8	53.93	1070 (latent); 15,570 (realised) [a]
Realised-emission cap, 95% [b]	39.07	31.67	94.9	–	64.59	5750
Intensity standard, −5% (payload 27.4 t) [c]	37.08	31.63	99.90	3.1	52.36	<0 (cost saving)

Note: [a] Increase in worst-case cost divided by the reduction in latent (contracted) and realised emissions, respectively. [b] Realised emissions are proportional to freight served; meeting the cap requires about 91 kt of demand to remain unserved, valued at $b - \bar{a}$. [c] Intensity falls from 48.4 to 45.9 gCO_2e/t-km.

4.6. Sensitivity and Ablation Analysis

The sensitivity and ablation analyses distinguish the effects of economic and operational parameters on expected cost, capacity, and service performance. Complete results are reported in Appendix B (Table A1).

Expected cost is substantially m ore responsive than fill rate to most parameter changes, with shortage penalty and truck payload producing the largest effects (Figure 6).

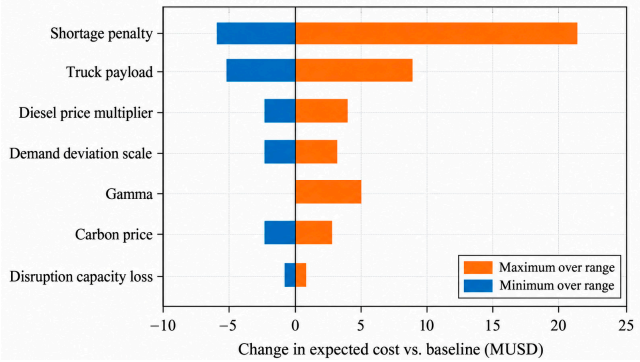


Figure 6. Tornado diagram of expected-cost changes relative to the baseline (70.45 MUSD) across the tested sensitivity ranges.

Truck payload is the strongest operational lever: increasing payload from 22 to 30 t reduces contracted trips by 26.6%, expected cost by 16.41 MUSD, and realised CO_2e emissions by 26.7%, with negligible change in fill rate. The shortage penalty produces the largest variation in expected cost, increasing it from 62.62 to 92.20 MUSD across the tested range, while contracted capacity changes by only approximately 0.3%. By contrast, carbon pricing has negligible influence on the capacity decision: increasing the carbon price to 150 USD/tCO_2e changes capacity by only eight trips (T04), although expected cost rises from 68.63 to 73.61 MUSD. Removing disruption-related capacity loss reduces contracted capacity from 97,030 to 94,285 trips (−2.8%; T08), while scenario-stress fill rate remains essentially unchanged.

4.7. Two-Factor Interactions

Two-factor experiments examine how the protection level Γ interacts with the shortage penalty, carbon price, demand-deviation scale, disruption severity, and carbon budget under scenario stress. Figure 7 illustrates the interaction with the shortage penalty.

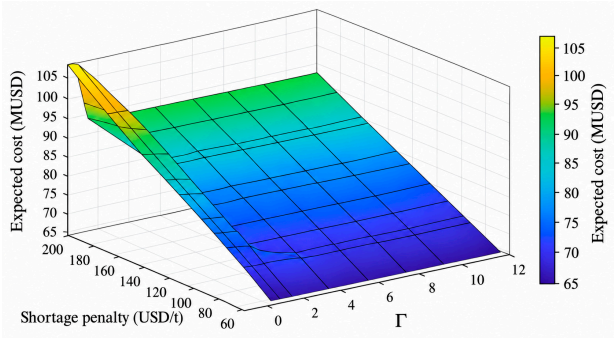


Figure 7. Expected annual cost under scenario stress as a function of Γ and the shortage penalty.

Γ × shortage penalty. A clear transition occurs between $\Gamma = 1$ and $\Gamma = 2$ across the tested shortage penalties (Figure 7). For $\Gamma \leq 1$, expected cost is highly sensitive to the penalty, reaching approximately 107 MUSD at 220 USD/t, whereas for $\Gamma \geq 2$ the response is smaller. At the highest penalty, the expected-cost difference between low- and high-protection solutions reaches approximately 15 MUSD.

Increasing the carbon price raises expected cost across protection levels but has negligible effects on contracted capacity or realised emissions, indicating that carbon pricing primarily affects cost rather than the capacity decision within the tested range.

Demand and disruption uncertainty generate different responses. Increasing the demand-deviation scale improves the fill rate mainly after the protection threshold is reached: service remains close to 90% for $\Gamma \leq 1$, whereas for $\Gamma \geq 2$ additional protective capacity raises it to approximately 95%. By contrast, increasing disruption-related capacity loss from 0.10 to 0.40 raises contracted capacity by approximately 4.1% with little change in fill rate. Thus, demand uncertainty interacts with the protection threshold, whereas disruption severity primarily affects the capacity required to maintain service.

The carbon-budget interaction also reveals two regimes. Low-protection solutions ($\Gamma \leq 1$) have lower contracted emissions but higher expected cost and shortage risk, whereas high-protection solutions ($\Gamma \geq 2$) reduce both cost and shortage probability at the expense of approximately 1.3 kt additional CO_2e. Within the high-protection regime, tighter carbon budgets progressively reduce the fill rate.

At low protection levels, by contrast, the emission constraint remains largely non-binding.

4.8. Rolling-Origin Evaluation, 2019–2025

The rolling-origin experiment provides a time-ordered comparison of the three planning approaches within the designed 2014–2025 series. For each year from 2019 to 2025, the models are re-estimated using only preceding observations and then evaluated against the corresponding held-out annual values of demand, disruption index, and diesel price from the designed series. Because both the estimation windows and the held-out years are taken from the designed series, this experiment checks the internal consistency of the planning procedure over time; it is not an empirical out-of-sample validation with observed corridor data.

Relative to the deterministic plan, the robust plan reduces the number of shortage months from 22 to 1 over the 84 evaluated months (−95%), while increasing mean annual cost by 1.38 MUSD (+2.8%).

Relative to the robust plan, the stochastic plan eliminates the remaining shortage month but increases mean annual cost by a further 3.25 MUSD (+6.3%).

Annual realised emissions are similar across the three plans, ranging from 31.19 to 31.45 kt CO_2e/year, indicating that most additional protective capacity is not ultimately operated.

The main rolling-origin performance indicators for the three planning approaches are summarised in Table 17.

Table 17. Rolling-origin evaluation performance, averages over 2019–2025 (7 planning years, 84 months).

