Truck electrification asset management under deep uncertainty

Raoul Nicolodi*[1] and Jeppe Rich[1]

[1]Transport Policy Section, Department of Technology, Management and Economics, Technical University of Denmark

Short summary

Truck operators replacing diesel fleets with electric alternatives face a sequential, irreversible investment problem whose structure makes a single optimal rollout plan inadequate for strategy. This paper develops a framework that distinguishes two sources of decision variance: *within-scenario flexibility*, the room to resequence or defer commitments inside a given future, and *cross-scenario sensitivity*, the degree to which the preferred strategy changes across futures. We show theoretically, using real options and mean-preserving spread arguments, that flexible operators with independent adoption choices should exhibit more within-scenario latitude, while inflexible operators constrained by bundled commitments should face more cross-scenario exposure. A Probabilistic Mixed-Integer Programming (PMIP) formulation makes the decomposition observable. Stylized experiments confirm the prediction: the inflexible operator's cross-scenario variance is twelve times larger than the flexible operator's, with virtually no within-scenario latitude, while the flexible operator retains meaningful room for manoeuvre within each future.
Keywords: deep uncertainty, operations research applications, stochastic programming, technology adoption timing, truck electrification, variance decomposition.

1 Introduction

Truck operators do not face a simple yes-or-no electrification decision. They face a finite-horizon optimisation problem over assets, infrastructure, and timing: which segments to electrify, when to commit, and how much optionality to preserve while technology, regulation, and demand evolve under uncertainty. Some vehicle classes can move early with current battery technology. Others only become attractive if range, charging infrastructure, or policy support improve further. The planning problem is therefore sequential, heterogeneous, and irreversible Feng & Figliozzi (2022); International Energy Agency (2025).

This creates two distinct sources of variance in the set of good solutions. Inside one future, several rollout paths may be close to optimal: assets can be resequenced, timing shifted, or some commitments deferred without much cost. Across different futures, the preferred rollout can change sharply as the technology frontier, carbon costs, or customer pressure develop differently. We call these dimensions *within-scenario flexibility* and *cross-scenario sensitivity*, respectively. The first captures operational latitude; the second captures strategic exposure.

Under deep uncertainty, an operator's relevant question is whether the future mainly creates room for manoeuvre within a robust strategic direction, or whether it forces fundamentally different asset strategies across plausible futures. The operational object is therefore a distribution of near-optimal plans, not a single incumbent Birge & Louveaux (2011); Shapiro et al. (2009); Marchau et al. (2019).

The operator's deployment architecture shapes which type of variance dominates. Flexible operators can stage adoption by segment, depot, or procurement cycle and preserve the option to learn before committing the rest of the fleet. Inflexible operators face bundled commitments, lumpy infrastructure investments, or contractual requirements that make partial adoption impractical. The same uncertainty can therefore create optionality in one setting and commitment risk in another Rothschild & Stiglitz (1970); Dixit & Pindyck (1994); McDonald & Siegel (1986).

This distinction carries immediate practical weight for two types of decision-makers. For a multinational logistics operator running fleets across several European countries under one corporate capital budget, the variance decomposition is a capital-allocation tool. Each country has its own charging infrastructure gap, road-pricing scheme, and network structure. Some branches will be inherently flexible; others face local coordination requirements that create bundled exposure. The relevant question is not only where returns are highest, but where cross-scenario exposure is manageable and where within-scenario latitude allows the operator to adapt as conditions evolve. A branch with high potential return but high cross-scenario sensitivity may be a worse capital destination than a branch with moderate return and genuine room for manoeuvre.
For policymakers who want the transition to extend beyond large operators with modular networks and strong balance sheets, these structural differences matter as well. Instruments calibrated to average subsidy intensity can inadvertently favour already-flexible operators. Infrastructure access, staged compliance windows, and instruments that reduce bundled commitments can matter as much as the nominal incentive for firms that face coordination-driven inflexibility.
The paper makes three linked contributions. First, it recasts truck electrification as an asset-management problem under deep uncertainty, using real options and Bellman-style reasoning to motivate why a distribution of good solutions is the right object for planning. Second, it derives a theoretical variance decomposition: flexible architectures should generate more within-scenario latitude; inflexible architectures more cross-scenario exposure. Third, it operationalises the decomposition using a Probabilistic Mixed-Integer Programming (PMIP) formulation, which acts as a distributed solver over near-optimal plans, and validates the theoretical predictions on a stylized but structured application.

