Stochastic Supply Planning for Wind Nacelles with Profit–Emission Trade-offs

ISBN: 978-1-943295-26-5

Thangamani G
Indian Institute of Management Kozhikode
(gtmani@iimk.ac.in)

Wind-energy manufacturers are adopting circular business models integrating new-build and remanufactured nacelles. This study formulates a multi-period stochastic MILP program for closed-loop supply planning that allocates capacity between manufacturing and remanufacturing under uncertain demand and core returns. Uncertainty is addressed via Sample Average Approximation, and sustainability is internalized through an ε-constraint on expected emissions to generate a profit–emissions Pareto frontier. The model, implemented using open-source Python and Excel tools, reveals how emission caps shift production toward remanufacturing and provides managerial insights on capacity allocation, channel interactions, and the marginal cost of carbon abatement.

Keywords: Closed-loop Supply Chain, Wind-Turbine Nacelles, Remanufacturing Optimization, Sample Average Approximation (SAA), Circular economy, Profit–Emissions Trade-Off

1. Introduction

The wind sector faces the dual imperative of scaling capacity and reducing lifecycle emissions. In parallel, many OEMs are experimenting with circular models—such as buybacks and remanufacturing of nacelles—to reduce embodied carbon and create new revenue streams from returned cores. However, planning such closed-loop systems requires simultaneously coordinating forward and reverse flows under uncertainty in demand, return availability and quality yields, and grade-dependent sales channels. Classic newsvendor or deterministic network design abstractions are too coarse to capture these intertwined effects. Instead, we adopt a stochastic, multi-period planning model that chooses production, remanufacturing throughput, inventory, and sales quantities per scenario, with an explicit environmental objective captured through an emissions cap. The resulting trade-off between financial performance and decarbonization is summarized by a profit–emissions frontier (Pareto set), which is particularly useful for communicating to executives and policy stakeholders.

Methodologically, our approach builds on fundamental results in stochastic programming and SAA, which approximate intractable expectations by sample averages and deliver asymptotic guarantees on optimality and consistency (Birge & Louveaux, 2011; Kleywegt, Shapiro, & Homem-de-Mello, 2002). From a supply-chain perspective, the paper aligns with the literature on closed-loop supply chains (CLSCs), remanufacturing economics, and channel design to mitigate cannibalization between remanufactured and new products (Guide & Van Wassenhove, 2009; Atasu, Sarvary, & Van Wassenhove, 2008). We operationalize these ideas in an implementable optimization template that reads scenarios from an Excel workbook and produces a frontier plot and optimized decision tables, making the method accessible to practitioners and scholars.

2. Literature Review

Closed-loop supply chains focus on recovering value from returned products through reuse, remanufacturing, and recycling, integrating reverse logistics with forward planning. Foundational reviews have tracked the field's evolution from engineering-focused case studies to quantitative decision models that weigh profitability against environmental outcomes and market impacts (Guide & Van Wassenhove, 2009; Govindan, Soleimani, & Kannan, 2015). In wind and other capital goods, remanufacturing quality grading and warranty considerations complicate channel strategy and pricing. Moreover, capacity coupling between new-build and reman lines can induce nontrivial trade-offs when emissions constraints are active (Ponte, Alvarez, & Atasu, 2021).

On the uncertainty side, stochastic programming and SAA provide a principled route to decision making under risk. The SAA framework replaces the true expected objective with a sample average over Monte Carlo scenarios, solves the resulting deterministic program, and, when repeated with independent batches, enables statistical validation and optimality-gap estimation (Kleywegt et al., 2002; Verweij, Ahmed, Kleywegt, Nemhauser, & Shapiro, 2003; Birge & Louveaux, 2011). For multi-objective settings, the ε-constraint method is a standard scalarization that enforces a hard bound on one objective (e.g., emissions) while optimizing another (profit), and sweeping ε yields a well-populated Pareto frontier suitable for managerial interpretation (Mavrotas, 2009). The broader circular-economy context underscores why such trade-offs matter in practice (Ellen MacArthur Foundation, 2013).

3. Model and Methods

Notation and Decision Variables are given below.