Plan	Annual Cost (MUSD/Year)	Fill Rate (%)	Shortage Month per/Year	Total Shortage Months	CO_2e (kt/Year)
Deterministic	50.06	99.18	3.14	22	31.19
Robust ($\Gamma = 4$)	51.43	99.98	0.14	1	31.44
Stochastic	54.68	100.00	0.00	0	31.45

The cost gap between the stochastic and robust plans increases from approximately 3 MUSD in 2019 to approximately 3.5 MUSD in 2025 (Figure 8, right). The deterministic plan exhibits the lowest annual fill rates, reaching approximately 98.5% in 2019 and 2023 (Figure 8, left).

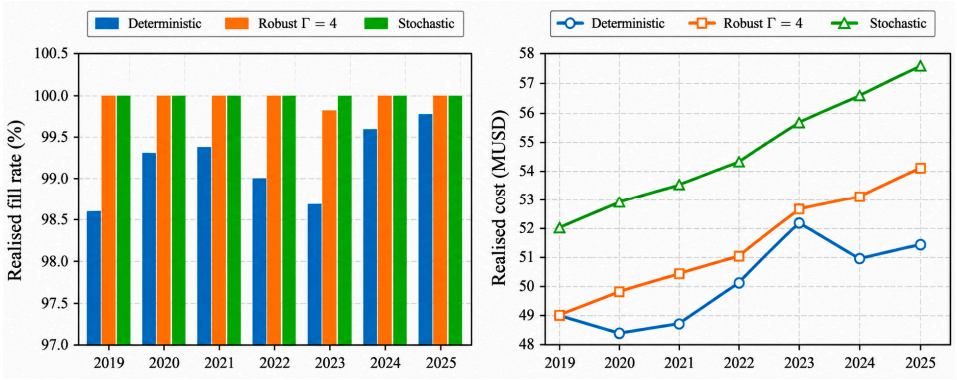


Figure 8. Rolling-origin evaluation of the deterministic, robust $\Gamma = 4$, and stochastic plans: annual fill rate (**left**) and annual cost (**right**), 2019–2025.

4.9. Alternative Forecast-Error Specifications

Table 18 re-evaluates the three plans when the i.i.d. lognormal error assumption (A11) is replaced by the alternatives described in Section 3.9. Expected costs and fill rates are insensitive to the error specification: across all specifications, expected normal-operation cost changes by at most 0.65 MUSD and the fill rate by at most 0.3 percentage points. The main effect is on tail risk. Under AR(1) errors, high-demand months cluster within years, and $CVaR_{95}$ under normal operations increases by 5.2% for the deterministic plan, 1.6% for the robust plan, and 0.8% for the stochastic plan; under the combined specification, the increases are 16.5%, 10.3%, and 5.2%, respectively. The robust plan retains the lowest normal-operation $CVaR_{95}$ in every single-departure specification but loses it to the stochastic plan under the combined specification (62.89 vs. 62.27 MUSD). Under scenario stress, the ranking of the three plans is unchanged in every specification.

The deviation calibration affects the size of the protection layer but not the threshold. Re-optimising M1 with $z_{SL} = 1.645$, 1.96, and 2.326 gives 95,127, 97,017, and 99,213 trips at $\Gamma = 4$; month-specific deviations give 96,851 trips, and the 97.5% quantile of the scaled t(4) distribution gives 97,036 trips. In all cases, $\Gamma^* = 1.92$ and the integer transition occurs between $\Gamma = 1$ and $\Gamma = 2$, as predicted by Proposition 1. Thus, misspecification of the error process mainly changes tail risk and the optimal size of the protection layer, whereas the location of the robustness threshold is determined by the cost structure.

Table 18. Performance of the three plans under alternative forecast-error specifications (MUSD; N = normal operations, S = scenario stress).

Error Specification	Deterministic			Robust ($\Gamma = 4$)			Stochastic		
	E[cost] N	$CVaR_{95}$ N	E[cost] S	E[cost] N	$CVaR_{95}$ N	E[cost] S	E[cost] N	$CVaR_{95}$ N	E[cost] S
i.i.d. lognormal (baseline)	54.02	58.82	75.92	55.00	57.01	70.41	58.44	59.18	67.75
AR(1), $\varphi = 0.5$	54.02	61.88	75.90	55.00	57.90	70.39	58.44	59.64	67.73
Heteroscedastic (winter SD +30%)	53.98	58.73	75.90	55.01	57.15	70.39	58.44	59.23	67.73
Student-t ($\nu = 4$)	53.75	60.79	75.84	55.23	60.20	70.37	58.57	61.60	67.72
Structural break (25%/year)	54.65	65.04	76.78	55.25	59.91	71.09	58.53	60.22	68.21
Combined (i)–(iv)	54.41	68.55	76.70	55.42	62.89	71.03	58.63	62.27	68.17

Note: All specifications have the same marginal error variance as the baseline, except the structural-break case, which adds a level shift. Values are from the Python re-implementation with a common random stream across specifications.

4.10. Comparison with Public Closure Records and Observed-Closure Replays

The designed panel contains closures in 58 of 144 months, each lasting at most 3 days (on average 0.60 days per month, or about 7 days per year), and its most severe scenarios assume 5–6 closure days (Table 4). Publicly reported closure records for the corridor are considerably more severe. Press reports citing the binational border committee indicate 58 days of closure during the 2023 winter season up to early September [20]; continuous closures reached 12 days in 2022 and 16 days in both 2023 and 2024 [21]; and in 2026 the corridor reopened on 18 August after a record closure of 34 days [22]. The designed disruption component therefore corresponds to a normal-to-moderate regime and understates the multi-week closures observed in recent winters, which reinforces the interpretation of the numerical results as a methodological demonstration.

Table 19 replays the two reported closure patterns described in Section 3.9 on top of the normal regime. No plan dominates across the two patterns. Under the 2026-type event, in which about 55% of capacity is lost in each of two consecutive months, no plausible capacity buffer prevents large shortages; the robust plan has the lowest expected cost (68.89 MUSD) and $CVaR_{95}$ (72.96 MUSD), and the stochastic plan has the highest expected cost of all plans (70.56 MUSD) because its additional capacity in the other months remains idle. Under the 2023-type season, in which closures are spread over four months, the broader buffer of the stochastic plan pays off: it has the lowest expected cost (77.41 MUSD), followed by the expected-value (77.72 MUSD) and robust (78.05 MUSD) plans. In both cases, the deterministic plan has the lowest fill rate. For closures of this depth, tactical capacity buffers alone cannot maintain service, and capacity planning should be complemented by contingency measures such as rerouting through alternative passes, inventory pre-positioning or flexible capacity contracts.

Table 19. Replays of closure patterns reported for the corridor (normal regime plus closure; costs in MUSD).