2 Methodology

In its most general form, the operator's problem is to choose a policy π from a feasible set Π to maximize expected discounted value across uncertain futures ω:

$$\max_{\pi \in \Pi} \mathbb{E}_{\omega} \left[V(\pi, \omega) \right], \tag{1}$$

while ensuring stability of π under various conditions.
We first approach the problem from a high-level, stylized continuous-time approach, giving us theoretical predictions on the structure of the solution space. We then solve the discretized, multi-segment problem using stochastic programming, combined with a sampling approach around the pointwise optimum, to describe the solution space and infer metrics that guide strategic decision-making.

A stylized continuous-time model
We adopt an approach grounded on the dynamic programming principle, enhanced by structural properties of the operator. The technology quality θ_t evolves as an arithmetic Brownian motion

$$d\theta_t = \mu dt + \sigma dW_t, \quad \theta_0 > 0, \tag{2}$$

over a finite time-horizon $[0, T]$, where μ is the drift (average improvement rate) and σ is the volatility (uncertainty).

The operator manages a set of strategic deployment units indexed $i = 1, \ldots, U$, which in the fleet-transition context represent different parts of the network with different specifications. Each unit i has a *hurdle level* h_i (minimum technology quality for feasible adoption) and a *unit contribution rate* δ_i (net benefit per unit of quality above the hurdle). Deploying unit i at quality $\theta \geq h_i$ yields per-period benefit

$$b_i(\theta) = \delta_i \left(\theta - h_i \right), \tag{3}$$

which may be negative for units expensive to operate at the adoption frontier. A *flexible* operator is able to deploy units individually. An *inflexible* operator must deploy some or all units jointly, for instance due to fleet composition, network design, or contractual agreements. The different deployment architectures lead to different total benefit functions $B_{\text{flex}}(\theta)$ and $B_{\text{inflex}}(\theta)$, the former being convex and the latter concave in the region described by hurdles with negative contributions.

The operator's payoff from adopting at time τ is:

$$J(\tau) = -L(\tau) + e^{-r\tau}\left[B_X(\theta_\tau)A(r, T-\tau) - C(\theta_\tau, p_\tau)\right] + e^{-rT}\mathbb{E}_\tau[S(\theta_\tau, \theta_T)], \tag{4}$$

where $L(\tau)$ is the discounted cost of delaying adoption, the benefit B_X ($X \in \{\text{flex, inflex}\}$) is discounted over the remaining horizon by a unit annuity factor $A(r, T-\tau)$, $C(\theta_\tau, p_\tau)$ is the capital cost of transitioning, and $S(\theta_\tau, \theta_T)$ is the terminal (scrap) value of the adopted technology. The value function $F(\theta, t) = \sup_{\tau\in[t,T]} \mathbb{E}_t[J(\tau)]$ satisfies the Hamilton-Jacobi-Bellman equation, grounding the problem in established theoretical territory.
Real-options theory and the mean-preserving spread logic of Rothschild and Stiglitz then lead to the following predictions:

- **Innovation speed:** Higher drift μ raises the option value of waiting, as faster improvements increase the continuation payoff more than the immediate payoff.
- **Mean-preserving spread:** Increasing volatility σ has opposite effects depending on architecture. For flexible operators, higher volatility increases option value; for inflexible operators, it reduces the continuation payoff in the concave region.
- **Level and sensitivity effects:** Flexible operators adopt earlier (lower effective hurdle) and face less scenario-dependent timing. Inflexible operators face higher adoption thresholds and greater strategic exposure: their timing is more sensitive to which future materialises.

Beyond the payoff comparison, a second important dimension is the variability of near-optimal plans. For each decision coordinate z_k (e.g. the adoption year of unit i), the law of total variance gives:

$$\underbrace{\mathrm{Var}(z_k)}_{\text{total}} = \underbrace{\mathbb{E}_s[\mathrm{Var}(z_k \mid s)]}_{\mathrm{VW}_k\text{: within-scenario flexibility}} + \underbrace{\mathrm{Var}_s(\mathbb{E}[z_k \mid s])}_{\mathrm{VC}_k\text{: cross-scenario sensitivity}}, \tag{5}$$

with s ranging over plausible futures. VW_k is operational latitude: room to choose among near-optimal timing sequences inside a given future. VC_k is strategic exposure: variation in the preferred plan as the uncertain future resolves. The two components are not symmetric. Within-scenario latitude is a form of optionality the operator can exploit without cost. Cross-scenario exposure is risk: the operator must commit to substantially different plans depending on which future materialises.

