

CARBON-AWARE HYBRID REPLENISHMENT UNDER UNCERTAINTY

GÁBOR NAGY¹–SZABOLCS SZENTESI²–GERGŐ ZEMLÉNYI³–TIHOMIR OPETUK⁴

Abstract: Inventory policies influence not only holding and shortage costs but also the transport emissions associated with replenishment frequency, shipment size, and the use of expedited delivery. Existing carbon-aware inventory models typically represent emissions as a cost term or regulatory constraint, while adaptive inventory approaches mainly respond to changes in demand and lead time. This study proposes a carbon-aware adaptive replenishment policy that extends a conventional adaptive (s, S) model by incorporating the anticipated environmental consequence of delayed ordering. Three policies – a static (s, S) policy, a conventional adaptive policy, and the proposed carbon-aware policy – were evaluated through Monte Carlo simulation under stable conditions, demand uncertainty, and combined demand and lead-time uncertainty. The results show that the additional value of carbon-aware adjustment is negligible in a stable environment but becomes more visible as uncertainty increases. Under combined uncertainty, the proposed policy reduced total logistics cost by 3.61%, expedited shipments by 28.48%, and transport-related emissions by 6.73% relative to the static benchmark, while improving the fill rate by 1.69 percentage points. These benefits were accompanied by a 15.97% increase in average inventory. Compared with the conventional adaptive policy, the additional gains were smaller, indicating that adaptive parameter updating remains the primary source of improvement. Carbon awareness therefore acts as a targeted refinement that is most useful when emergency replenishment is frequent and substantially more emission-intensive than regular supply.

Keywords: carbon-aware inventory control; hybrid replenishment; emergency shipment; demand uncertainty; lead-time uncertainty; sustainable logistics

1. INTRODUCTION

Inventory management is traditionally concerned with maintaining product availability at an acceptable level of capital commitment and operating cost. The underlying decision, however, has consequences beyond the warehouse. Order frequency, shipment size, replenishment timing, and transport-mode selection influence vehicle utilization, delivery consolidation, energy consumption, and greenhouse-gas emissions. A policy that appears economical from a narrow inventory perspective may therefore produce inefficient transport operations. Frequent small deliveries can reduce average stock, but they may increase vehicle movements and emissions per delivered unit. Conversely, larger and less frequent orders can improve transport utilization while increasing holding cost, storage requirements, and exposure to obsolescence. The relationship between inventory decisions and environmental performance has been recognized in several foundational studies. Early carbon-aware supply-chain models demonstrated that conventional operational decisions can be modified to incorporate emissions without abandoning their original analytical structure [1]. Carbon trading was subsequently introduced into inventory models to examine how allowance prices

¹ PhD, Institute of Logistics, University of Miskolc, Hungary
gabor.nagy4@uni-miskolc.hu

² PhD, University of Miskolc, Hungary
szabolcs.szentesi@uni-miskolc.hu

³ PhD student, University of Miskolc, Hungary
gergo.zemlenyi@student.uni-miskolc.hu

⁴ associate professor, University of Zagreb, Croatia
tihomir.opetuk@fsb.unizg.hr

affect order quantities and total cost [2]. Related research extended the classical EOQ formulation through explicit carbon constraints [3] and broader sustainability criteria covering transportation, storage, and material handling [4]. These contributions established that economic order quantities and environmentally preferable order quantities need not coincide. More recent reviews confirm that sustainable inventory research has largely developed around carbon taxes, emission caps, cap-and-trade systems, and offsetting arrangements [5].

A second research stream addresses the joint planning of inventory and transportation. Inventory-routing models coordinate replenishment timing, delivery quantities, and vehicle routes within a common decision problem [6]. Environmental extensions have incorporated uncertain demand, product deterioration, explicit fuel consumption, and service constraints [7]. Later models considered heterogeneous fleets and load-dependent emissions [8], as well as alternative carbon-control policies under demand uncertainty [9]. The sustainable inventory-routing literature therefore provides a stronger representation of the relationship between storage and transportation than isolated inventory models. Nevertheless, its primary focus remains the optimization of delivery schedules and routes. The activation of a replenishment decision, the selection between alternative inventory policies, and the adaptive generation of policy parameters are generally embedded in the optimization model rather than treated as distinct decision functions [10]. Uncertainty further complicates the coordination of inventory and transport decisions. Demand changes may result from seasonality, promotional activity, structural market shifts, or short-term volatility. Supply-side uncertainty may arise from variable lead times, capacity restrictions, quality problems, or transport disruption. Inventory-routing research under uncertainty shows that current replenishment decisions must account for both present inventory states and future demand realizations [11]. Robust and risk-averse models extend this logic by protecting the system against adverse outcomes instead of optimizing only expected performance [12].

At the inventory-policy level, changes in lead time may invalidate previously calibrated reorder points and safety-stock levels, requiring a controlled transition to a new operating state [13]. The growing availability of operational data has also encouraged the integration of predictive and reactive decision mechanisms. Predictive components estimate future demand, lead time, or disruption risk, whereas reactive mechanisms respond to realized deviations. Hybrid control research stresses that these functions should be coordinated but not conceptually merged [14]. Learning-based inventory models offer additional flexibility in sequential decision problems, although their practical application requires careful state definition, constrained action spaces, reliable training data, and comparison with strong analytical policies [15]. The value of adaptation therefore depends not only on predictive accuracy but also on how predictive information modifies the actual replenishment logic. The distinction between regular and emergency replenishment is especially important in this context. A regular supply mode usually offers lower acquisition and transportation costs but involves a longer lead time. An emergency mode can reduce the duration or magnitude of a stockout, but it normally incurs a higher fixed cost, a higher unit cost, or both. Periodic-review systems with emergency replenishment have shown that a second supply mode can generate substantial savings when it is activated selectively in response to imminent shortages [16]. More recent research distinguishes time-based, quantity-based, and hybrid emergency-shipment policies and demonstrates that the trigger condition strongly affects cost and service performance [17].

The environmental consequences of emergency replenishment have received less systematic attention. Expedited orders are often transported in smaller quantities, with lower consolidation rates and tighter delivery requirements. Their carbon intensity per delivered unit may therefore be considerably higher than that of planned replenishment. A decision rule that waits until a shortage becomes imminent can preserve inventory efficiency in the short term but may force the use of a carbon-intensive transport option later. Earlier activation of standard replenishment may avoid this outcome, although it can increase average inventory. The resulting decision is not a simple choice between lower cost and lower emissions. It involves a three-way trade-off among inventory exposure, service reliability, and transport-related environmental performance. Recent integrated green logistics research has combined predictive updating, multi-objective routing, service performance, and auditable emissions accounting within a shared simulation-based workflow [18]. That approach provides a basis for connecting operational decisions with environmental records, but its prescriptive core is routing and scheduling. It does not explicitly distinguish replenishment-trigger activation, inventory-policy selection, and parameter generation as separate inventory-control functions.

This study addresses the resulting research gap by developing a carbon-aware hybrid replenishment framework under demand and lead-time uncertainty. The framework separates three decision functions. The trigger layer determines whether the observed and predicted system state requires intervention. The policy layer evaluates regular, adaptive, and expedited replenishment alternatives. The model layer generates and updates the parameters used by these policies. Realized cost, service, emergency-shipment, and emission outcomes are subsequently fed back into the next decision cycle.

The research is guided by the following questions:

- RQ1. How does carbon-aware hybrid replenishment affect total logistics cost, fill rate, inventory level, emergency-shipment frequency, and transport-related emissions compared with conventional inventory policies?
- RQ2. Under which combinations of demand variability and lead-time uncertainty does carbon-aware adaptation provide a measurable operational and environmental advantage?
- RQ3. What are the individual contributions of trigger fusion, policy selection, and adaptive emission feedback to the performance of the complete framework?