Plan	Trips	2026-Type Event (34-Day Closure)				2023-Type Season (58 Closure Days)			
		E[cost]	$CVaR_{95}$	Fill (%)	P(short) (%)	E[cost]	$CVaR_{95}$	Fill (%)	P(short) (%)
Deterministic	89,519	69.06	74.72	91.01	36.2	79.80	86.35	85.30	49.0
Robust ($\Gamma = 4$)	97,017	68.89	72.96	92.42	19.3	78.05	83.58	87.40	35.4
M1-DC ($\hat{\rho} = 0.14$, $\Gamma_c = 2$)	98,106	69.25	73.23	92.46	18.4	78.29	83.77	87.48	34.7
M1-DC ($\hat{\rho} = 0.30$, $\Gamma_c = 2$)	99,467	69.77	73.70	92.47	17.8	78.71	84.16	87.55	34.1
Expected value (EV)	100,634	69.48	73.38	92.79	17.4	77.72	83.16	88.16	33.9
Stochastic	106,131	70.56	74.42	93.25	16.7	77.41	82.83	89.24	33.4

Note: 2026-type: capacity multiplier 14/31 in July and August; 2023-type: 58 closure days distributed as 12, 18, 16, and 12 days in June–September (assumption). Capacity multiplier = share of operating days.

5. Discussion and Managerial Implications

This section discusses the theoretical, managerial, and policy implications of the results, together with their transferability and implementation.

5.1. Theoretical Implications

The results yield three main theoretical implications:

Threshold behaviour of budgeted robustness. Budgeted robustness under convex piecewise-linear recourse can exhibit pronounced threshold behaviour. Although the uncertainty budget Γ is commonly analysed through sensitivity sweeps [6,23], the present results show that the capacity decision can remain close to the nominal solution over a range of low protection levels and then change sharply. Proposition 1 shows that, for the continuous relaxation, this change occurs at an exact budget, which can be written in three progressively simpler forms:

$$\Gamma^* = \sum_{t \in T} \frac{c_{fix}/\pi_t + h}{b + h - a_t} \approx \frac{T\left(c_{fix}/\pi + h\right)}{b + h - a} \approx \frac{T\left(c_{fix}/q + h\right)}{p_{sh}\,\varphi_{DCI}} \tag{31}$$

The first expression in Equation (31) is exact for the continuous relaxation under the conditions of Proposition 1; the second replaces monthly coefficients by their averages; and the third, a simplified screening version, additionally neglects the expected disruption loss ($\underline{\pi} = q$) and the term $a - h$ relative to b. At baseline, the three expressions give 1.92, 1.92, and 1.77 months, respectively. Because the minimum reserve shifts the low-protection capacity but cancels from the marginal comparison, it does not affect Γ* (Table 15).

The result extends newsvendor-type critical-ratio logic from robust inventory theory [24] to the allocation of an uncertainty budget across periods. Reporting Γ^* together with the conventional Γ-sweep can therefore improve interpretation when broad regions of the uncertainty-budget range produce little change in the optimal decision. Formally, Γ*/T equals the average critical shortage probability k_t/κ_t of a risk-neutral newsvendor (Appendix C), so the robust budget can be compared directly with the number of shortage months that a risk-neutral planner would accept.

Robust and stochastic planning may play complementary roles. Rather than acting as direct substitutes, robust and stochastic formulations may provide complementary protection depending on the uncertainty represented [23,25]. The out-of-sample experiments sharpen this conclusion: the advantage of the stochastic plan persists on stress sets that are not generated from its own scenarios, while part of the gap is due to the uncertainty represented rather than to the modelling paradigm, as shown by the joint demand–capacity model M1-DC.

Contracted and realised emissions respond differently to resilience buffers. Green road-freight research has largely focused on operational emissions [26,27], while resilience research emphasises the value of capacity buffers [28,29]. The present results connect these perspectives by distinguishing emissions associated with contracted capacity from those generated by actual vehicle operation. Contracted capacity differs by up to 18.6% across plans, whereas realised emissions in the rolling-origin evaluation differ by less than 1% (31.19–31.45 kt/year).

The apparent resilience–sustainability trade-off therefore depends on the accounting basis: contracted-fleet measures can indicate a larger environmental cost than measures based on realised vehicle operation.

5.2. Managerial Insights

Four managerial insights emerge from the analysed calibration. Because the calibration is designed and its disruption component understates recent closures (Section 4.10), the

numerical magnitudes below are illustrative; the qualitative mechanisms, particularly the threshold structure, constitute the principal methodological results, whereas their empirical transferability requires local recalibration and validation with observed data.

Insight 1—The break-even budget can screen candidate protection levels.

Equation (21) provides a screening device for candidate protection levels using fixed trip cost, usable trip capacity, idle-capacity cost, variable service cost, and the effective shortage penalty.

At baseline, $\Gamma^* = 1.92$ months (1.77 months with the simplified version), consistent with the model transition between $\Gamma = 1$ and $\Gamma = 2$. The deterministic plan averages 3.14 shortage months per year in the rolling-origin evaluation, exceeding this threshold.

The threshold decreases as effective shortage cost increases (Table 20), ranging from 2.36 months for industrial inputs to 1.39 months for chemicals and pharmaceuticals; in every case, the relaxation sweep places the jump within 0.05 of the value given by Equation (21). Because Γ^* refers to the worst-case objective, it indicates when protection becomes worthwhile rather than the exact protection level required in a given year.

Table 20. Break-even robustness budget Γ^* under alternative shortage-cost assumptions: exact value, simplified screening value, and numerical transition.

Case	Effective Shortage Cost b (USD/t)	Γ^* Exact, Equation (21)	Γ^* Simplified	Transition in Relaxed Γ-Sweep
Shortage penalty 50 USD/t	66.6	3.91	3.36	[3.90, 3.95]
Shortage penalty 75 USD/t	99.9	2.48	2.24	[2.45, 2.50]
Baseline product mix (95 USD/t)	126.5	1.92	1.77	[1.90, 1.95]
Shortage penalty 130 USD/t	173.2	1.37	1.29	[1.35, 1.40]
Shortage penalty 220 USD/t	293.0	0.79	0.76	[0.75, 0.80]
Industrial inputs (delay-cost index 1.1)	104.5	2.36	2.14	[2.35, 2.40]
Consumer goods (1.3)	123.5	1.97	1.81	[1.95, 2.00]
Machinery and equipment (1.4)	133.0	1.82	1.68	[1.80, 1.85]
Food and agricultural products (1.6)	152.0	1.58	1.47	[1.55, 1.60]
Chemicals and pharmaceuticals (1.8)	171.0	1.39	1.31	[1.35, 1.40]

Note: Simplified values use Equation (31), third expression. Transition = interval of Γ (step 0.05) in which the capacity of the continuous relaxation jumps.