Proposition 1 (Variance decomposition ordering)**.** *For each deployment coordinate k:*

1. *Flexible operators have weakly larger expected within-scenario width* VW_k *than inflexible operators.*
2. *Inflexible operators have weakly larger cross-scenario sensitivity* VC_k *than flexible operators.*

The intuition follows from the curvature argument above. A flexible operator's independent adoption choices allow near-optimal substitution within a scenario: small cost perturbations shift the preferred timing without invalidating the whole plan. An inflexible

operator's bundle constraint forces adoption decisions to covary, so changes across scenarios tend to be all-or-nothing.

From theory to practice: PMIP and Gibbs sampling

While the continuous-time model provides theoretical guidance, it is intractable for multi-segment fleets with finite scenario sets. We operationalise the framework using a Probabilistic Mixed-Integer Programming (PMIP) formulation. For each scenario s, we first solve a strategic LP to maximize the expected discounted payoff over discretized time $[0, \dots, T]$. Let $x_{i,t} \in \{0, 1\}$ indicate whether unit i has adopted by epoch t and $y_j \in \{0, 1\}$ indicate whether infrastructure option j is commissioned.

We solve the strategic LP

$$\max_{x,y} \sum_{t=1}^{T} \sum_{i=1}^{R} c_{i,t}^{(s)} x_{i,t} - \sum_{j=1}^{M} G_j y_j, \tag{6}$$

subject to:

$$x_{i,t} \le x_{i,t+1} \quad \forall\, i,\ t, \qquad \text{(irreversibility)} \tag{7}$$

$$\sum_{i=1}^{R} K_{i,t}^{(s)} (x_{i,t} - x_{i,t-1}) + \sum_{j=1}^{M} \Gamma_{j,t} y_j \le B_t \qquad \forall\, t, \qquad \text{(budget constraint)} \tag{8}$$

$$x_{i,t} \le y_j \qquad \forall\, (i, j) \in \mathcal{A},\ t, \quad \text{(infrastructure coupling)} \tag{9}$$

$$x_{i,t} = 0 \qquad \text{if } \theta_t^{(s)} < h_i, \qquad \text{(technology threshold)} \tag{10}$$

where $c_{i,t}^{(s)}$ is the net discounted contribution of unit i at time t in scenario s, G_j is the fixed commissioning cost of option j, $K_{i,t}^{(s)}$ is the scenario-dependent capital requirement, $\Gamma_{j,t}$ is the per-period monetary load of option j, and $\mathcal{A}$ denotes infrastructure-to-unit coupling.

To recover a distribution of near-optimal plans, we augment the LP objective with entropy regularization:

$$\max_{z \in \mathcal{F}^{(s)}} \Psi_\beta^{(s)}(z) = f^{(s)}(z) + \beta H(z), \tag{11}$$

where $z = (x, y)$, $f^{(s)}(z)$ are the coefficients from (6), $\mathcal{F}^{(s)}$ is the feasible space defined by (7)–(10), the temperature β controls the mass spread over near-optimal solutions, and $H(z)$ is the discrete entropy. The regularised objective induces a Gibbs distribution $\mathbb{P}_\beta^{(s)}(z) \propto \exp(f^{(s)}(z)/\beta)$ over feasible transition plans. We sample from this distribution via a Metropolis-Hastings algorithm, obtaining a collection of feasible plans approximating the near-optimal set $\mathcal{F}_\varepsilon^{(s)}$ for each scenario.

Empirical variance decomposition

The PMIP output provides the empirical counterpart of the theoretical predictions. For each decision coordinate k, we compute:

$$\widehat{\mathrm{VW}}_k = \frac{1}{S} \sum_{s=1}^{S} \widehat{\mathrm{Var}}\left(z_k \mid s\right), \tag{12}$$

$$\widehat{\mathrm{VC}}_k = \widehat{\mathrm{Var}}_s\left(\bar{z}_k^{(s)}\right), \tag{13}$$

where $\bar{z}_k^{(s)}$ is the average PMIP draw for coordinate k in scenario s. Flexible operators are expected to display large $\widehat{\mathrm{VW}}_k$ (many near-optimal strategies for a given future) and moderate $\widehat{\mathrm{VC}}_k$. Inflexible operators should display small $\widehat{\mathrm{VW}}_k$ (the bundle leaves little room for near-optimal variation) and large $\widehat{\mathrm{VC}}_k$ (whether and when the bundle clears viability varies sharply across scenarios).