The study makes three main contributions. First, it extends carbon-aware inventory control by allowing transport-related carbon information to influence the activation and selection of replenishment policies rather than appearing only in the objective function or as an external constraint. Second, it connects regular and expedited replenishment through a layered trigger-policy-model architecture. Third, it introduces a closed feedback mechanism in which realized cost, service, emergency-shipment, and carbon indicators modify subsequent decision parameters. The framework is evaluated through a controlled Monte Carlo simulation comparing a static (s, S) policy, a conventional adaptive (s, S) policy, and the proposed carbon-aware adaptive replenishment policy under three uncertainty scenarios. The remainder of the paper is structured as follows. Section 2 reviews carbon-aware inventory models, sustainable inventory-routing research, adaptive replenishment approaches, and emergency-shipment policies, and identifies the related research gap. Section 3 presents the proposed carbon-aware replenishment model, including the investigated problem, the compared policies, and the cost and emission structure. Section 4 describes the simulation

design, scenario settings, model parameters, and performance indicators. Section 5 presents and discusses the simulation results, including the scenario-specific findings and the sensitivity analysis. Section 6 summarizes the main contributions, managerial implications, limitations, and directions for future research.

2. LITERATURE REVIEW AND RESEARCH GAP

2.1. Carbon-aware inventory control

Carbon-aware inventory research extends conventional replenishment models by associating environmental effects with purchasing, ordering, transportation, storage, and production decisions. Benjaafar et al. [1] demonstrated that widely used supply-chain models can incorporate carbon emissions through additional accounting terms or constraints. Their analysis also showed that operational changes may reduce emissions without causing proportionate cost increases. This result is important because it establishes inventory policy as a potential environmental intervention rather than treating emissions exclusively as a technological or reporting issue. Hua et al. [2] examined inventory decisions under a carbon-trading mechanism. Their formulation showed that the price of allowances and the firm's initial carbon allocation can alter the economic order quantity and total cost. Chen et al. [3] considered an explicit carbon constraint within the EOQ model and examined the relationship between order frequency, operating cost, and emissions. Battini et al. [4] expanded the scope further by incorporating environmental effects associated with transportation, storage, and material handling into a sustainable EOQ formulation. Taken together, these studies demonstrate that shipment size and replenishment frequency influence both economic and environmental performance.

Despite their analytical value, most carbon-aware EOQ formulations retain a static decision structure. Demand, lead time, emission factors, and policy parameters are generally assumed to be known or fixed over the relevant horizon. Carbon is commonly included as a cost coefficient, a cap, or a tradable resource. It rarely functions as state-dependent information that can modify whether a replenishment decision should be activated or which policy should be used. The taxonomy developed by Sebatjane [5] confirms the strong emphasis on formal carbon-regulation regimes. The reviewed literature mainly addresses carbon taxes, direct emission limits, cap-and-trade arrangements, and offsetting mechanisms. These instruments are relevant for long-term policy evaluation, but they do not fully capture operational settings in which the carbon consequences of replenishment depend on the current stock position, shipment urgency, available transport capacity, and selected delivery mode. A further limitation concerns the treatment of service reliability. A carbon-efficient policy may lower shipment frequency or improve load consolidation, but it can also increase replenishment lead time or reduce responsiveness. Conversely, a policy designed exclusively around service may rely heavily on expedited transport. Carbon-aware inventory control therefore requires simultaneous consideration of holding cost, shortage risk, service performance, and transport emissions.

2.2. Sustainable inventory-routing models

The inventory-routing problem provides the principal operations-research framework for coordinating replenishment and distribution. Instead of treating customer orders as fixed

inputs, the supplier determines delivery timing, quantities, and routes based on inventory states and operational constraints. The literature reviewed by Coelho et al. [6] covers a wide range of deterministic and stochastic formulations, including single- and multi-product settings, different replenishment policies, and alternative routing structures. Environmental considerations have progressively expanded this framework. Soysal et al. [7] formulated a multi-period inventory-routing model for perishable products that incorporates uncertain demand, product shelf life, load-dependent fuel consumption, and a service-level constraint. The model connects transportation emissions with inventory and delivery decisions rather than assigning a constant environmental value to each order. Its application to food distribution also illustrates the importance of product characteristics when evaluating the cost–emission relationship.

Cheng et al. [8] introduced a green inventory-routing problem with a heterogeneous fleet. Their model considers vehicle-dependent operating conditions, fuel consumption, and carbon emissions. The heterogeneous-fleet setting is relevant because the environmental consequence of a replenishment decision depends not only on delivery quantity and distance but also on the available vehicle type and its utilization. Micheli and Mantella [9] extended environmentally oriented inventory routing by combining uncertain demand, heterogeneous vehicles, a detailed emissions model, and alternative carbon-control policies. Their results show that the effectiveness and cost of environmental regulation depend on the structure of the operating system and on the selected control mechanism. A carbon cap, carbon tax, cap-and-trade scheme, and offsetting arrangement can lead to different inventory and routing decisions even within the same logistics network. The review by Soysal et al. [10] demonstrates the growing breadth of sustainable inventory-routing research. The literature includes fuel consumption, carbon emissions, perishability, waste, reverse flows, collaboration, and closed-loop distribution. At the same time, the review indicates that sustainability-related inventory-routing models remain methodologically diverse and often problem-specific. The principal strength of inventory-routing models is their ability to coordinate inventory and transportation within a unified optimization problem. Their limitation for the present research lies elsewhere. Replenishment activation and policy choice are usually implicit in the optimization structure. The models determine an optimal or near-optimal delivery plan, but they do not generally distinguish between:

- the signal that activates reconsideration of the inventory decision;
- the policy that generates a candidate replenishment response;
- and the model that updates the policy parameters.

This distinction becomes important when several data sources and supply modes operate simultaneously. A low inventory position, an adverse demand forecast, a worsening lead-time estimate, and a high expected carbon intensity may produce different or conflicting signals. A layered architecture is needed to determine how these signals should influence the replenishment decision.

2.3. Uncertainty and adaptive replenishment

Inventory and distribution systems operate under uncertainty in both demand and supply. Cui et al. [11] provide a general treatment of inventory routing under uncertain conditions. Their work emphasizes the intertemporal relationship between inventory status, delivery decisions, and future demand. This relationship limits the effectiveness of static policies when demand

distributions or operating conditions change over time. Risk-neutral optimization based on expected cost may be insufficient when adverse scenarios have severe operational consequences. Feng et al. [12] address this issue through a robust and risk-averse inventory-routing formulation. Their approach gives greater attention to unfavorable demand realizations and illustrates how protection against extreme outcomes can alter inventory and routing decisions. Lead-time uncertainty is particularly relevant for replenishment control because the reorder point and safety stock depend directly on the quantity of demand expected during replenishment. Axsäter [13] showed that a change in lead time modifies the optimal steady-state solution and requires a managed transition between inventory states. Simply retaining the original parameters after the lead time has changed can generate unnecessary cost or service deterioration. Adaptation may occur at several levels. The system can update a forecast, revise a parameter, change an activation threshold, or select a different policy. These forms of adaptation should not be treated as equivalent.

Cardin et al. [14] distinguish predictive planning from reactive control within hybrid architectures. Predictive mechanisms anticipate future conditions, while reactive mechanisms manage deviations that have already occurred. Effective integration requires clear interfaces and decision authority between the two functions. Learning-based inventory methods extend adaptation across multiple decision periods. Boute et al. [15] describe the potential of deep reinforcement learning in inventory control, but they also highlight its demanding design choices and computational requirements. State representation, reward construction, policy constraints, training stability, and benchmark selection directly influence whether a learned policy is useful in practice. These concerns support a restrained approach in which data-driven components update parameters or select among admissible policies rather than operating without structural safeguards. For the proposed study, adaptation is therefore interpreted as controlled parameter and policy adjustment. The system retains explicit replenishment rules, but their activation thresholds, target levels, and selection weights may change in response to persistent variations in demand, lead time, service performance, and transport-related emissions.