Insight 2—Protective capacity exhibits option-like economic characteristics. The robust plan comprises a base layer of approximately 89,500 trips and a protection layer of about 7500 trips (+8.4%), concentrated in higher-demand months.

The protection layer increases expected normal-operation cost by 0.98 MUSD but reduces scenario-stress expected cost by 5.55 MUSD and $CVaR_{95}$ by approximately 8.05 MUSD.

Under a simplified two-regime interpretation, these cost differences imply an illustrative break-even stress probability of approximately 15%; this value is not an empirically estimated disruption probability.

This asymmetric cost profile resembles a capacity option and motivates evaluation of reservation, standby, or call-off contracts [30]. The normal-operation premium of approximately 130 USD per additional contracted trip provides an illustrative benchmark for such arrangements.

Insight 3—Load efficiency has a stronger effect on capacity and emissions than carbon pricing within the tested ranges. Carbon prices up to 150 USD/tCO_2e produce negligible changes in capacity (tests T04 and T09), reflecting the small carbon-cost component relative to the effective shortage penalty of 126.5 USD/t.

Increasing effective payload from 22 to 30 t simultaneously reduces contracted trips, expected cost, and realised CO_2e without materially affecting service (T05), consistent with the importance of load factor in freight decarbonisation [31].

Accordingly, load planning, freight matching, consolidation, and higher-capacity vehicle configurations may provide more direct operational levers than carbon pricing alone.

Insight 4—Delay cost is a key economic driver. The shortage penalty produces the largest variation in expected cost (Figure 6), making empirical estimation of product-specific delay costs particularly important.

Product-specific delay costs may incorporate spoilage, production interruption, and lost sales and could support differentiated service policies when protective capacity is scarce.

The protection layer is directly linked to the calibrated demand deviation in Equation (4).

Reducing the demand-deviation scale from 1.0 to 0.5 lowers the capacity requirement by approximately 5900 trips (T02), indicating the potential value of information sharing when it genuinely improves forecast precision.

By contrast, mechanically narrowing the assumed uncertainty range without improved forecast accuracy increases expected stress-regime cost by approximately 4.0 MUSD.

5.3. Policy Implications

The results suggest policy considerations for corridor reliability, freight efficiency, and emissions management, although the numerical magnitudes require validation with observed corridor data before application.

Corridor reliability directly affects tactical capacity requirements: each 0.1 increase in disruption-related capacity loss adds approximately 1300 contracted trips per year (T07), whereas omitting disruption reduces capacity by 2745 trips (2.8%; T08). These effects provide an additional criterion for evaluating measures such as winter maintenance, avalanche protection, and digital border processing.

Because carbon pricing has little effect on capacity within the tested range, complementary measures that improve load efficiency—such as higher-capacity vehicle permits, shared consolidation facilities, and digital freight matching—may have greater operational leverage.

Emissions policy should also distinguish between contracted-fleet and realised emissions. A cap on contracted-fleet emissions implies an abatement cost of approximately 1070 USD/tCO_2e of latent emissions, becomes infeasible beyond a 7% reduction under the 5% reserve requirement (beyond 3–12% for reserves between 10% and 0%; Table 15), and reduces realised emissions only marginally while raising shortage risk (Table 16). Such a cap would therefore penalise carriers for holding resilience buffers. A cap on realised emissions avoids this distortion but, for a given vehicle technology, acts as a cap on freight volume when it binds.

An emission-intensity standard (gCO_2e/t-km) is neutral with respect to both capacity buffers and volume, and it rewards the payload and utilisation improvements that the model identifies as the strongest levers. It is therefore the more appropriate primary instrument for tactical capacity decisions, possibly complemented by reporting of contracted capacity for transparency.

5.4. Generalisability and Boundary Conditions

The mechanisms identified may be relevant to freight systems in which capacity is committed before uncertain demand, shortage costs exceed holding costs, and infrastructure disruptions affect usable capacity, including mountain passes, ferry links, and port-hinterland corridors.

Application to other settings requires local re-estimation of the parameters entering Equation (31), together with the associated capacity, cost, and service effects. Empirical transferability cannot be established from the present designed calibration alone. Because the model is homogeneous of degree one in demand and Γ^* does not depend on demand

volume or volatility, transfer of the threshold primarily requires the cost and capacity coefficients in Equation (21), whereas the size of the protection layer requires local estimates of demand volatility and disruption severity.

5.5. Implementation Pathway

The framework can be implemented progressively through data integration, calibration and pilot testing, contracting, and operational monitoring (Figure 9). Each stage should be completed using observed organisational and corridor data before model-based decisions are introduced.

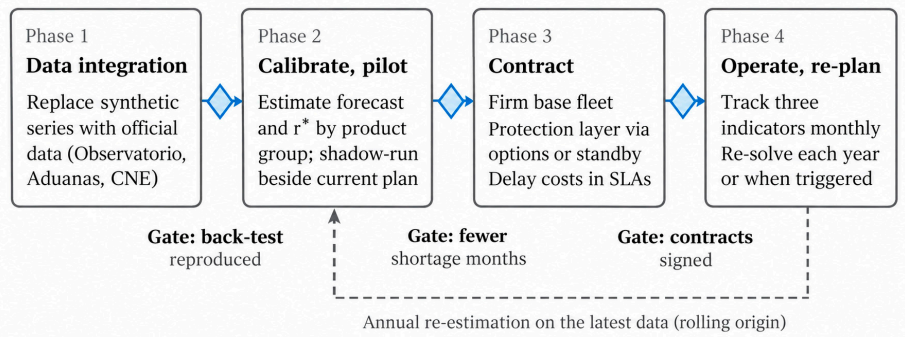


Figure 9. Implementation pathway with decision gates and periodic rolling-origin recalibration.

6. Conclusions, Limitations, and Future Research

6.1. Conclusions

This study develops a budgeted robust mixed-integer programme for tactical truck-capacity planning in a seasonal and disruption-prone cross-border corridor. Convex piecewise-linear recourse represents shortage, idle-capacity, operating, and carbon costs, while linear-programming duality yields a compact MILP robust counterpart. The formulation is benchmarked against deterministic and two-stage stochastic models and extended through a joint demand–capacity uncertainty set and an ε-constraint on contracted-fleet emissions. Performance is assessed through Monte Carlo simulation, independent out-of-sample stress sets, alternative error specifications, replays of reported closure patterns and a seven-year rolling-origin evaluation of the designed series. Because the dataset is designed, the results are a methodological demonstration rather than an empirical validation.