Sets and indices:
Periods$t \in \mathcal{T} = \{1, \dots, T\}$;
Scenarios $\omega \in \Omega$ with probabilities p_ω;
Customer segments$k \in \mathcal{K}$;
grades $g \in \mathcal{G} = \{\mathrm{N}, \mathrm{A}, \mathrm{B}\}$ (new, reman Grade-A, reman Grade-B).

Parameters (data):
Demand $D_{kt\omega}$;
Available cores for buyback$\bar{B}_{t\omega}$;
Reman yields $y^A_{t\omega}$, $y^B_{t\omega}$, $y^S_{t\omega}$ (with$y^A + y^B + y^S = 1$);
Prices P_{gkt} (exogenous);
Unit costsc^N_t, c^{bb}_t, c^{rm}_t;
Holding costs h_{gt};
Unmet-demand penaltyπ_{kt};
Capacities$\bar{M}^N_t$, $\bar{M}^{\mathrm{rm}}_t$;
Emission factorse^N, e^A, e^B;
Initial inventories I^N_0, I^A_0, I^B_0 and initial coresC_0.

Decision variables per(t, ω):
New build $x^N_{t\omega}$;
Cores bought $b_{t\omega}$;
reman throughput $r_{t\omega}$;
Reman outputs by grade $q^A_{t\omega} = y^A_{t\omega}\, r_{t\omega}$, $q^B_{t\omega} = y^B_{t\omega}\, r_{t\omega}$;
End-of-period inventories$I^N_{t\omega}, I^A_{t\omega}, I^B_{t\omega}$;
Core inventory$C_{t\omega}$;
Sales$s^N_{kt\omega}, s^A_{kt\omega}, s^B_{kt\omega}$;
Unmet demand$u_{kt\omega}$.

Flow Balance and Capacity Constraints:
Core balance and availability

$$C_{t\omega} = C_{t-1,\omega} + b_{t\omega} - r_{t\omega}, \qquad b_{t\omega} \le \bar{B}_{t\omega}.$$

Finished-goods inventories (for $g \in \{\mathrm{N}, \mathrm{A}, \mathrm{B}\}$):

$$I^N_{t\omega} = I^N_{t-1,\omega} + x^N_{t\omega} - \sum_{k \in \mathcal{K}} s^N_{kt\omega}\,,$$
$$I^A_{t\omega} = I^A_{t-1,\omega} + q^A_{t\omega} - \sum_{k \in \mathcal{K}} s^A_{kt\omega}\,,$$
$$I^B_{t\omega} = I^B_{t-1,\omega} + q^B_{t\omega} - \sum_{k \in \mathcal{K}} s^B_{kt\omega}\,.$$

Demand satisfaction by segment

$$\sum_{g \in \{\mathrm{N},\mathrm{A},\mathrm{B}\}} s^g_{kt\omega} + u_{kt\omega} = D_{kt\omega}, \quad \forall k \in \mathcal{K}.$$

Production Capacities

$$x^N_{t\omega} \le \bar{M}^N_t, \qquad r_{t\omega} \le \bar{M}^{\mathrm{rm}}_t.$$

Optional channel exposure limits (to mitigate cannibalization):

$$s^A_{kt\omega} \le \alpha^A_k D_{kt\omega}, s^B_{kt\omega} \le \alpha^B_k D_{kt\omega}.$$

Objective and ε-Constraint (Profit–Emissions Frontier)

We maximize expected profit across scenarios while enforcing a bound on expected emissions (the ε-constraint). The expected-profit objective is

$$\max \sum_{\omega\in\Omega} p_\omega \sum_{t\in\mathcal{T}} [\sum_{k\in\mathcal{K}} \sum_{g\in\{\mathrm{N,A,B}\}} P_{gkt}\, s^g_{kt\omega} - \left(c^N_t x^N_{t\omega} + c^{\mathrm{bb}}_t b_{t\omega} + c^{\mathrm{rm}}_t r_{t\omega}\right) - \sum_{g\in\{\mathrm{N,A,B}\}} h_{gt}\, I^g_{t\omega} - \sum_{k\in\mathcal{K}} \pi_{kt}\, u_{kt\omega}].$$