2.4. Regular and emergency replenishment

Emergency replenishment provides a direct response to stockout risk when the next regular order cannot arrive in time. Tagaras and Vlachos [16] examined a periodic-review inventory system with two replenishment modes. Regular orders follow a base-stock policy and arrive after a deterministic lead time. Emergency orders have shorter lead times and higher acquisition costs. Their results show that selective emergency replenishment can generate considerable cost savings compared with a single-mode system. Poormoaiied et al. [17] developed a quantity-based emergency-shipment policy in which an emergency order is triggered when the inventory position falls below a defined level during the review period. They compared quantity-based, time-based, and hybrid activation rules and demonstrated that the trigger structure materially influences the resulting cost. Continuous inventory monitoring can support faster intervention, whereas fixed-time checks may respond too late to unusually high demand. These studies provide a strong basis for policy-level hybridization. Regular and emergency replenishment are not simply two versions of the same order. They differ in activation conditions, lead time, cost structure, and operational role. However, their environmental consequences are typically not included in the trigger logic. The emergency mode is evaluated mainly through shortage reduction and additional cost. A carbon-aware

extension requires the transport consequences of urgency to be represented explicitly. Let the standard mode offer a longer lead time, lower fixed transport cost, and lower emissions per delivered unit because of better consolidation. Let the expedited mode offer a shorter lead time but higher cost and carbon intensity. Under this structure, a delayed decision can create a need for expedited replenishment even when an earlier standard shipment would have produced a more favorable combined result. The relevant policy question is therefore not merely whether emergency replenishment should be available. It is whether the system should modify its regular ordering behaviour when it anticipates that postponement will lead to a high-cost and high-emission emergency shipment.

2.5. Research gap and positioning of the present study

The reviewed literature provides established models for carbon-constrained inventory management, sustainable inventory routing, uncertainty treatment, adaptive control, and emergency replenishment. These research streams are closely related, but they remain only partially integrated. Carbon-aware inventory models generally place emissions in the objective function or impose them through a regulatory constraint. Sustainable inventory-routing models coordinate delivery quantities and transport decisions, but replenishment activation and policy selection are often embedded in a single prescriptive model. Adaptive inventory studies respond to changing demand and lead time, yet their feedback variables are mainly cost, forecast accuracy, and service performance. Emergency-replenishment models distinguish regular and expedited supply modes but rarely allow transport-related carbon intensity to influence the activation threshold or subsequent policy parameters. A recent integrated green logistics framework connected predictive updating, multi-objective routing, cooperative allocation, and auditable emissions accounting [18]. Its contribution lies in workflow-level integration around a routing and scheduling core. The present study shifts the analytical focus to replenishment control. It examines how carbon performance can influence the trigger, policy, and model levels of an inventory decision (Table I).

Table I.
Positioning of the proposed framework

Research stream	Inventory decision	Transport integration	Adaptive parameters	Emergency mode	Carbon feedback
Carbon-aware EOQ and inventory models [1–5]	Yes	Limited	Usually no	Usually no	No
Sustainable inventory routing [6–10]	Yes	Yes	Limited	Model-specific	Usually no
Uncertain and adaptive replenishment [11–15]	Yes	Limited	Yes	Limited	No
Emergency replenishment [16, 17]	Yes	Partial	Policy-specific	Yes	No
Proposed framework	Yes	Yes	Yes	Yes	Yes

The novelty of the study does not lie in introducing a new standalone inventory policy or emissions model. It lies in the decision architecture connecting these components. The

framework treats carbon performance as operational information that can influence future replenishment behaviour, rather than solely as an outcome reported after the decision.

3. PROPOSED CARBON-AWARE REPLENISHMENT MODEL

3.1. Problem description

The study considers a single-item inventory system supplied through regular and expedited replenishment. Demand occurs in discrete periods and is uncertain. Regular replenishment has a longer lead time but lower transport cost and lower carbon intensity. Expedited replenishment arrives sooner but is more expensive and generates higher emissions per delivered unit. The investigated decision problem concerns the timing and quantity of replenishment. A policy that reacts too late to changing demand or lead-time conditions may require an expedited shipment to avoid a stockout. Earlier regular replenishment may reduce this risk, although it can also increase the average inventory level. The decision therefore involves a trade-off among holding cost, service performance, emergency-shipment cost, and transport-related emissions. Backorders are permitted, meaning that demand not satisfied immediately is fulfilled after replenishment arrives. The system is evaluated over a finite planning horizon under different demand and lead-time conditions.

3.2. Compared replenishment policies

Three replenishment policies are compared.

Static (s, S) policy. The benchmark policy uses fixed reorder-point and order-up-to parameters. An order is placed when the inventory position falls to or below the reorder point:

$$Q_t = \begin{cases} S - IP_t, & IP_t \leq s, \\ 0, & IP_t > s, \end{cases} \quad (1)$$

where IP_t denotes the inventory position, s is the reorder point, and S is the order-up-to level.

The policy does not respond to changes in demand variability or replenishment lead time.

Adaptive (s, S) policy. The adaptive policy periodically updates the reorder point using the latest demand and lead-time information:

$$s_t = \hat{\mu}_{D,t} \hat{L}_t + z \hat{\sigma}_{DL,t}, \quad (2)$$

where $\hat{\mu}_{D,t}$ is the estimated average demand, $\hat{L}_t$ is the estimated regular lead time, $\hat{\sigma}_{DL,t}$ represents demand uncertainty during the lead time, and z is the service-level factor.

The corresponding order quantity is:

$$Q_t = \begin{cases} s_t - IP_t, & IP_t \leq s_t, \\ 0, & IP_t > s_t. \end{cases} \quad (3)$$

This policy responds to changing operating conditions but does not directly consider transport emissions.

Carbon-aware adaptive policy. The third policy extends the adaptive model by considering the expected consequences of postponing regular replenishment. When the probability of a future stockout is high, waiting may create a need for expedited transportation.

The carbon-aware reorder point is defined as:

$$s_t^C = s_t + \Delta_t^C, \quad (4)$$

where Δ_t^C is a carbon-related adjustment. The adjustment increases when postponing replenishment is likely to require an expedited shipment with higher transport emissions.

The resulting policy may initiate regular replenishment earlier than the conventional adaptive policy. Its objective is not simply to minimize emissions. It seeks to avoid situations in which a small reduction in inventory holding leads to a disproportionately expensive and carbon-intensive emergency delivery.

3.3. Cost and emission model

Total logistics cost consists of holding, shortage, ordering, regular transport, and expedited transport costs:

$$TC = HC + SC + OC + RC + EC. \quad (5)$$

Transport-related emissions are calculated separately for regular and expedited replenishment:

$$TE = \sum_{t=1}^T (e^R Q_t^R + e^E Q_t^E), \quad (6)$$

where e^R and e^E denote the emission factors of regular and expedited transport, respectively. The model assumes:

$$e^E > e^R, c^E > c^R, L^E < L^R. \quad (7)$$

Thus, expedited replenishment provides a service advantage but creates an economic and environmental disadvantage.

3.4. Performance measures

Only four principal indicators are required:

- Total logistics cost
- Fill rate
- Number or proportion of expedited shipments
- Transport-related carbon emissions

Average inventory can be reported as a supporting measure, because earlier replenishment may reduce emergency deliveries at the cost of carrying additional stock.

