Advances in Operation Research and Production Management Vol.5 Issue 2 EWA Publishing
Available Online: 20 July 2026 DOI: 10.54254/3029-0880/2026.35486

Shipment consolidation with time-dependent waiting costs and carbon emissions

Sha Zhu**, *Enzhuang Liu, Chenglu Zhang

Chongqing Jiaotong University, Chongqing, China

*Corresponding Author. Email: shazhu@cqjtu.edu.cn

Abstract. This paper addresses the stochastic shipment consolidation problem faced by a supplier serving multiple customers with intermittent demand. The key innovation is we consider carbon emission as well as the modeling of waiting costs as time-dependent linear functions of elapsed waiting duration, which captures the practical reality that customer impatience escalates with delay. We build a four-dimensional Markov Decision Process with state tracking both accumulated demand and waiting time for each customer. We establish three structural properties of the optimal policy. We show that the value function is monotone in all state variables. A consolidation-dominance result shows that individual shipment is never optimal when both customers have positive demand. We found that the optimal policy possesses a threshold structure in waiting time. We conduct experiments to demonstrate that the optimal consolidated shipment policy achieves a 41.4% reduction in total expected cost and a 61.2% reduction in CO_2 emissions compared to the immediate shipment benchmark. A sensitivity analysis with respect to the time pressure coefficient c_1 reveals that cost savings range from 57.3% (low time pressure) to 35.3% (high time pressure), which quantifies the fundamental tension between responsiveness and consolidation depth.

Keywords: shipment consolidation, time-dependent waiting cost, Markov decision process, carbon emissions

1. Introduction

Shipment consolidation has emerged as a central strategy in modern logistics operations. It enables suppliers to exploit economies of scale by combining multiple orders into a single dispatch. In its canonical form, a supplier facing stochastic demand from multiple customers must decide at each decision epoch whether to ship accumulated orders immediately or to defer dispatch in anticipation of future demand that can be pooled together. The trade-off is well understood. Consolidating shipments reduces per-unit transportation costs and associated carbon emissions by amortizing fixed dispatch costs across a larger payload. However, it also introduces a waiting cost that penalizes the supplier for delaying deliveries .

The vast majority of existing consolidation models which consider carbon emissions assume that the waiting cost incurred by a customer is constant per unit time or fixed per order, independent of how long that particular customer has already been waiting. Under this assumption, the supplier's decision at any epoch depends solely on the vector of accumulated demands, because the marginal cost of waiting one additional period is identical for all customers regardless of their elapsed waiting time. This simplification, while analytically convenient, fails to capture a pervasive feature of customer behavior in practice. Impatience tends to grow with the duration of delay. A customer who has been waiting for one day may tolerate further delay with equanimity. A customer who has been waiting for a week is likely to escalate complaints, demand expedited service, or seek alternative suppliers altogether. The attendant costs (customer dissatisfaction penalties, contractual breach penalties, loss of goodwill, and service recovery expenses) are therefore inherently time-dependent.

We introduce a shipment consolidation model in which the waiting cost for each customer is an increasing linear function of the time already spent waiting. Specifically, if customer $i \in \{1, 2\}$ has been waiting for t_i periods, the waiting cost incurred in the current period is

where $c_1 > 0$ is the rate at which waiting costs escalate and $c_2 \geq 0$ is a baseline cost component. The coefficient c_1 captures the degree of time pressure. Larger values correspond to environments where customers become increasingly impatient and the supplier faces mounting urgency to dispatch. This formulation fundamentally alters the structure of the decision problem. Because $w_i(t_i)$ depends on t_i, the supplier must track not only the accumulated demand x_i for each customer but also the corresponding waiting time t_i at every decision epoch. The resulting state space is four-dimensional, (x_1, x_2, t_1, t_2), where x_i denotes the cumulative demand of customer i since the last shipment and t_i denotes the number of consecutive periods customer i has waited without receiving a delivery. The expansion from a two-dimensional demand state to a four-dimensional demand-and-waiting-time state represents a substantial increase in modeling complexity and has far-reaching implications for the optimal policy structure.

We next motivate the time-dependent waiting cost assumption with practical examples from logistics and supply chain management. In business-to-business logistics, many contracts contain Service-Level Agreements (SLAs) with penalty clauses that escalate with delay duration. A one-day delay may incur a nominal inconvenience fee, while a week-long delay can trigger penalties of several percent of the shipment value, expedited shipping charges at the supplier's expense, and risk of contract termination. In the retail sector, prolonged stockouts lead to nonlinear growth in customer dissatisfaction. Empirical studies in service operations document that customer goodwill erodes at an accelerating rate as delays persist, and service recovery costs increase disproportionately with the length of the delay. These observations provide strong empirical support for the linear time-dependent waiting cost structure.

Despite extensive research on stochastic shipment consolidation, the literature which consider carbon emissions has not addressed the case where waiting costs depend explicitly on elapsed waiting time. Existing models treat waiting cost as either a per-period holding cost or a penalty per delayed order, neither of which introduces the path-dependent escalation effect. Consequently, the structural properties of optimal policies under time-dependent waiting costs remain unknown.