A key contrast to *Sample Average Approximation* (SAA) is worth stating explicitly. SAA computes a single solution that maximises the expected objective over scenario realisations. That gives a point estimate: the plan that, on average, performs well. The PMIP formulation instead recovers the distribution of near-optimal plans, which is precisely the object that contains the VW/VC decomposition. SAA averages away the cross-scenario spread; PMIP makes it visible. This is the computational sense in which the two approaches answer different questions.

3 Results and discussion

We validate the framework using a stylized truck-electrification model with four deployment units across a 3×3 scenario grid defined by two dimensions: *technology* (slow, mid, or fast improvement) and *policy* (weak, mid, or strong carbon pressure). The four units represent the segment heterogeneity typical of a mixed freight operator:

- *Core urban service* (unit 1): Low hurdle, high base contribution — early and consistently attractive.
- *Regional network* (unit 2): Medium hurdle, medium base contribution — viable once technology reaches a moderate threshold.
- *Heavy operations* (unit 3): High hurdle, small positive contribution — requires strong combined progress to be worth committing.
- *Frontier long-range* (unit 4): Very high hurdle, negative contribution rate ($\delta_4 < 0$) — remains commercially unattractive until deep into the technology frontier.

The *flexible operator* adopts each unit independently, subject to the budget constraint. The *inflexible operator* is constrained to transition units 3 and 4 together (bundled infrastructure), which is the structural device that creates the concave region of B_{inflex}.

Incumbent adoption patterns

Figure 1 shows the optimal adoption years for both operator types across the 3×3 scenario grid.

For the flexible operator, the pattern is orderly and robust. Core urban service is adopted in Year 0 in all nine scenarios: its hurdle is low and its contribution is unconditionally positive. Regional network follows in Years 1–2, with timing shifting by about one year between weak and strong policy environments. Heavy operations is adopted in Years 2–4 in all scenarios with mid or high technology readiness, but defers in the slow-technology, weak-policy corner. Frontier long-range is never adopted in any scenario: the negative contribution rate means adoption would reduce total payoff at any technology level reachable in the planning horizon. Taken together, the flexible operator adopts three of four units in most scenarios, with timing variation that is modest and systematic.

The inflexible operator presents a starkly different picture. Core urban service and Regional network behave similarly to the flexible case, since they are not part of the bundle. Heavy operations and Frontier long-range, however, can only be activated together. The bundle is infeasible in low-technology scenarios because the technology threshold for Frontier long-range is never crossed. In mid-technology or mid-policy scenarios, it is technically feasible