The fill rate is calculated as:

$$FR = 1 - \frac{\text{unfulfilled demand}}{\text{total demand}}. \quad (8)$$

A policy is considered superior only when it reduces cost or emissions without causing an unacceptable decline in the fill rate.

4. SIMULATION DESIGN

4.1. Experimental setup

The three replenishment policies were evaluated through a discrete-time Monte Carlo simulation. Each simulation period represents one week of operation, and the planning horizon consists of 52 periods. A single product is considered, while unmet demand is backordered and fulfilled when the required inventory becomes available.

The simulation compares the following policies:

- **P1: Static (s, S) policy**, using fixed replenishment parameters;
- **P2: Adaptive (s, S) policy**, periodically updating the reorder point according to observed demand and lead-time conditions;
- **P3: Carbon-aware adaptive policy**, extending the adaptive policy with an additional adjustment intended to reduce the future need for expedited replenishment.

The three policies are exposed to the same demand and lead-time realizations in each replication. This common-random-number approach ensures that differences in performance result from the applied decision rules rather than from different stochastic input sequences.

A total of 1,000 independent replications are performed for each scenario. This number provides a sufficiently stable basis for comparing mean cost, service, emergency-shipment frequency, and transport-related emissions. Results are reported as simulation means together with 95% confidence intervals.

The initial inventory position is identical for all policies. Outstanding orders and backorders are set to zero at the beginning of each replication. The initial reorder point is calculated from expected demand during the regular replenishment lead time:

$$s_0 = \mu_D L^R + z \sigma_{DL}, \quad (9)$$

where μ_D denotes expected weekly demand, L^R is the regular lead time, z is the service-level factor, and σ_{DL} represents demand uncertainty during the replenishment lead time.

The initial order-up-to level is:

$$S_0 = s_0 + q, \quad (10)$$

where q is the planned replenishment quantity above the reorder point.

The service-level factor is set to $z = 1.65$, corresponding to an initial target service probability of approximately 95%. This value is used consistently across the three policies so that differences do not arise from different initial service targets.

4.2. Demand and lead-time scenarios

Three scenarios are used to represent increasing levels of operational uncertainty.

SC1: Stable operating environment. The first scenario represents a relatively predictable system. Weekly demand fluctuates around a stable mean with limited variation, while the regular replenishment lead time remains constant.

This scenario provides the reference setting. It indicates whether the additional complexity of adaptive and carbon-aware replenishment produces any meaningful advantage when the operating environment is already stable.

SC2: Demand uncertainty. The second scenario retains a stable regular lead time but introduces substantially higher demand variability. The expected demand remains unchanged, allowing the effects of variability to be distinguished from those of long-term demand growth.

This scenario examines whether adaptive parameter updating can reduce the occurrence of shortages and emergency replenishment when demand becomes less predictable.

SC3: Combined demand and lead-time uncertainty. The third scenario combines high demand variability with a stochastic regular replenishment lead time. Regular orders may therefore arrive later than expected, increasing the probability that the available inventory becomes insufficient before the next replenishment.

This represents the most challenging operating environment. The scenario is particularly relevant for assessing whether the carbon-aware adjustment can initiate regular replenishment early enough to avoid later use of expedited transport (Table II).

Table II.
Main simulation scenarios

Scenario	Weekly demand	Regular lead time	Purpose
SC1	Stable, low variability	Constant	Baseline comparison
SC2	Stable mean, high variability	Constant	Demand-risk assessment
SC3	High variability	Stochastic	Combined demand and supply uncertainty

Weekly demand is generated from a non-negative normal distribution. Negative values, although unlikely under the selected parameters, are replaced by zero. The general demand process is:

$$D_t = \max\{0, \mu_D + \varepsilon_t\}, \varepsilon_t \sim N(0, \sigma_D^2). \quad (11)$$

The expected weekly demand is set to 200 units in all scenarios. Demand variability is represented by a standard deviation of 30 units in SC1 and 70 units in SC2 and SC3.

The regular lead time is fixed at two weeks in SC1 and SC2. In SC3, it follows a discrete distribution:

$$L_t^R = \begin{cases} 2, & \text{with probability 0.50,} \\ 3, & \text{with probability 0.30,} \\ 4, & \text{with probability 0.20.} \end{cases} \quad (12)$$

Expedited replenishment has a one-week lead time in all scenarios. Its shorter lead time allows the system to respond to imminent shortages, but at a higher cost and emission intensity.

4.3. Policy parameterization

The static policy retains its initial values of s and S throughout the entire planning horizon. It therefore does not respond to changes in observed demand variability or regular lead time.

The adaptive policy updates its reorder point every four periods. The update is based on the most recent demand observations and the current estimate of the regular lead time:

$$s_t = \hat{\mu}_{D,t} \hat{L}_t^R + z \hat{\sigma}_{DL,t}. \quad (13)$$

The corresponding order-up-to level is:

$$S_t = s_t + q. \quad (14)$$

A four-period updating interval is used to prevent the policy from reacting excessively to individual observations. It provides a balance between responsiveness and parameter stability.

The carbon-aware policy applies the same adaptive calculation but introduces an additional adjustment:

$$s_t^C = s_t + \Delta_t^C. \quad (15)$$

The adjustment depends on the estimated probability that postponing regular replenishment would create a need for expedited delivery:

$$\Delta_t^C = \kappa P_t^E \hat{\mu}_{D,t}, \quad (16)$$

where P_t^E is the estimated probability of requiring expedited replenishment and κ is the strength of the carbon-aware adjustment.

A higher value of P_t^E shifts the effective reorder point upward. The policy may consequently initiate a regular shipment earlier than the conventional adaptive policy. When the risk of expedited replenishment is low, the adjustment approaches zero and the policy behaves similarly to the standard adaptive model.

The baseline value of κ is set to 0.25. This means that the maximum additional inventory protection corresponds to approximately one quarter of expected weekly demand. The parameter is kept deliberately moderate to prevent the carbon-aware policy from achieving lower emergency-shipment frequency solely by maintaining excessive inventory.

4.4. Cost and emission parameters

The simulation uses normalized monetary and emission units. This approach allows the comparison to focus on structural policy behaviour without claiming calibration to a specific company or transport network.

Regular and expedited replenishment differ in fixed cost, variable cost, lead time, and carbon intensity (Table III).

The baseline model assumes that expedited transport generates twice the emissions per delivered unit compared with regular replenishment:

$$e^E = 2e^R. \quad (17)$$

The purpose of this assumption is not to represent a universal physical emission ratio. Expedited shipments may differ considerably according to transport mode, shipment size, distance, and vehicle utilization. The ratio therefore serves as an experimental parameter

whose influence is examined separately. Holding and shortage costs are applied at the end of each period. Regular and expedited ordering costs are charged whenever the corresponding order type is activated. Total emissions are calculated from the delivered quantities and their assigned emission factors.

Table III.
Baseline simulation parameters

Parameter	Symbol	Baseline value
Mean weekly demand	μ_D	200 units
Demand standard deviation, SC1	σ_D	30 units
Demand standard deviation, SC2–SC3	σ_D	70 units
Regular lead time, SC1–SC2	L^R	2 weeks
Expedited lead time	L^E	1 week
Service-level factor	z	1.65
Order-up-to increment	q	300 units
Holding cost	h	1 cost unit/unit/week
Backorder cost	b	8 cost units/unit/week
Regular order cost	K^R	100 cost units/order
Expedited order cost	K^E	250 cost units/order
Regular emission factor	e^R	1 emission unit/unit
Expedited emission factor	e^E	2 emission units/unit
Carbon-adjustment coefficient	κ	0.25

4.5. Performance evaluation

The policy comparison is based on five performance indicators:

1. Total logistics cost
2. Fill rate
3. Average inventory
4. Frequency of expedited replenishment
5. Transport-related emissions

The frequency of expedited replenishment is calculated as the proportion of periods in which an expedited order is placed:

$$ER = \frac{1}{T} \sum_{t=1}^T \mathbf{1}\{Q_t^E > 0\}. \quad (18)$$

The main comparisons are conducted between:

- P1 and P2, showing the effect of adaptation;
- P2 and P3, showing the additional effect of carbon awareness;
- P1 and P3, showing the total improvement relative to the static benchmark.