Four findings summarise the analysis.

First, robust protection exhibits pronounced threshold behaviour. The principal capacity adjustment occurs between $\Gamma = 1$ and $\Gamma = 2$, consistent with the exact break-even budget $\Gamma^* \approx 1.9$ months. Beyond this threshold, additional protection produces only marginal changes in contracted capacity and simulation-based performance. The threshold depends only on cost and capacity parameters; it is invariant to the scale of demand uncertainty and to the minimum reserve.

Second, the robust plan substantially improves reliability at a moderate cost premium. At $\Gamma = 4$, expected normal-operation cost increases by 1.8%, whereas scenario-stress expected cost and $CVaR_{95}$ decrease by 6.9% and 7.8%, respectively. In the rolling-origin evaluation, shortage months decrease from 22 to 1 over 84 months, with a 2.8% increase in mean annual cost.

Third, robust and stochastic formulations provide different forms of protection. The robust plan captures 67% of the deterministic-to-stochastic reduction in scenario-stress expected cost while requiring 8.6% fewer contracted trips. The stochastic formulation provides additional protection under joint demand–capacity shocks, whereas the baseline

robust model M1 protects adversarially against demand deviations. The robust plan's share of the deterministic-to-stochastic stress-cost reduction remains 67–68% on independent out-of-sample stress sets but falls to 58% under a shifted regime with more frequent and persistent closures. Treating capacity loss as a second adversarial source (M1-DC) recovers part of this gap with fewer trips than the stochastic plan and, beyond its own threshold Γ_c^*, matches the stress performance of the stochastic plan at a higher normal-operation cost.

Fourth, operational efficiency has a stronger influence on capacity than carbon pricing within the tested ranges. Carbon prices up to 150 USD/tCO_2e produce negligible capacity changes, whereas higher payload reduces contracted capacity, expected cost, and realised emissions. The ε-constraint analysis further shows that tighter contracted-fleet emission limits introduce a service–emission trade-off and become infeasible below approximately 93% of the baseline budget. This limit is set by the minimum reserve rather than by the robustness threshold, and intensity-based instruments avoid penalising resilience buffers.

Overall, the study contributes an exact and interpretable characterisation of the critical robustness budget, a common framework for comparing deterministic, robust, and stochastic capacity decisions, and an explicit distinction between contracted-fleet and realised operating emissions. The numerical magnitudes remain conditional on the designed calibration and therefore require empirical validation before operational or policy application. This caution is reinforced by public records, which indicate more severe closures than those represented in the designed panel.

6.2. Limitations

The results should be interpreted in light of three groups of limitations: data, model structure, and evaluation design.

Data. The monthly panel, scenario set, and product-mix parameters are designed to be economically plausible and reproducible but remain synthetic. Consequently, Γ^*, capacity requirements, the price of robustness, and sensitivity effects are calibration-dependent. The analysis is also restricted to one corridor, one freight direction, and monthly aggregation; return flows and intra-month dynamics are not represented. Public closure records indicate that the designed disruption component understates recent multi-week closures (Section 4.10).

Model. In the baseline robust model, the uncertainty set covers independent, upward demand deviations only, and disruptions enter through expected capacity loss; the joint extension M1-DC treats capacity loss adversarially, but with a fixed depth $\hat{\rho}$ and without temporal dependence. Unserved freight is treated as lost within the month, excluding backlog, rerouting, and modal substitution. The model assumes a homogeneous fleet, advance capacity contracting without spot-market recourse, and an ε-constraint defined on contracted rather than realised emissions. Finally, Γ^* is exact for the continuous relaxation under the conditions of Proposition 1; integrality, heterogeneity across months and fractional budgets make the transition slightly smoother in the integer model.

Evaluation. Independent out-of-sample stress sets are included, but they remain synthetic. The baseline Monte Carlo experiments assume independent lognormal forecast errors; alternative specifications mainly affect tail risk (Section 4.9), and observed freight demand may exhibit further dependence and regime changes. The rolling-origin procedure preserves temporal ordering but remains based on the designed series and therefore does not constitute independent empirical validation. In addition, the analysis adopts a single-planner perspective and does not represent competition, contract negotiation, or cross-border institutional coordination.

6.3. Future Research

The principal research extensions arising from these limitations are summarised in Table 21.

Table 21. Main limitations and associated directions for future research.

Limitation	Research Direction	Candidate Approach
Synthetic calibration	Empirical validation using observed corridor data	Official freight, customs, fuel, disruption, and emission series
Fixed-depth capacity uncertainty without temporal dependence	Joint demand–capacity uncertainty with temporal dependence	Distributionally robust or hybrid robust–stochastic models
Static, here-and-now capacity	Dynamic re-planning as information becomes available	Adjustable robust optimisation
Firm contracts only	Design and pricing of flexible capacity contracts	Real options; shipper–carrier contract models
Single corridor, single vehicle type	Network, modal and technology choice	Multi-corridor network design with fleet-technology decisions
Stationary disruption risk	Climate-driven change in closure frequency and severity	Non-stationary scenario generation; climate-risk scenarios

Empirical validation is the immediate priority. Re-estimating the framework using observed freight, customs, fuel, disruption, and emission data would establish whether the threshold behaviour persists and permit corridor-specific calibration of Γ^*.

Richer uncertainty representations could extend the joint set of M1-DC with temporal dependence, closure durations, and data-driven uncertainty sets using robust optimisation [32] or Wasserstein distributionally robust optimisation [33]. These extensions would permit joint representation of demand and capacity uncertainty and facilitate hybrid robust–stochastic formulations.

Adjustable robust optimisation [34] could allow part of the protection layer to be revised as information becomes available, quantifying the value of dynamic planning flexibility.

Explicit shipper–carrier contracting models could endogenise reservation fees, exercise prices, and delay-cost provisions for flexible capacity [30].

Finally, multi-corridor and multimodal extensions could represent rerouting and modal substitution, while fleet-technology decisions could connect tactical capacity planning with longer-term decarbonisation pathways [31]. Non-stationary disruption scenarios could also be incorporated to evaluate the stability of capacity decisions under evolving corridor conditions [5,35].

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/math14234373/s1, Workbook S1: designed nine-sheet research workbook (monthly panel, model parameters, scenarios, product mix, sensitivity design, source registry and documentation); Code S1: Python implementation of Models M0–M3 and M1-DC, including Monte Carlo simulation, out-of-sample stress tests, forecast-error and observed-closure experiments, and threshold verification.