Expected cradle-to-gate emissions are constrained by

$$\sum_{\omega\in\Omega} p_\omega \sum_{t\in\mathcal{T}} (e^N x^N_{t\omega} + e^A q^A_{t\omega} + e^B q^B_{t\omega}) \;\le\; \varepsilon.$$

Sweeping ε over a grid yields the Pareto-efficient profit–emissions frontier. An alternative scalarization uses an internal carbon price λ in the objective

$$\max\ \mathbb{E}[\text{Revenue} - \text{Cost}] - \lambda \cdot \mathbb{E}[\text{Emissions}];$$

Under convexity, varying λ recovers the same efficient set (Mavrotas, 2009).

Uncertainty via Sample Average Approximation (SAA)

Let $f(x,\xi)$ denote the one-period profit contribution given decision vector x and random vector ξ (demands, core availability, yields). The true objective is the expectation

$$F(x) = \mathbb{E}[f(x,\xi)].$$

SAA draws an i.i.d. sample $\{\xi^1, \dots, \xi^N\}$ and solves the sample-average problem

$$\hat{F}_N(x) = \frac{1}{N}\sum_{i=1}^{N} f\,(x,\xi^i).$$

As$N \to \infty$, SAA solutions converge to true optima under mild regularity, and optimal values satisfy a law of large numbers; batching and out-of-sample evaluation provide empirical confidence on solution quality (Birge & Louveaux, 2011; Kleywegt et al., 2002; Verweij et al., 2003).

3.1 Data and Experimental Design

We construct a realistic, anonymized dataset to illustrate the method: three periods, two customer segments (premium and utility), and three grades (N, A, B). Prices are exogenous by grade and segment (higher in premium), and costs include new-build, core buyback, reman processing, and inventory holding. Scenario uncertainty spans segment demands, available cores for buyback, and reman yield shares for A/B/Scrap that sum to one. Capacities differ by period. Emission factors reflect cradle-to-gate accounting (highest for new build, lower for reman grades). All inputs are provided in an Excel workbook with sheets for periods, segments, grades, prices, costs, holding, penalties, alpha-limits, capacities, emissions, initial inventories, scenarios, demand, cores, yields, and a grid of emissions caps.

For numerical illustration, we sweep seven cap values $\varepsilon \in \{400{,}500{,}600{,}700{,}800{,}900{,}1000\}$. For each ε, the model is solved with the open-source CBC solver via PuLP, and expected profit/emissions are recorded to trace the frontier. While the dataset is synthetic, parameter magnitudes are selected to be plausible for heavy equipment and aligned with CLSC economics reported in the literature (Guide & Van Wassenhove, 2009; Govindan et al., 2015).

4. Results

Figure 1 and Table 1 show expected profit versus expected emissions across ε. The curve is concave and decreasing, revealing efficient trade-offs. At tight caps ($\varepsilon \le 600$), the solution prioritizes reman throughput, uses premium-channel exposure limits to protect new-product margins, and tolerates some unmet demand in the utility segment if penalties are modest. As ε increases, new-build output ramps up, inventories are used more aggressively to satisfy premium orders, and core purchases decouple from reman capacity as scrap rates bind. Kinks on the frontier correspond to binding capacity constraints or alpha limits, which are managerially interpretable milestones of decarbonization.