A policy is not regarded as superior solely because it produces the lowest emissions. The interpretation also considers its fill rate and average inventory. A large emission reduction achieved through excessive inventory accumulation would not represent a balanced improvement.

For each performance measure, the relative difference between policy j and the static benchmark is calculated as:

$$\Delta KPI_j = \frac{KPI_j - KPI_{P1}}{KPI_{P1}} \times 100. \quad (19)$$

For cost, emissions, average inventory, and emergency-shipment frequency, a negative percentage represents an improvement. For the fill rate, a positive percentage indicates improved service performance.

4.6. Sensitivity analysis

A focused sensitivity analysis is conducted for the emission intensity of expedited replenishment. Three ratios are examined:

$$\frac{e^E}{e^R} \in \{1.5, 2.0, 3.0\}. \quad (20)$$

The first setting represents a moderate environmental disadvantage of expedited delivery. The second is the baseline case, while the third represents a highly carbon-intensive emergency mode. This analysis addresses whether the carbon-aware policy remains useful when the environmental difference between the two replenishment modes is relatively small. It also helps identify the point at which earlier regular ordering becomes justified by the avoided emissions from emergency transport. The simulation design deliberately limits the number of scenarios and sensitivity parameters. Its purpose is to provide a transparent comparison of the three policies rather than to reproduce every possible demand, cost, and transport condition. This focused structure also allows the results to be interpreted directly in relation to the two research questions.

5. RESULTS AND DISCUSSION

5.1. Overall policy performance

Table IV summarizes the mean simulation results obtained for the three replenishment policies under the investigated operating scenarios. Each reported value is based on 1,000 independent replications over a 52-period planning horizon. The same demand and lead-time realizations were applied to all three policies within each replication. Total logistics cost and transport-related emissions are expressed in normalized simulation units.

Table IV.
Main simulation results for the compared replenishment policies

Scenario	Policy	Total logistics cost	Fill rate (%)	Average inventory (units)	Expedited shipments (%)	Transport-related emissions
SC1	P1: Static (s, S)	14,470.75	95.99	135.28	14.17	11,095.75
SC1	P2: Adaptive (s, S)	14,506.85	96.00	134.88	14.75	11,130.50
SC1	P3: Carbon-aware adaptive	14,466.29	95.98	134.22	14.38	11,095.53
SC2	P1: Static (s, S)	18,923.70	93.72	183.52	16.19	11,887.43
SC2	P2: Adaptive (s, S)	18,124.76	94.82	198.67	9.58	11,161.53
SC2	P3: Carbon-aware adaptive	18,044.60	94.84	197.74	9.13	11,105.04
SC3	P1: Static (s, S)	20,517.43	92.35	167.31	29.24	13,721.59
SC3	P2: Adaptive (s, S)	19,779.73	93.86	190.67	22.13	12,936.05
SC3	P3: Carbon-aware adaptive	19,777.14	94.04	194.03	20.91	12,797.73

Note: SC1 represents stable operating conditions, SC2 represents high demand uncertainty, and SC3 represents combined demand and regular lead-time uncertainty. P1 is the static benchmark, P2 is the conventional adaptive policy, and P3 is the proposed carbon-aware adaptive policy.

As shown in Table IV, the relative performance of the three policies depended strongly on the level and source of uncertainty. Under SC1, the outcomes were nearly identical. Total logistics cost ranged from 14,466.29 to 14,506.85 cost units, while the fill rate remained close to 96% for all three policies. The static policy therefore remained effective when demand variability was limited and the regular replenishment lead time was stable. The adaptive policy did not provide a material benefit in the stable scenario. Its periodic parameter updates produced a slightly higher expedited-shipment frequency and transport-related emissions than the static benchmark. The carbon-aware policy largely corrected this difference, but its performance remained practically equivalent to P1. This result indicates that additional adaptation is not automatically advantageous when the assumptions underlying the initial inventory parameters remain valid. More substantial differences emerged under SC2. The higher variability of weekly demand reduced the fill rate of the static policy to 93.72% and increased total logistics cost to 18,923.70 cost units. The adaptive policy improved the fill rate to 94.82% and reduced the frequency of expedited replenishment from 16.19% to 9.58%. The carbon-aware policy achieved the strongest overall result in SC2. Its total logistics cost was 18,044.60 cost units, while transport-related emissions decreased to 11,105.04 emission units. The frequency of expedited replenishment fell to 9.13%, and the fill rate increased to

94.84%. These improvements were accompanied by an increase in average inventory from 183.52 to 197.74 units. The differences became most visible under SC3. The combination of high demand variability and uncertain regular lead time reduced the fill rate of the static policy to 92.35%. P1 also required expedited replenishment in 29.24% of the simulated periods and generated 13,721.59 emission units. Both adaptive policies performed more effectively under these conditions. P2 increased the fill rate to 93.86% and reduced expedited shipments to 22.13%. P3 provided a further improvement, reaching a fill rate of 94.04% and reducing the expedited-shipment frequency to 20.91%. Transport-related emissions were lowest under P3, at 12,797.73 emission units. The carbon-aware policy did not achieve this improvement without a trade-off. Its average inventory under SC3 was 194.03 units, compared with 167.31 units for the static policy. Earlier regular replenishment therefore reduced reliance on emergency transport but increased the amount of stock held within the system. To provide a more direct comparison, Table V reports the relative performance changes achieved by the carbon-aware policy against the static benchmark. Percentage differences were calculated separately within each scenario.

Table V.
Relative performance change of the carbon-aware policy compared with the static benchmark

Scenario	Total cost change (%)	Emission change (%)	Expedited shipment change (%)	Fill-rate change (percentage points)	Average inventory change (%)
SC1	−0.03	0.00	+1.52	−0.01	−0.79
SC2	−4.65	−6.58	−43.60	+1.12	+7.75
SC3	−3.61	−6.73	−28.48	+1.69	+15.97

Note: Negative values for total cost, emissions, and expedited shipments indicate a reduction relative to P1. A positive fill-rate change represents improved service performance. A positive average-inventory change indicates additional stock resulting from earlier regular replenishment.

The relative comparison confirms that the value of carbon-aware adaptation depends on operating uncertainty. Under SC1, the performance differences were negligible. P3 reduced total cost by only 0.03%, while transport-related emissions remained effectively unchanged. The frequency of expedited replenishment increased marginally by 1.52%, and the fill-rate difference was only −0.01 percentage points. Under SC2, P3 reduced total logistics cost by 4.65% and transport-related emissions by 6.58% relative to the static benchmark. The frequency of expedited replenishment decreased by 43.60%, while the fill rate improved by 1.12 percentage points. The environmental and service improvements required a 7.75% increase in average inventory. Under SC3, the carbon-aware policy reduced total logistics cost by 3.61%, transport-related emissions by 6.73%, and expedited shipments by 28.48%. The fill rate improved by 1.69 percentage points. However, average inventory increased by 15.97%, illustrating the trade-off between earlier regular ordering and lower dependence on emergency transportation. Taken together, the results indicate that the proposed policy provides little additional value in a stable operating environment. Its advantage becomes more pronounced when demand variability and lead-time uncertainty increase the likelihood of delayed intervention. Under such conditions, the carbon-aware adjustment reduces the

future need for expedited replenishment by initiating selected regular orders earlier. The resulting reduction in cost and emissions must nevertheless be evaluated against the associated increase in average inventory.