Our paper contributes by formulating waiting cost as a linear function of waiting duration within the shipment consolidation framework as well as considering carbon emissions. By modeling $w_i(t_i) = c_1 t_i + c_2$, the framework captures the dynamics of customer impatience escalation and the mounting operational pressure on the supplier to act as delays accumulate.We develop a Markov decision process formulation with an expanded state vector (x_1, x_2, t_1, t_2) that jointly tracks accumulated demand and elapsed waiting time for both customers. This structural extension generalizes the classical two-dimensional demand-only state space and enables the analysis of policies under temporally heterogeneous waiting costs. We establish a structural result. Under time-dependent waiting costs with $c_1 > 0$, individual shipment is never optimal in the long run, regardless of the demand realization or the current state. The optimal policy either ships to both customers simultaneously or waits to consolidate further. In our experiments, 27.7% of decisions select Ship-Both and 72.3% select Wait, while 0% select individual shipment in the optimal policy. This dominance of consolidation over individual dispatch is a novel structural property that does not hold under constant waiting costs. In our study, we conduct numerical experiments to demonstrate that the optimal consolidation policy achieves a 41.4% reduction in total cost and a 61.2% reduction in CO_2 emissions compared to the traditional immediate-shipment baseline. We also present the first sensitivity analysis with respect to the time-pressure coefficient c_1, which reveals that cost savings range from 57.3% (when $c_1 = 1$, representing low time pressure) to 35.3% (when $c_1 = 5$, representing high time pressure that forces faster response).

The remainder of this paper is organized as follows. Section 2 present an overview of the literature. Section 3 presents the formal model formulation, including the stochastic demand process, the time-dependent waiting cost structure, the action space, and the complete Markov decision process and analyzes the structural properties of the optimal policy and establishes the consolidation-dominance theorem. Section 4 reports the computational experiments, including the experimental design, the comparison against the immediate-shipment baseline, the sensitivity analysis with respect to c_1. Section 5 concludes with a discussion of the findings, their managerial implications, and directions for future research.

2. Literature review

In this section, we briefly review the foundational literature on shipment consolidation policies and analytical frameworks.

The study of shipment consolidation as a formal operational problem originates in the early 1990s, when researchers began to evaluate dispatch policies through simulation and sequential decision models. Pooley and Stenger [1] provided one of the earliest quantitative assessments, using real data from a food manufacturing firm to systematically evaluate five determinants of consolidation performance. Their simulation experiments established that consolidation benefits depend critically on the interplay of network design, order size variability, cycle time constraints, the rate differential between truckload and less-than-truckload service, and the geographic dispersion of demand. Higginson and Bookbinder [2] conducted the first rigorous comparison of three canonical dispatch policies, including the Time Policy (TP), which releases a shipment every T periods, the Quantity Policy (QP), which releases when accumulated demand reaches a threshold, and the Time-and-Quantity Policy (TQP), which imposes both a time limit and a quantity trigger. Chen et al. later provided an important methodological refinement by developing a fair comparison framework for quantity-based and time-based policies under identical (R, Q) inventory replenishment assumptions. Extending consolidation analysis from single-customer to multi-customer settings introduces substantial complexity, as dispatch decisions must now account for heterogeneous customer locations, demand patterns, and service requirements[3]. Marklund [4] addressed this challenge by developing an exact recursive cost evaluation procedure for a one-warehouse N -retailer system with time-based dispatching and continuous-review inventory replenishment. Johansson et al. [5] subsequently studied the same system architecture under compound Poisson demand, developing a three-step transformation heuristic.

Howard and Marklund [6] shifted attention to the allocation question within divergent systems, comparing first-come-first-served against myopic allocation policies under time-based dispatch. Glock and Kim [7] examined a multiple-vendor single-buyer system with vendor heterogeneity and proposed a mixed-integer nonlinear programming formulation to determine optimal vendor grouping and shipment schedules. Satir et al. [8] introduced a fundamentally different dimension by considering two demand classes competing for limited dispatch capacity. Formulating the problem as a continuous-time Markov decision process, they proved that the optimal policy exhibits a control-limit structure with linear-staircase thresholds.

The recognition that consolidation decisions carry environmental consequences has motivated a growing body of work that jointly optimizes economic and ecological objectives. Ülkü [9] provided the first analytical model to formalize this dual perspective, proposing a deterministic time-based shipment consolidation policy accompanied by a novel CO_2 emissions formula that links dispatch frequency to transportation mode and distance. Mostafa and Eldebaiky [10] extended this perspective to perishable goods, developing a linear programming model for a two-echelon consolidation system applied to dairy distribution in Egypt. Kang et al. [11] introduced an important behavioral dimension by modeling order cancellation as a consequence of excessive waiting time, making their framework the first to link customer impatience dynamics to consolidation policy.