Author Contributions: Conceptualization, A.K.Y. and C.D.; methodology, A.K.Y. and I.D.; software, A.K.Y. and C.D.; validation, C.D. and I.D.; formal analysis, A.K.Y. and C.D.; investigation, A.K.Y. and C.D.; resources, C.D. and I.D.; data curation, A.K.Y. and I.D.; writing—original draft preparation, A.K.Y. and C.D.; writing—review and editing, C.D. and I.D.; visualization, A.K.Y. and C.D.; supervision, A.K.Y., C.D. and I.D.; project administration, A.K.Y., C.D. and I.D.; funding acquisition, C.D. and I.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received institutional support from Universidad de Santiago de Chile, the Faculty of Engineering of Universidad de Santiago de Chile, and the Department of Industrial Engineering of Universidad de Santiago de Chile.

Data Availability Statement: The original contributions presented in this study are included in the supplementary material (Workbook S1 and Code S1, which will also be deposited in Zenodo upon acceptance). Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors gratefully acknowledge the support of the Center for Innovation in Technology and Design of Materials for the Built Environment—CIMAC_ANID CTI250005.

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

Abbreviation

The following abbreviations are used in this manuscript:

CO2e	Carbon dioxide equivalent
CVaR	Conditional value-at-risk
DRO	Distributionally robust optimisation
LP	Linear programme
M1-DC	Joint demand–capacity robust model
OOS-1/2/3	Out-of-sample stress sets
EEV	Expected result of the expected-value solution
EV	Expected-value solution
$gCO_2e/t\text{-}km$	Grams of carbon dioxide equivalent per tonne-kilometre
MC	Monte Carlo
MILP	Mixed-integer linear programme
MUSD	Million US dollars
OOS	Out-of-sample
PoR	Price of robustness
RO	Robust optimisation
RP	Recourse problem
SP	Stochastic programming
USD/tCO_2e	US dollars per tonne of carbon dioxide equivalent
VaR	Value-at-risk
VSS	Value of the stochastic solution

Appendix A

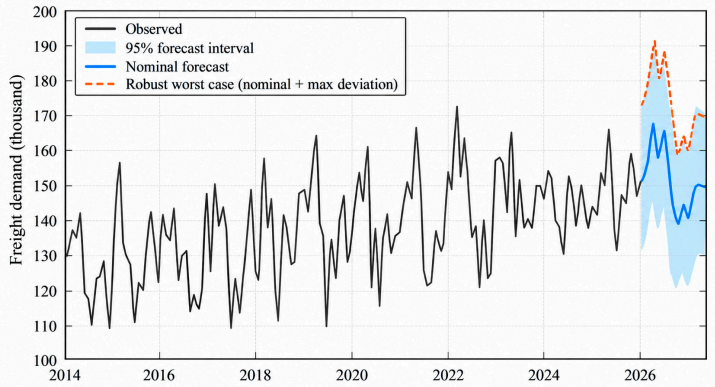


Figure A1. Designed monthly import-freight series used for calibration (2014–2025), the 2026 nominal forecast with its 95% prediction interval, and the corresponding worst-case demand path used in the Γ-robust formulation.

Appendix B

Table A1. Sensitivity and ablation results (robust model, scenario-stress evaluation; values given as low → high level).

Test	Factor	Range	Trips	E[Z] (MUSD)	Fill (%)	CO2e (kt)	Hypothesis
T01	Γ	0 → 12	89,527 → 97,086	76.00 → 70.42	89.51 → 93.25	31.86 → 33.19	Supported, with a threshold at Γ ≈ 2
T02	Demand deviation scale	0.5 → 1.5	91,149 → 102,920	74.49 → 68.31	90.45 → 95.18	32.19 → 33.87	Supported
T03	Shortage penalty (USD/t)	50 → 220	96,770 → 97,065	62.62 → 92.20	93.13 → 93.24	33.14 → 33.19	Weakly supported (+0.3% capacity)
T04	Carbon price (USD/t)	0 → 150	97,031 → 97,023	68.63 → 73.61	93.23 → 93.23	33.18 → 33.18	Not supported
T05	Truck payload (t)	22 → 30	114,640 → 84,105	79.92 → 63.51	93.22 → 93.24	39.21 → 28.76	Supported
T06	Diesel price multiplier	0.8 → 1.4	97,038 → 97,014	68.11 → 75.14	93.23 → 93.22	33.18 → 33.18	Cost only; plan unchanged
T07	Disruption capacity loss	0.1 → 0.4	95,515 → 99,407	70.23 → 70.79	93.23 → 93.23	33.18 → 33.18	Supported
T08	Disruption state (off → on)	0 → 1	94,285 → 97,030	70.05 → 70.45	93.23 → 93.23	33.18 → 33.18	Supported
T09	Emissions in objective (off → on)	0 → 1	97,031 → 97,030	68.63 → 70.45	93.23 → 93.23	33.18 → 33.18	Not supported
T10	Demand uncertainty (off → on)	0 → 1	89,527 → 97,030	76.00 → 70.45	89.51 → 93.23	31.86 → 33.18	Supported

Note: Arrows indicate the change from the lowest to the highest tested level of each factor. $E[Z]$ denotes expected annual cost under scenario stress. CO_2e values represent expected realised emissions.

Appendix C. Proofs of Propositions 1 and 2

Proof of Proposition 1. Step 1 (dual form of the worst case). For fixed x, let $\Delta_t = {g_t}^1 - {g_t}^0 \geq 0$. By linear-programming duality (Equations (18) and (19)) and elimination of μ_t,

$$\max_z \left\{ \sum_t \Delta_t z_t : \sum_t z_t \leq \Gamma,\ 0 \leq z_t \leq 1 \right\} = \min_{\lambda \geq 0} \left\{ \Gamma\lambda + \sum_t (\Delta_t - \lambda)^+ \right\} \tag{A1}$$

Hence, the continuous relaxation of M1 is equivalent to $\min_{\lambda \geq 0} \varphi(\lambda)$, with $\varphi(\lambda) = \Gamma\lambda + \Sigma_t\ \psi_t(\lambda)$ and $\psi_t(\lambda) = \min_{C \geq L} \{c_{fix} C/\pi_t + {g_t}^0(C) + (\Delta_t(C) - \lambda)^+\}$. The function φ is convex because it is the partial minimum of a jointly convex function of (C, λ).