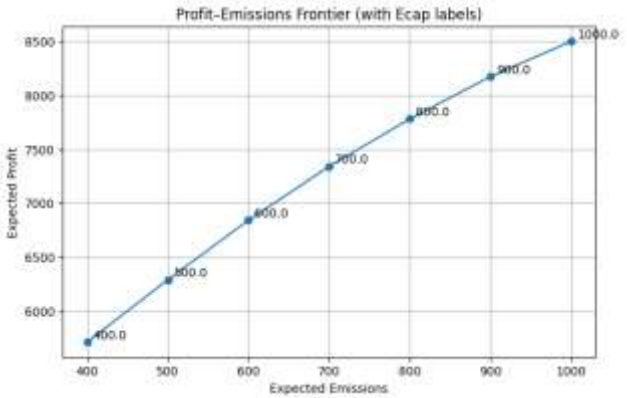

Figure 1 *Profit-Emissions Frontier*

Table 1 *Profit–Emissions Frontier*

Emissions Cap	Status	Expected Profit	Expected Emissions
400	Optimal	5712.464985	400.0000034
500	Optimal	6290.186317	500.0000029
600	Optimal	6840.186328	600.0000054
700	Optimal	7339.06324	700.0000054
800	Optimal	7779.063242	800.0000059
900	Optimal	8171.037703	900.0000054
1000	Optimal	8496.55894	1000.000005

The production mix graph (Figure 2) shows how the optimized supply chain shifts strategically from new manufacturing to remanufacturing across the planning horizon. In the early period, the model relies more on new production to meet demand, but this quickly declines as returned cores become available. Remanufacturing remains consistently high across all periods, indicating that the optimization prioritizes cost efficiency and emission reduction by leveraging recovered components whenever possible. This pattern reflects the closed-loop structure of the problem, where returns in earlier periods enhance remanufacturing capacity in subsequent periods.

The sales and emissions graphs together (Figure 3 and 4) illustrate how this production strategy affects both market output and sustainability performance. Sales by grade reveal that high-quality Grade N units dominate initially, while Grade A and B remanufactured units take on a greater role later, ensuring demand fulfilment while minimizing penalties. The emissions breakdown further confirms the environmental advantage of remanufacturing: although new production contributes the largest single share of emissions, Grades A and B together represent a substantial portion of output at significantly lower emission intensity. Overall, the optimized plan achieves a balanced trade-off between profitability, demand satisfaction, and compliance with emission constraints