5.2. Performance under stable conditions

Under stable operating conditions, the differences among the three replenishment policies remained negligible. The static (s, S) policy achieved a total logistics cost of 14,470.75 cost units and a fill rate of 95.99%. The adaptive policy produced a slightly higher total cost of 14,506.85 cost units, while its fill rate remained almost unchanged at 96.00%. The carbon-aware adaptive policy achieved the lowest total cost, although the difference relative to the static benchmark was only 4.46 cost units, corresponding to a reduction of 0.03%. The environmental results followed a similar pattern. Transport-related emissions amounted to 11,095.75 units under P1 and 11,095.53 units under P3. This difference is operationally negligible and indicates that carbon-aware adjustment provides little additional environmental benefit when demand and lead time remain predictable. The frequency of expedited replenishment was 14.17% for P1, 14.75% for P2, and 14.38% for P3. The carbon-aware policy therefore did not reduce emergency replenishment relative to the static benchmark in SC1. Its expedited-shipment frequency was approximately 1.5% higher than that of P1, although it remained lower than the value obtained for the conventional adaptive policy. Average inventory decreased slightly from 135.28 units under P1 to 134.22 units under P3. This result confirms that the carbon-aware model did not achieve its performance through systematic inventory accumulation in the stable scenario. Instead, all three policies operated under similar replenishment conditions and generated nearly identical outcomes. The results obtained for SC1 suggest that adaptation should not be introduced solely because more recent demand and lead-time information is available. When the operating environment remains close to the assumptions used for initial policy calibration, a conventional static policy may provide an adequate balance among cost, inventory, service, and emissions. The additional decision logic of P2 and P3 becomes useful only when changes in the operating environment materially reduce the suitability of fixed replenishment parameters.

5.3. Performance under demand uncertainty

Higher demand variability produced substantially larger differences among the replenishment policies. Under SC2, the total logistics cost of the static policy increased to 18,923.70 cost units, while its fill rate declined to 93.72%. The fixed reorder point was no longer sufficiently responsive to short-term demand fluctuations, leading to more frequent inventory shortages and a greater need for expedited replenishment. The conventional adaptive policy reduced total cost to 18,124.76 cost units. This represents a reduction of approximately 4.22% relative to the static benchmark. Its fill rate increased by 1.10 percentage points to 94.82%, while expedited shipments decreased from 16.19% to 9.58%. The carbon-aware adaptive policy provided the best overall result in SC2. Its total logistics cost was 18,044.60 cost units, corresponding to a reduction of 4.65% relative to P1. Compared with P2, the additional cost reduction was relatively modest, at approximately 0.44%. This indicates that most of the cost improvement resulted from adaptive parameter updating, while the explicit consideration of transport emissions provided a smaller but still favourable incremental effect. The frequency of expedited replenishment declined from

16.19% under the static policy to 9.13% under P3. This represents a reduction of 43.60% relative to P1 and approximately 4.70% relative to P2. The lower emergency-shipment frequency indicates that the carbon-aware adjustment changed the timing of selected regular orders before the inventory position reached a critical shortage state.

Transport-related emissions decreased from 11,887.43 units under P1 to 11,105.04 units under P3, representing a reduction of 6.58%. Relative to the conventional adaptive policy, P3 achieved an additional emission reduction of approximately 0.51%. The relatively small difference between P2 and P3 shows that adaptive ordering already prevents a substantial share of unnecessary emergency shipments. Nevertheless, the carbon-aware extension further improves the selection of replenishment timing when postponement is likely to create a high-emission response. The environmental and service improvements were accompanied by higher average inventory. P3 maintained an average inventory of 197.74 units, compared with 183.52 units under P1. The resulting increase of 7.75% represents the principal trade-off associated with the proposed approach. Earlier regular replenishment reduces the likelihood of a stockout and the subsequent use of expedited delivery, but it also causes some units to enter inventory before they would have arrived under the static policy. Despite this increase, the fill rate improved by 1.12 percentage points and total logistics cost decreased. The additional inventory was therefore not merely an environmental sacrifice. It also contributed to lower shortage exposure and reduced emergency-order costs. Under demand uncertainty, the proposed policy achieved a more favourable balance among inventory availability, transport emissions, and total cost than the static benchmark.

5.4. Performance under combined uncertainty

The strongest performance differences were observed under SC3, where high demand variability was combined with stochastic regular replenishment lead times. This scenario created two simultaneous sources of operational risk. Demand could exceed its expected level, while regular orders could arrive later than initially anticipated. Consequently, fixed inventory parameters were less capable of maintaining the intended service level. The static policy generated the highest total logistics cost, amounting to 20,517.43 cost units. Its fill rate fell to 92.35%, the lowest value observed across all scenarios and policies. Expedited replenishment was required in 29.24% of the simulated periods, while transport-related emissions reached 13,721.59 units. The conventional adaptive policy improved performance by updating the reorder point according to the observed demand and lead-time conditions. Total logistics cost decreased to 19,779.73 cost units, while the fill rate increased to 93.86%. The frequency of expedited replenishment declined to 22.13%, and transport-related emissions decreased to 12,936.05 units. The carbon-aware policy produced a total logistics cost of 19,777.14 cost units. This represents a reduction of 3.61% relative to P1. The difference between P2 and P3 was very small in cost terms, amounting to only 2.59 cost units. The principal additional benefit of carbon awareness was therefore not a substantial further cost reduction, but a better environmental and service outcome. P3 achieved a fill rate of 94.04%, improving service performance by 1.69 percentage points relative to the static policy and by 0.18 percentage points relative to P2. Expedited shipments decreased to 20.91%, representing a reduction of 28.48% compared with P1 and approximately 5.51% compared with the conventional adaptive policy. Transport-related emissions were reduced from 13,721.59 units under P1 to 12,797.73 units under P3. The resulting reduction of 6.73% was the largest environmental improvement observed in the three scenarios. Compared with P2,

carbon-aware adaptation generated an additional emission reduction of approximately 1.07%. These improvements were associated with a substantial increase in average inventory. P3 maintained an average of 194.03 units, compared with 167.31 units under P1 and 190.67 units under P2. Relative to the static benchmark, average inventory increased by 15.97%. Compared with P2, the increase was more limited, at approximately 1.76%. The results indicate that most of the additional inventory was caused by adaptation to demand and lead-time uncertainty rather than by carbon awareness alone. The difference between P2 and P3 was substantially smaller than the difference between P1 and either adaptive policy. Carbon-aware adjustment therefore refined an already adaptive decision rather than replacing the primary role of demand- and lead-time-based parameter updating. SC3 also demonstrates the main practical value of the proposed model. When supply uncertainty increases, waiting until the conventional reorder point is reached can create a situation in which regular replenishment can no longer arrive before a stockout occurs. Earlier intervention may increase inventory temporarily, but it reduces the need for more expensive and emission-intensive emergency transportation.

5.5. Sensitivity to expedited-transport emissions

The sensitivity analysis examined whether the advantage of the carbon-aware policy changes when the emission intensity of expedited replenishment differs from the baseline assumption. Three relative emission factors were considered:

$$\frac{e^E}{e^R} \in \{1.5, 2.0, 3.0\}. \quad (21)$$

The ratio of 1.5 represents a moderate environmental difference between regular and expedited replenishment. The baseline ratio of 2.0 assumes that expedited transport generates twice the emissions per delivered unit. The ratio of 3.0 represents a setting in which emergency deliveries are particularly carbon-intensive because of limited shipment consolidation, low vehicle utilization, or the use of an emission-intensive transport mode (Table VI).