Step 2 (monthly problem). For $C_t \geq U_t$, the increment Δ_t equals $(a_t - h)\ \hat{d}_t\ \hat{d}_t$ and does not decrease further, whereas the cost increases; hence, $C_t > U_t$ is never optimal. For $C_t \in [L_t, U_t]$, nominal demand lies in the excess regime and worst-case demand in the shortage regime of Equation (13), which gives

$$\frac{c_{fix}}{\pi_t} C_t + g_t^0 = k_t C_t + (a_t - h)\, d_t, \qquad \Delta_t(C_t) = A_t - \kappa_t C_t, \qquad A_t = bU_t - (a_t - h)\, d_t \tag{A2}$$

For fixed λ, the monthly objective $k_t\ C_t + (A_t - \kappa_t\ C_t - \lambda)^+$ decreases with slope $k_t - \kappa_t < 0$ until $\Delta_t(C_t) = \lambda$ and increases with slope $k_t > 0$ thereafter. Its unique minimiser is

$$C_t(\lambda) = min\left\{ U_t,\ max\left\{ L_t,\ \frac{A_t - \lambda}{\kappa_t} \right\} \right\}, \qquad \Delta_t(U_t) = (a_t - h)\hat{d_t}, \qquad \Delta_t(L_t) = b\hat{d_t} - \kappa_t r\, d_t \tag{A3}$$

which is non-increasing in λ; condition (ii) ensures $L_t < U_t$.

Step 3 (derivative). From (A3), ${\psi_t}'(\lambda) = -1$ if $\lambda < \Delta_t(U_t)$ (capacity at U_t and positive excess $\Delta_t - \lambda$), ${\psi_t}'(\lambda) = -k_t/\kappa_t$ if $\Delta_t(U_t) < \lambda < \Delta_t(L_t)$ (interior capacity), and ${\psi_t}'(\lambda) = 0$ if $\lambda > \Delta_t(L_t)$. Therefore,

$$\varphi'(\lambda) = \Gamma - \sum_{t:\, \Delta_t(U_t) < \lambda < \Delta_t(L_t)} \frac{k_t}{\kappa_t} - |\{t :\ \lambda < \Delta_t(U_t)\}| \tag{A4}$$

Step 4 (threshold). Under condition (iii), every month is interior for $\lambda \in (\lambda_{lo}, \lambda_{hi})$, so $\varphi'(\lambda) = \Gamma - \Gamma^*$ on this interval. If $\Gamma < \Gamma^*$, φ is strictly decreasing on the interval and,

by convexity, every minimiser satisfies $\lambda^* \geq \lambda_{hi}$. For a minimiser (C^*, λ^*), C^* minimises the monthly problems at λ^*, so $C_t^* = C_t(\lambda^*) \leq C_t(\lambda_{hi})$, which proves (a). If $\Gamma > \Gamma^*$, φ is strictly increasing on the interval, every minimiser satisfies $\lambda^* \leq \lambda_{lo}$, and $C_t^* \geq C_t(\lambda_{lo})$, which proves (b). For identical months, $\lambda_{hi} = \Delta(L)$ and $\lambda_{lo} = \Delta(U)$, so (A3) gives $C(\lambda_{hi}) = L$ and $C(\lambda_{lo}) = U$; at $\Gamma = \Gamma^*$, φ is constant on $[\lambda_{lo}, \lambda_{hi}]$ and every $C \in [L, U]$ is attained, which proves (c). □

Newsvendor interpretation. Consider a risk-neutral planner who chooses C_t to minimise $c_{fix}C_t/\pi_t + E[f_t(d_t, C_t)]$ for a continuous demand distribution. Since $\partial E[f_t]/\partial C_t = (a_t - b)P(d_t > C_t) + hP(d_t < C_t)$, the first-order condition is

$$P(d_t > C_t^*) = \frac{c_{fix}/\pi_t + h}{b - a_t + h} = \frac{k_t}{\kappa_t} \tag{A5}$$

Summing over months, $\Gamma^* = \Sigma_t P(d_t > C_t^*)$ is the expected number of shortage months per year at the risk-neutral optimum.

Proof of Proposition 2. For identical months and $C \in [D^1, D^1/(1-\hat{\rho})]$, states $(0, 0)$ and $(1, 0)$ are in the excess regime, whereas state $(1, 1)$ is in the shortage regime. Hence, $\delta^{10} = g^{10} - g^{00} = (a - h)\hat{d}$ is constant in C, and $\delta^{11} = g^{11} - g^{00}$ has slope $-[(1-\hat{\rho})(b - a) + h]$. Moreover, $\delta^{11} - \delta^{10} - \delta^{01} = (b - a + h)(D^1 - y) > 0$, with $y = \max\{D^0, (1-\hat{\rho})C\} < D^1$, so shortages caused by the demand surge and by the capacity loss compound. Replacing one month in state $(1, 0)$ and another in state $(0, 1)$ by a single month in state $(1, 1)$ uses the same budgets and does not decrease the adversary's value. With $\Gamma_c \leq \Gamma_d \leq T$, the adversary therefore assigns state $(1, 1)$ to Γ_c months and state $(1, 0)$ to $\Gamma_d - \Gamma_c$ further months, and the worst-case value is $T(c_{fix}C/\pi + g^{00}) + \Gamma_c\delta^{11} + (\Gamma_d - \Gamma_c)\delta^{10}$, whose slope in C is $Tk - \Gamma_c[(1-\hat{\rho})(b - a) + h]$. □