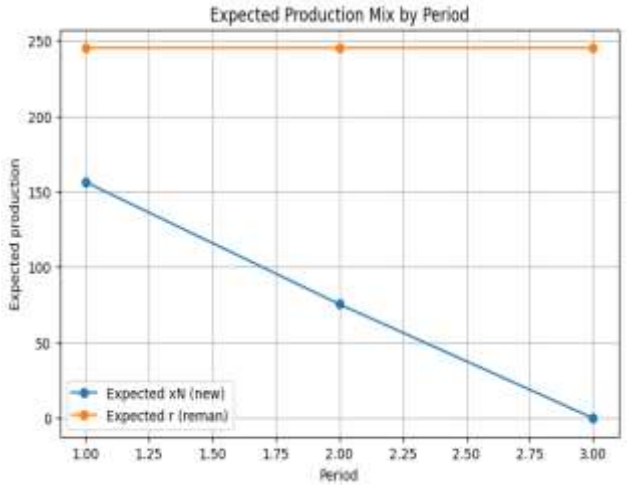

Figure 2 *Production Mix Graph*

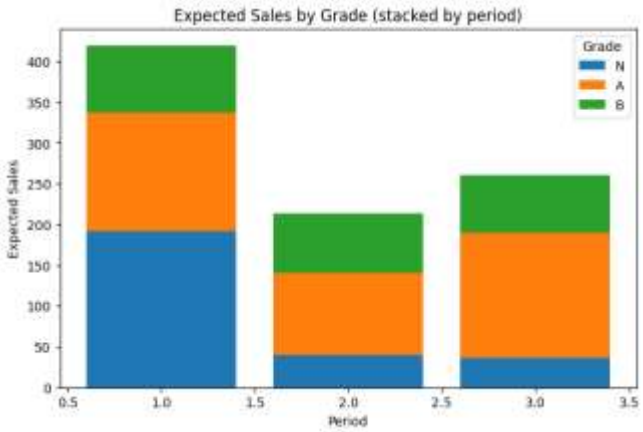

Figure 3 *Expected Sales by Grades*

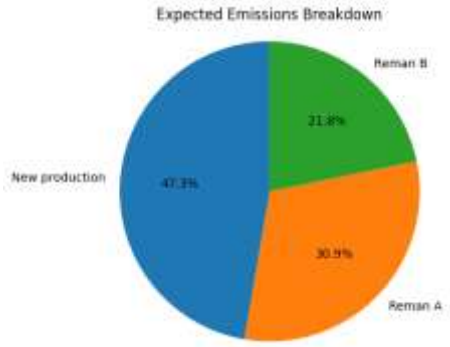

Figure 4 *Expected Emissions Breakdown*

4.1 Discussion and Managerial Implications

The frontier reveals how the marginal cost of abatement evolves as emissions are tightened. Near $\varepsilon = 600$, moving one step left on the curve requires reallocating capacity from new build to reman and relaxing premium exposure limits, which can be costly in foregone margins. Conversely, beyond $\varepsilon \geq 900$, additional emissions headroom mostly produces new units and yields diminishing profit gains because demand saturation in the utility segment and inventory carrying costs start to bind. From a policy perspective, the frontier makes transparent the range of internal carbon prices (shadow prices of the ε-constraint) that are consistent with observed decarbonization choices, offering a way to benchmark carbon budgets and procurement incentives (Mavrotas, 2009; Ellen MacArthur Foundation, 2013).

4.2 Limitations and Future Research

Our linear model abstracts from price endogeneity and strategic competition; prices are treated as exogenous. Incorporating price–demand curves and market equilibrium would turn the planning problem into a mathematical program with equilibrium constraints. We also assume deterministic lead times and ignore warranty feedbacks. Extensions could include dynamic pricing, endogenous core-supply incentives, multi-plant networks, and robust or distributionally robust formulations for heavy-tailed uncertainties. Finally, empirical validation on proprietary wind fleet data would strengthen external validity (Birge & Louveaux, 2011; Guide & Van Wassenhove, 2009).

5. Conclusion

We presented a practical stochastic programming template for closed-loop planning of wind-turbine nacelles that integrates SAA for uncertainty and an ε-constraint for emissions. The resulting profit–emissions frontier helps managers and policymakers visualize efficient decarbonization pathways and quantify shadow prices of carbon. Because the model reads Excel scenarios and is implemented with open-source tools, it is immediately replicable and extensible to richer settings, including policy analysis, multi-plant networks, and price-sensitive demands.

6. References

1. Atasu, A., Sarvary, M., & Van Wassenhove, L. N. (2008). Remanufacturing as a marketing strategy. *Management Science*, *54*(10), 1731–1746. https://doi.org/10.1287/mnsc.1080.0893
2. Birge, J. R., & Louveaux, F. (2011). *Introduction to Stochastic Programming* (2nd ed.). Springer. https://doi.org/10.1007/978-1-4614-0237-4

3. Ellen MacArthur Foundation. (2013). *Towards the circular economy Vol. 1: An economic and business rationale for an accelerated transition*. https://www.ellenmacarthurfoundation.org/towards-the-circular-economy-vol-1-an-economic-and-business-rationale-for-an
4. Govindan, K., Soleimani, H., & Kannan, D. (2015). Reverse logistics and closed-loop supply chain: A comprehensive review to explore the future. *European Journal of Operational Research, 240*(3), 603–626. https://doi.org/10.1016/j.ejor.2014.07.012
5. Guide, V. D. R., & Van Wassenhove, L. N. (2009). The evolution of closed-loop supply chain research. *Operations Research, 57*(1), 10–18. https://doi.org/10.1287/opre.1080.0628
6. Kleywegt, A. J., Shapiro, A., & Homem-de-Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization, 12*(2), 479–502. https://doi.org/10.1137/S1052623499363220
7. Mavrotas, G. (2009). Effective implementation of the ε-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation, 213*(2), 455–465. https://doi.org/10.1016/j.amc.2009.03.037
8. Borja Ponte, Salvatore Cannella, Roberto Dominguez, Mohamed M. Naim,Aris A. Syntetos (2021). Quality grading of returns and the dynamics of remanufacturing. *International Journal of Production Economics, 236*, 108129.
9. Verweij, B., Ahmed, S., Kleywegt, A. J., Nemhauser, G., & Shapiro, A. (2003). The sample average approximation method applied to stochastic routing problems: A computational study. *Computational Optimization and Applications, 24*, 289–333.