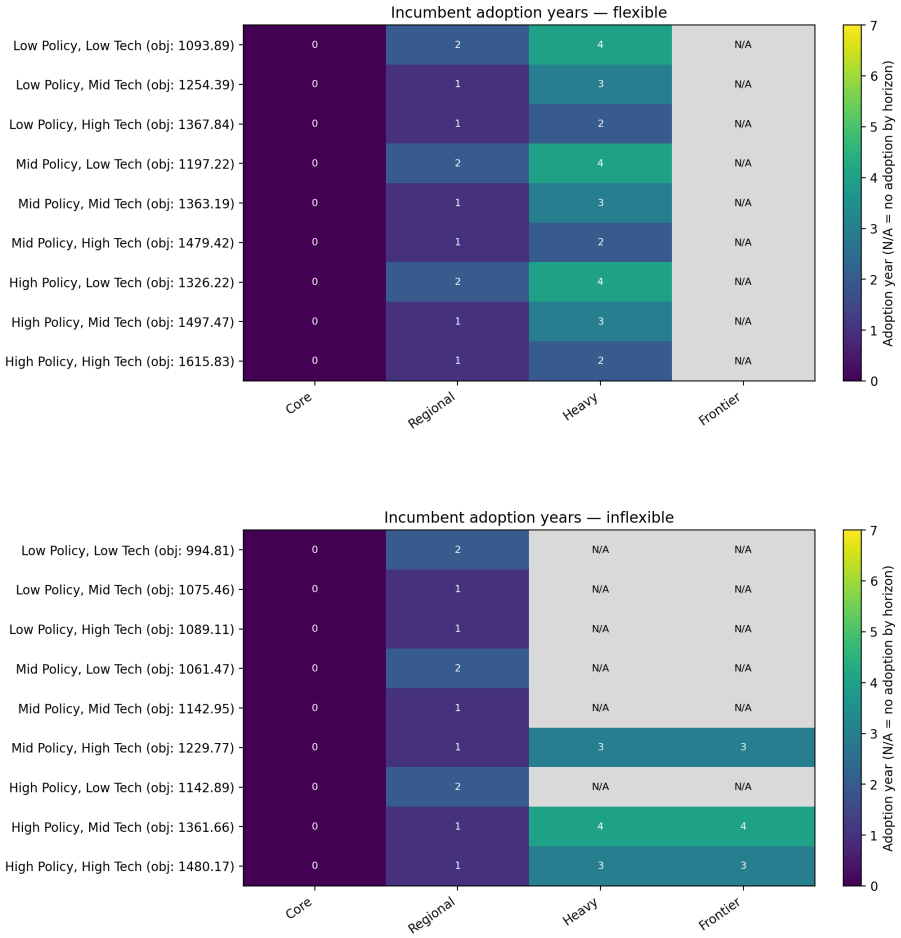


Figure 1: Incumbent adoption years for flexible (top) and inflexible (bottom) operators across 3×3 technology-policy scenarios. Colour indicates adoption year; white/grey cells denote non-adoption within the planning horizon.

but economically unattractive, because the negative contribution rate of Frontier long-range pulls the bundle payoff below the adoption threshold. Adoption occurs only in scenarios with both sufficient technology progress and strong policy support. The result is a binary pattern: no adoption in the lower-left of the scenario grid, adoption in Years 3–5 in the upper-right. This is exactly the cross-scenario sensitivity predicted by Proposition 1(ii): whether and when the bundle activates is highly scenario-dependent, while the flexible operator's decisions are comparatively stable.

Objective values and the cost of inflexibility

The flexible operator achieves higher discounted payoffs across all scenarios, ranging from 1,093 to 1,616 objective units. The inflexible operator achieves 995–1,480, with the largest gaps in scenarios where the bundle is either delayed or infeasible. The gap is not symmetric: in favourable scenarios (high technology, strong policy) the two operators converge because the bundle eventually clears and adoption timing aligns. In unfavourable scenarios the gap widens, because the flexible operator has already staged partial adoption while the inflexible operator waits for conditions that may not arrive within the horizon.

Variance decomposition: quantifying the theoretical predictions

Table 1 reports the empirical variance decomposition from the PMIP runs at $\varepsilon = 0.05$. The results confirm Proposition 1 sharply. For the inflexible operator, essentially all variance is cross-scenario: the cross share is 100% and $\widehat{\mathrm{VC}} = 14.64$. For the flexible operator,

Table 1: Empirical variance decomposition by operator regime ($\varepsilon = 0.05$, $\tau = 18$, $S = 9$ scenarios). $\widehat{\mathrm{VW}}$: within-scenario flexibility; $\widehat{\mathrm{VC}}$: cross-scenario sensitivity; cross share: $\widehat{\mathrm{VC}}/(\widehat{\mathrm{VW}} + \widehat{\mathrm{VC}})$.

Regime	$\widehat{\mathrm{VW}}$	$\widehat{\mathrm{VC}}$	Cross share
Flexible	0.83	1.22	60%
Inflexible	0.00	14.64	100%

the cross share drops to 60% and $\widehat{\mathrm{VC}} = 1.22$, making the inflexible operator's cross-scenario variance twelve times larger in absolute terms. The flexible operator also retains meaningful within-scenario latitude ($\widehat{\mathrm{VW}} = 0.83$), reflecting the room to stage commitments or shift timing by a year without substantial loss. The inflexible operator has zero within-scenario latitude: once the bundle constraint is active, the adoption plan is fully determined by the scenario.

These numbers have a practical interpretation. A flexible operator entering the planning problem knows that roughly 40% of its total decision variance ($\widehat{\mathrm{VW}}/\text{total} \approx 40\%$) reflects room for manoeuvre that it can exploit operationally — by adjusting procurement cycles, piloting segments, or staging depot upgrades without committing the full capital budget. An inflexible operator faces no such latitude. Its entire variance is driven by which future materialises. That creates two concrete risks: the operator either commits capital at a time that turns out to be wrong for the scenario that materialises, or it waits so long that it misses the adoption window entirely in scenarios where the bundle remains economically marginal.