Table VI.
Sensitivity of the carbon-aware policy to expedited-transport emission intensity

Emission ratio	Total logistics cost	Fill rate (%)	Average inventory (units)	Expedited shipments (%)	Transport emissions	Emission change vs. P2 (%)
1.5	19,777.14	94.04	194.03	20.91	11,166.75	-0.39
2.0	19,777.14	94.04	194.03	20.91	12,797.73	-1.07
3.0	19,777.14	94.04	194.03	20.91	16,059.69	-2.01

Note: Results refer to policy P3 under scenario SC3. The emission ratio represents the emission intensity of expedited replenishment relative to regular replenishment. Since the emission ratio affects emissions accounting but does not directly modify the replenishment decision rule in the present model, total logistics cost, fill rate, average inventory, and expedited-shipment frequency remain unchanged. The final column compares the transport emissions of P3 with those of the conventional adaptive policy P2 under the same emission ratio.

At an emission ratio of 1.5, the expected advantage of carbon-aware adaptation is likely to be smaller because avoiding one expedited shipment produces only a moderate environmental benefit. Under this condition, the additional holding cost associated with earlier replenishment may offset a substantial share of the environmental improvement. At the baseline ratio of 2.0, P3 reduced emissions by 1.07% relative to P2 and by 6.73% relative to the static benchmark. The result indicates that carbon-aware adjustment provides an incremental benefit beyond conventional demand- and lead-time adaptation, although the majority of the overall improvement continues to originate from adaptive inventory control. When the emission intensity of expedited transport increases to three times that of regular replenishment, earlier regular ordering is expected to become more favourable. Under such conditions, the environmental penalty associated with late intervention increases, strengthening the justification for the carbon-related reorder-point adjustment. The sensitivity analysis is important because the emission difference between regular and emergency transport is not universal. It depends on shipment size, distance, transport mode, vehicle type, capacity utilization, and the degree of consolidation. The proposed policy is therefore most relevant in logistics systems where urgent replenishment produces a clearly higher carbon intensity per delivered unit.

5.6. Discussion

The simulation results show that the value of carbon-aware replenishment depends strongly on the operating environment. Under stable conditions, all three policies produced similar cost, service, and environmental outcomes. The static (s, S) policy remained effective because the demand and lead-time assumptions used during its calibration continued to reflect the realized system conditions. This finding is consistent with the role of conventional inventory policies as efficient decision rules in relatively stationary environments. The results do not support the assumption that a more complex or adaptive policy is automatically superior. When parameter instability is limited and emergency replenishment is already infrequent, carbon-related adjustment adds little measurable value. The situation changed when demand variability increased.

Under SC2, adaptive updating reduced total logistics cost, improved the fill rate, and lowered the need for expedited replenishment. P3 achieved the strongest overall result, but the difference between P2 and P3 was smaller than the difference between either adaptive policy and the static benchmark. This suggests that demand adaptation remains the primary driver of operational improvement, while carbon awareness acts as a secondary refinement. Combined uncertainty strengthened the relevance of carbon-related information. Under SC3, the carbon-aware policy reduced expedited shipments by 28.48% and emissions by 6.73% relative to the static policy. It also improved the fill rate by 1.69 percentage points. These gains were achieved with a 15.97% increase in average inventory. The observed trade-off is central to the interpretation of the model. The carbon-aware policy does not reduce emissions without operational consequences. It shifts selected replenishment decisions from later expedited orders toward earlier regular orders. This reduces emergency transport but increases the time during which some units remain in inventory. The suitability of the policy therefore depends on whether the avoided shortage, emergency-order, and emission consequences justify the additional holding cost. Earlier carbon-aware inventory studies mainly represented emissions through cost coefficients, carbon prices, or explicit regulatory

constraints [1-5]. Within such formulations, the environmental parameter modifies the economic attractiveness of a particular order quantity.

The proposed approach differs by allowing the anticipated emission consequence of delayed replenishment to influence the timing of the order. Sustainable inventory-routing models provide a more integrated representation of inventory and transport decisions [6–10]. These models can determine delivery timing, quantity, routing, and fleet use within a common optimization problem. Their analytical scope is broader than that of the present study, but this breadth is accompanied by higher data and computational requirements. The proposed policy focuses on a narrower operational decision: whether regular replenishment should be initiated earlier to avoid a future expedited shipment. The results also extend the literature on regular and emergency replenishment. Previous studies demonstrated that access to a faster secondary supply mode can improve cost and service performance when emergency orders are activated selectively [16, 17]. The present findings indicate that the environmental burden of the emergency mode should also be considered. In some cases, the most effective way to reduce emergency-transport emissions is not to optimize the expedited order after the shortage risk has emerged, but to modify the preceding regular-order decision. The limited difference between P2 and P3 deserves particular attention. Carbon-aware adjustment produced an additional improvement, but it did not transform the overall performance of the adaptive policy. Under SC2, P3 reduced emissions by approximately 0.51% relative to P2. Under SC3, the corresponding reduction was approximately 1.07%. These values suggest that the environmental extension should remain proportionate and should not be presented as a substitute for accurate demand and lead-time adaptation.

From a managerial perspective, the results support the use of carbon-aware replenishment in systems where emergency deliveries are frequent, expensive, and emission-intensive. Examples may include express road transport, dedicated small-load shipments, premium courier services, and emergency air freight. In such settings, the environmental and economic consequences of delayed replenishment can be substantial. Implementation does not necessarily require a complete inventory-routing optimization system. A company can begin by estimating the historical frequency, cost, shipment size, and carbon intensity of regular and expedited replenishment. These values can then be used to determine whether selected regular orders should be initiated earlier when the probability of emergency transport becomes sufficiently high. The results should nevertheless be interpreted within the limits of the simulation design. The model considers one product, one stocking location, and two replenishment modes. Transport emissions are represented through normalized factors rather than route-, vehicle-, and load-specific calculations. The findings therefore demonstrate structural policy behaviour rather than direct performance expectations for a particular company. Overall, the analysis indicates that carbon-aware adaptation is most useful when uncertainty weakens the performance of fixed inventory parameters and increases dependence on emergency transport. Its main contribution is not the elimination of the inventory–transport trade-off, but the explicit recognition that postponing a regular order can create future economic and environmental consequences that should be considered at the time of the inventory decision.

6. CONCLUSIONS

This study examined whether transport-related carbon information can improve replenishment decisions in inventory systems supplied through regular and expedited modes.

A carbon-aware adaptive policy was developed by extending a conventional adaptive (s, S) policy with an additional reorder-point adjustment. The purpose of this adjustment is to anticipate situations in which postponing a regular order would increase the likelihood of a later expedited shipment with higher cost and carbon intensity. The proposed policy was compared with a static (s, S) benchmark and a conventional adaptive policy under three operating conditions.

The first scenario represented stable demand and lead time, the second introduced higher demand variability, and the third combined demand uncertainty with stochastic regular replenishment lead times. The comparison considered total logistics cost, fill rate, average inventory, expedited-shipment frequency, and transport-related emissions. The results show that the usefulness of carbon-aware adaptation depends strongly on the level of operational uncertainty. Under stable conditions, the three policies produced nearly identical outcomes. The proposed policy reduced total cost by only 0.03% relative to the static benchmark, while its effect on emissions and service performance was negligible. This finding indicates that a properly calibrated static inventory policy may remain sufficient when demand and lead-time conditions are predictable. More substantial differences emerged under demand uncertainty.

In the second scenario, the carbon-aware policy reduced total logistics cost by 4.65%, transport-related emissions by 6.58%, and expedited shipments by 43.60% relative to the static policy. The fill rate improved by 1.12 percentage points, although average inventory increased by 7.75%. These results suggest that earlier regular replenishment can provide both economic and environmental benefits when demand variability increases the risk of late intervention.