References

1. Koç, Ç.; Bektaş, T.; Laporte, G. Decarbonizing road freight transportation: Recent advances and future trends. *J. Oper. Res. Soc.* **2026**, *77*, 43–63. [CrossRef]
2. Müllerklein, D.; Fontaine, P. Resilient transportation network design with disruption uncertainty and lead times. *Eur. J. Oper. Res.* **2025**, *322*, 827–840. [CrossRef]
3. Wei, X.; Wang, M.; Meng, Q. Vulnerability assessment and reduction for intermodal freight transportation networks. *Reliab. Eng. Syst. Saf.* **2024**, *245*, 109969. [CrossRef]
4. Hsu, S.; Balthrop, A.; Pellathy, D.; Kulpa, T.; Ferrada, G.A.; Fu, J. Climate shock impacts on supply chains: The case of the truckload spot market. *Transp. Res. Part E Logist. Transp. Rev.* **2026**, *208*, 104609. [CrossRef]
5. Cordero, R.R.; Feron, S.; Damiani, A.; MacDonell, S.; Carrasco, J.; Pizarro, J.; Karas, C.; Jorquera, J.; Sepulveda, E.; Cabello, F.; et al. Rapid decline in extratropical Andean snow cover driven by the poleward migration of the Southern Hemisphere westerlies. *Sci. Rep.* **2024**, *14*, 26365. [CrossRef] [PubMed]
6. Bertsimas, D.; Sim, M. The price of robustness. *Oper. Res.* **2004**, *52*, 35–53. [CrossRef]
7. Liu, X.; Wang, D.; Yin, Y.; Cheng, T. Robust optimization for the electric vehicle pickup and delivery problem with time windows and uncertain demands. *Comput. Oper. Res.* **2023**, *151*, 106119. [CrossRef]
8. Yu, V.F.; Anh, P.T.; Baldacci, R. A robust optimization approach for the vehicle routing problem with cross-docking under demand uncertainty. *Transp. Res. Part E Logist. Transp. Rev.* **2023**, *173*, 103106. [CrossRef]
9. Crainic, T.G.; Perboli, G.; Rei, W.; Rosano, M.; Lerma, V. Capacity planning with uncertainty on contract fulfillment. *Eur. J. Oper. Res.* **2024**, *314*, 152–175. [CrossRef]
10. Gbadegoye, J.; Camur, M.C.; Li, X. A two-stage stochastic model for road–rail intermodal freight transportation under demand and capacity uncertainty. *Int. J. Transp. Sci. Technol.* **2025**, *23*, 280–300. [CrossRef]
11. Xia, D.; Ma, J.; Azadeh, S.S.; Zhang, W. Data-driven distributionally robust timetabling and dynamic-capacity allocation for automated bus systems with modular vehicles. *Transp. Res. Part C Emerg. Technol.* **2023**, *155*, 104314. [CrossRef]
12. Shekarabi, S.A.H.; Mavi, R.K.; Macau, F.R. An extended robust optimisation approach for sustainable and resilient supply chain network design: A case of perishable products. *Eng. Appl. Artif. Intell.* **2025**, *152*, 110846. [CrossRef]

13. Jin, Z.; Zhang, C.; Ng, K.K.H.; Tian, X.; Wang, H.; Xu, G. Humanitarian logistics planning under uncertain demand and third-party supply capacity: A distributionally robust optimisation approach. *Int. J. Prod. Res.* **2026**, *64*, 3309–3335. [CrossRef]
14. Wu, C.; Zhang, Y.; Xiao, Y.; Mo, W.; Xiao, Y.; Wang, J. Optimization of multimodal paths for oversize and heavyweight cargo under different carbon pricing policies. *Sustainability* **2024**, *16*, 6588. [CrossRef]
15. Zhu, L.; Liu, Z.; Jian, W. Contribution assessment of carbon tax on the reduction of freight corridor carbon emissions through modal shift. *Transp. Res. Rec.* **2023**, *2677*, 167–179. [CrossRef]
16. Solakivi, T.; Ojala, L. Environmental and economic potential of high-capacity trucks. *Case Stud. Transp. Policy* **2024**, *18*, 101284. [CrossRef]
17. Xu, X.; Ou, S.; Peng, T.; Lin, Z.; Ou, X.; Xu, Z.; Gan, M.; Li, D.; Liu, X. The potential role of truck-hailing and operational efficiency improvement in China's road freight decarbonization. *Nat. Commun.* **2026**, *17*, 4714. [CrossRef] [PubMed]
18. Wang, Y.; Xiao, L.; Luo, X.-G. On the optimal solution of a distributionally robust multi-product newsvendor problem. *J. Oper. Res. Soc.* **2023**, *74*, 2578–2592. [CrossRef]
19. Soyster, A.L. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* **1973**, *21*, 1154–1157. [CrossRef]
20. Gabrel, V.; Murat, C.; Thiele, A. Recent advances in robust optimization: An overview. *Eur. J. Oper. Res.* **2014**, *235*, 471–483. [CrossRef]
21. Bertsimas, D.; Thiele, A. A robust optimization approach to inventory theory. *Oper. Res.* **2006**, *54*, 150–168. [CrossRef]
22. Birge, J.R.; Louveaux, F. *Introduction to Stochastic Programming*, 2nd ed.; Springer: New York, NY, USA, 2011. [CrossRef]
23. Bektaş, T.; Laporte, G. The pollution-routing problem. *Transp. Res. Part B Methodol.* **2011**, *45*, 1232–1250. [CrossRef]
24. Demir, E.; Bektaş, T.; Laporte, G. A review of recent research on green road freight transportation. *Eur. J. Oper. Res.* **2014**, *237*, 775–793. [CrossRef]
25. Sheffi, Y.; Rice, J.B. A supply chain view of the resilient enterprise. *MIT Sloan Manag. Rev.* **2005**, *47*, 41–48.
26. Tang, C.S. Perspectives in supply chain risk management. *Int. J. Prod. Econ.* **2006**, *103*, 451–488. [CrossRef]
27. Spinler, S.; Huchzermeier, A. The valuation of options on capacity with cost and demand uncertainty. *Eur. J. Oper. Res.* **2006**, *171*, 915–934. [CrossRef]
28. McKinnon, A. *Decarbonizing Logistics: Distributing Goods in a Low Carbon World*; Kogan Page: London, UK, 2018.
29. Bertsimas, D.; Gupta, V.; Kallus, N. Data-driven robust optimization. *Math. Program.* **2018**, *167*, 235–292. [CrossRef]
30. Mohajerin Esfahani, P.; Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Math. Program.* **2018**, *171*, 115–166. [CrossRef]
31. Ben-Tal, A.; Goryashko, A.; Guslitzer, E.; Nemirovski, A. Adjustable robust solutions of uncertain linear programs. *Math. Program.* **2004**, *99*, 351–376. [CrossRef]
32. Hosseini, S.; Ivanov, D.; Dolgui, A. Review of quantitative methods for supply chain resilience analysis. *Transp. Res. Part E Logist. Transp. Rev.* **2019**, *125*, 285–307. [CrossRef]
33. Notife. Más de 9000 Camiones Varados Tras el Cierre del Paso Internacional Cristo Redentor. September 2023. Available online: https://www.notife.com/mas-de-9000-camiones-varados-tras-el-cierre-del-paso-internacional-cristo-redentor (accessed on 13 November 2026).
34. Diario UNO. El Paso Cristo Redentor Cumplió 20 Días Cerrado y no Hay Certeza Sobre su Reapertura. 3 August 2026. Available online: https://www.diariouno.com.ar/sociedad/el-paso-cristo-redentor-cumplio-20-dias-cerrado-y-no-hay-certeza-su-reapertura-n1581521 (accessed on 13 November 2026).
35. Reuters. Key Border Crossing Between Argentina and Chile Reopens After Weather Closure. 18 August 2026. Available online: https://www.internazionale.it/ultime-notizie-reuters/2026/08/18/key-border-crossing-between-argentina-and-chile-reopens-after-weather-closure (accessed on 13 November 2026).

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.