What PMIP reveals that SAA conceals

The variance decomposition in Table 1 is invisible to a standard SAA approach. SAA would compute a single weighted-average incumbent plan across the nine scenarios. For the inflexible operator, that plan would specify some expected adoption year for the bundle, averaging over the scenarios where the bundle activates and those where it does not. The result would be an adoption recommendation that is rarely optimal in any specific scenario: too early for the low-technology futures, too late (or non-adoption) for the high-technology ones. The cross-scenario spread — which is the operative measure of strategic exposure — is averaged away by construction.

The PMIP distribution makes that spread observable and quantifiable. The near-optimal set $\mathcal{F}_\varepsilon^{(s)}$ differs across scenarios: in favourable futures, it contains many plans with similar payoffs (within-scenario latitude); in unfavourable futures, it may contain only one or two feasible plans near the optimum, all of which avoid the bundle (no latitude, extreme cross-scenario contrast). This is the information that an operator needs to decide whether inflexibility is a structural liability under deep uncertainty, and which instrument — infrastructure access, phased compliance, or bundle-breaking incentives — would most directly reduce it.

4 Conclusions

This paper has developed a framework for truck electrification as an asset-management problem under deep uncertainty, in which the relevant planning object is a distribution of near-optimal strategies rather than a single optimal incumbent. The central finding is that deployment architecture — flexible versus inflexible — determines which type of decision variance dominates. Real-options logic and mean-preserving spread arguments predict the ordering; PMIP-based sampling makes it observable; and the stylized application confirms it starkly. The inflexible operator's cross-scenario variance is twelve times larger than the

flexible operator's, with virtually no within-scenario latitude, while the flexible operator retains meaningful room for manoeuvre within each future.

The methodological contribution is the pairing of a theoretically motivated variance decomposition with a computational engine that delivers empirical estimates of both components. PMIP goes beyond what SAA can offer: rather than returning one plan that is best on average, it recovers the distribution of near-optimal plans and thereby makes the scenario-contingent spread directly measurable. This connects to a broader result in stochastic integer programming: SAA provides a lower bound on the true optimum, while the Boltzmann-weighted PMIP distribution provides a complementary upper bound, together framing the solution quality problem from both sides.

Several extensions follow naturally from the framework. The most immediate is to replace the LP benchmark with a full PMIP sweep across the uncertainty regime, connecting the within-scenario and cross-scenario components across all nine scenarios simultaneously. This would allow computation of the Value of Stochastic Solution (VSS) for inflexible operators — how much is lost by treating cross-scenario exposure as if it were within-scenario latitude — and identification of the specific scenario combinations where bundle lock-in causes the greatest loss. A second extension moves toward a multi-branch operator with a shared corporate budget, where Proposition 1 becomes a capital-allocation guide: branches with high within-scenario latitude are relatively safe capital destinations, while branches with high cross-scenario sensitivity warrant smaller commitments or instrument design that reduces bundling. A third direction is to use the near-optimal neighbourhood structure from PMIP to test robustness explicitly: identifying plans that sacrifice a small fraction of expected payoff in exchange for substantially reduced cross-scenario exposure is directly relevant for risk-averse operators and for policymakers designing staged compliance frameworks.

Acknowledgements

The authors thank the Transport Policy Section at DTU for support.

References

Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming* (2nd ed.). Springer.

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty.* Princeton: Princeton University Press.

Feng, W., & Figliozzi, M. (2022). Electrification of urban freight transport: Policy, technology, and logistics. *Transportation Research Part C: Emerging Technologies*, *134*, 103463. (Focuses on the logistical constraints of fleet transition—useful for your "Network Flexibility" sections.)

International Energy Agency. (2025). *Global ev outlook 2025: Battery innovation and mineral supplies* (Tech. Rep.). Paris: IEA. (Provides the most recent data on battery density gains and price drops (your θ_t process).)

Marchau, V. A. W. J., Walker, W. E., Bloemen, P. J. T. M., & Popper, S. W. (2019). *Decision making under deep uncertainty: From theory to practice.* Springer.

McDonald, R., & Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, *101*(4), 707–728.

Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. a definition. *Journal of Economic Theory*, *2*(3), 225–243. doi: 10.1016/0022-0531(70)90038-4

Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2009). *Lectures on stochastic programming: Modeling and theory.* SIAM.