The strongest overall effect was observed under combined demand and lead-time uncertainty. In this setting, the proposed policy reduced total logistics cost by 3.61%, expedited-shipment frequency by 28.48%, and transport-related emissions by 6.73% compared with the static benchmark. The fill rate increased from 92.35% to 94.04%, representing an improvement of 1.69 percentage points. These gains were accompanied by a 15.97% increase in average inventory. The comparison with the conventional adaptive policy provides a more nuanced interpretation. Most of the improvement resulted from updating the reorder point according to changing demand and lead-time conditions. Carbon-aware adjustment produced an additional emission reduction of approximately 0.51% under demand uncertainty and 1.07% under combined uncertainty. Its incremental effect on total cost was limited, particularly in the most uncertain scenario. Carbon awareness should therefore be interpreted as a targeted extension of adaptive inventory control rather than as a replacement for accurate demand and lead-time parameterization.

The theoretical contribution of the study lies in treating transport-related carbon intensity as operational decision information. Previous carbon-aware inventory models generally include emissions as a fixed cost, carbon price, cap, or reporting measure. In the proposed approach, the expected environmental consequence of postponing replenishment directly affects the timing of a regular order. The model therefore connects inventory adaptation with the future risk of emission-intensive emergency transport. From a managerial perspective, the findings are particularly relevant to systems in which expedited deliveries are frequent, costly, and poorly consolidated. Examples include dedicated small-load transport, premium courier services, urgent road shipments, and emergency air freight. In such environments, earlier regular replenishment may reduce both service risk and environmental burden. However, the additional inventory created by earlier ordering must be considered explicitly. A lower number of emergency shipments does not automatically imply a superior policy if

the resulting holding cost or storage requirement becomes excessive. The study has several limitations. It considers a single product, one stocking location, and two replenishment modes. Demand and lead-time conditions are represented through controlled stochastic scenarios, while transport emissions are calculated using normalized factors rather than route-, load-, and vehicle-specific data. The findings should therefore be interpreted as simulation-based evidence of policy behaviour rather than as direct estimates for a particular company.

Future research should validate the proposed policy using company-level inventory, shipment, and emissions data. Extensions could include multi-item systems, limited supplier capacity, shipment consolidation, partial backordering, perishability, and dynamic transport-emission factors. A broader inventory-routing model could also examine how carbon-aware replenishment timing interacts with vehicle routing and delivery scheduling. Overall, the results indicate that carbon-aware adaptation is most valuable when uncertainty creates a substantial risk of expensive and emission-intensive emergency replenishment. Its contribution lies in recognizing that the decision to postpone a regular order can generate future economic, service, and environmental consequences that should be considered before the shortage becomes critical.

ACKNOWLEDGEMENTS

"The research presented in this article was supported in part by Project No. 2023-1.2.4-TÉT-2023-00027, implemented with the support of the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund and financed under the 2023-1.2.4-TÉT funding scheme, and by the project Simulation and Modelling of Material Handling Systems (SIM_MFS), co-funded by the Ministry of Science, Education and Youth of the Republic of Croatia within the bilateral Hungarian–Croatian cooperation programme."

REFERENCES

- [1] Benjaafar, S., Li, Y. & Daskin, M. (2013). Carbon footprint and the management of supply chains: Insights from simple models. *IEEE Transactions on Automation Science and Engineering*, **10**(1), 99–116, <https://doi.org/10.1109/TASE.2012.2203304>
- [2] Hua, G., Cheng, T. C. E. & Wang, S. (2011). Managing carbon footprints in inventory management. *International Journal of Production Economics*, **132**(2), 178–185, <https://doi.org/10.1016/j.ijpe.2011.03.024>
- [3] Chen, X., Benjaafar, S. & Elomri, A. (2013). The carbon-constrained EOQ. *Operations Research Letters*, **41**(2), 172–179, <https://doi.org/10.1016/j.orl.2012.12.003>
- [4] Battini, D., Persona, A. & Sgarbosa, F. (2014). A sustainable EOQ model: Theoretical formulation and applications. *International Journal of Production Economics*, **149**, 145–153, <https://doi.org/10.1016/j.ijpe.2013.06.026>
- [5] Sebatjane, M. (2025). Sustainable inventory models under carbon emissions regulations: Taxonomy and literature review. *Computers & Operations Research*, **173**, 106865, <https://doi.org/10.1016/j.cor.2024.106865>
- [6] Coelho, L. C., Cordeau, J.-F. & Laporte, G. (2014). Thirty years of inventory routing. *Transportation Science*, **48**(1), 1–19, <https://doi.org/10.1287/trsc.2013.0472>
- [7] Soysal, M., Bloemhof-Ruwaard, J. M., Haijema, R. & van der Vorst, J. G. A. J. (2015). Modeling an inventory routing problem for perishable products with environmental considerations and demand uncertainty. *International Journal of Production Economics*, **164**, 118–133, <https://doi.org/10.1016/j.ijpe.2015.03.008>

- [8] Cheng, C., Yang, P., Qi, M. & Rousseau, L.-M. (2017). Modeling a green inventory routing problem with a heterogeneous fleet. *Transportation Research Part E: Logistics and Transportation Review*, **97**, 97–112, <https://doi.org/10.1016/j.tre.2016.11.001>
- [9] Micheli, G. J. L. & Mantella, F. (2018). Modelling an environmentally-extended inventory routing problem with demand uncertainty and a heterogeneous fleet under carbon control policies. *International Journal of Production Economics*, **204**, 316–327, <https://doi.org/10.1016/j.ijpe.2018.08.018>
- [10] Soysal, M., Çimen, M., Belbağ, S. & Toğrul, E. (2019). A review on sustainable inventory routing. *Computers & Industrial Engineering*, **132**, 395–411, <https://doi.org/10.1016/j.cie.2019.04.026>
- [11] Cui, Z., Long, D. Z., Qi, J. & Zhang, L. (2023). The inventory routing problem under uncertainty. *Operations Research*, **71**(1), 378–395, <https://doi.org/10.1287/opre.2022.2407>
- [12] Feng, Y., Che, A. & Tian, N. (2024). Robust inventory routing problem under uncertain demand and risk-averse criterion. *Omega*, **127**, 103082, <https://doi.org/10.1016/j.omega.2024.103082>
- [13] Axsäter, S. (2011). Inventory control when the lead-time changes. *Production and Operations Management*, **20**(1), 72–80, <https://doi.org/10.1111/j.1937-5956.2010.01125.x>
- [14] Cardin, O., Trentesaux, D., Thomas, A., Castagna, P., Berger, T. & Bril El-Haouzi, H. (2017). Coupling predictive scheduling and reactive control in manufacturing hybrid control architectures: State of the art and future challenges. *Journal of Intelligent Manufacturing*, **28**(7), 1503–1517, <https://doi.org/10.1007/s10845-015-1139-0>
- [15] Boute, R. N., Gijsbrechts, J., van Jaarsveld, W. & Vanvuchelen, N. (2022). Deep reinforcement learning for inventory control: A roadmap. *European Journal of Operational Research*, **298**(2), 401–412, <https://doi.org/10.1016/j.ejor.2021.07.016>
- [16] Tagaras, G. & Vlachos, D. (2001). A periodic review inventory system with emergency replenishments. *Management Science*, **47**(3), 415–429, <https://doi.org/10.1287/mnsc.47.3.415.9770>
- [17] Poormoaiied, S., Atan, Z. & van Woensel, T. (2022). Quantity-based emergency shipment policies. *IIE Transactions*, **54**(12), 1186–1198, <https://doi.org/10.1080/24725854.2022.2041773>
- [18] Nagy, G., Umetaliev, A. & Szentesi, S. (2026). A simulation-based integrated decision-support framework for auditable green logistics. *Logistics*, **10**(5), Article 98, <https://doi.org/10.3390/logistics10050098>