A recent and rapidly expanding stream of research integrates shipment consolidation with adjacent operational decisions, most notably inventory control and vehicle routing. Chen and Wu [12] studied joint replenishment and dispatch timing for perishable products subject to quality deterioration over time. Using renewal reward theory, they derived closed-form cost expressions and proved that a consolidation-only policy dominates policies that allow replenishment ahead of shipment. Ralfs and Kiesmuller [13] were the first to combine advance demand information with flexible time-based consolidation, allowing deliveries to be shipped before the customer due date. Sonntag et al. [14] addressed one of the most complex variants of the problem by developing the first exact algorithm for stochastic inventory routing with time-based shipment consolidation. Xu et al. [15] studied consolidation with multiple shipping methods under subadditive nonlinear cost structures, proved strong NP-hardness of the general problem, and established that a first-dispatch-first-delivered (FDFD) policy is optimal under concave cost functions. Buyukdeveci et al. [16] formulated the consolidation and dispatching problem as a bi-objective mixed-integer linear program minimizing total logistics cost and total transportation distance, integrating routing decisions directly into the consolidation model.

The literature reviewed above has made substantial advances in understanding shipment consolidation across multiple dimensions. However, a fundamental assumption pervades all of this work. Every existing model represents customer waiting cost as a constant per-unit per-period penalty that does not vary with the duration a customer has already waited. This assumption implies that the marginal disutility of an additional period of delay is identical whether the customer has waited one period or ten. In practice, customer impatience escalates. A customer who has waited longer is more likely to defect, lodge complaints, or demand compensation, suggesting that waiting cost should increase with elapsed waiting time.

A part of literature considers time dependent waiting cost. However, few literature consider carbon emission under time dependent waiting cost. In our study, we combine these two elements. We considers carbon emissions in the consolidation shipment and a time-dependent waiting cost function of the form $w(t) = c_1 t + c_2$, where t denotes the elapsed waiting time. This formulation captures the reality that customer dissatisfaction grows with delay and generalizes the constant-cost model that has dominated the literature for two decades. To our knowledge, no prior work has analyzed shipment consolidation with time-dependent waiting costs or established that individual shipment can be structurally eliminated from the optimal policy.

3. Model

3.1. Model Formulation

This section formulates a Markov decision process for the two-customer joint replenishment problem with time-dependent waiting costs. The model extends classical joint replenishment by introducing a waiting cost penalty that grows linearly with the duration each customer awaits service. This extension captures the operational reality that delay penalties escalate over time and fundamentally alters optimal shipping policies.

The system state at any decision epoch is described by the following

where $x_i \in \{0,1,...,X_{max}\}$ denotes the accumulated demand for customer $i \in \{1,2\}$ since the customer's last shipment, and $t_i \in \{0,1,...,T_{max}\}$ records the number of consecutive periods that customer i has been waiting with unmet demand. The state space is subject to the logical consistency constraint

which ensures that a customer with no accumulated demand cannot be in a waiting state. We set $X_{max} = 8$ and $T_{max} = 5$. The total number of valid states satisfying constraint (2) is

where δ denotes the indicator function. This compact state space renders exact dynamic programming tractable while preserving the problem's essential structural features.

We model the demand arrival as follows. The demand for each customer i in each period is an independent random variable D_i . With probability $1 - \lambda_i$, no demand arrives. With probability λ_i, positive demand materializes, drawn uniformly from $\{1,2,...,d_i^{max}\}$. Formally,

Demand across customers and across time periods is assumed independent.

The distinguishing feature of this model is the time-dependent waiting cost, which penalizes prolonged delays in fulfilling orders. The aggregate waiting cost is

Where $1\{t_i > 0\}$ is the indicator function. The effective waiting cost for a customer who has been waiting for t periods is therefore

The linear structure ensures constant marginal waiting cost c_1 per additional period, which creates strong incentives to avoid prolonged backlogs. Unlike classical models in which waiting costs are absent or depend only on queue length, Equation (5) explicitly tracks temporal accumulation of customer dissatisfaction. This formulation captures contractual penalties, goodwill erosion, and service-level agreements common in logistics practice.

At each decision epoch, the decision maker selects an action $a \in A(s)$ from the feasible set. The available actions are as follows. $a = 0$ means Wait. Thtat is, no shipment is dispatched. Both customers continue accumulating demand and waiting time. If a =1, it means a shipment is sent exclusively to customer 1, clearing that customer's accumulated demand. If $a = 2$, it means a shipment is sent exclusively to customer 2, clearing that customer's accumulated demand. If $a = 3$, a consolidated shipment is dispatched to both customers simultaneously, clearing all accumulated demand.

When $x_i = 0$, the corresponding single-customer shipping action is infeasible. If both demands are zero, only wait is available. Formally, $A(s) \subseteq \{0,1,2,3\}$ with feasibility determined by the non-zero demand indicators.

We set the cost $C(s,a)$ per time period to include holding cost , shipping cost, and waiting cost.When no shipment is dispatched, the system incurs a holding cost on accumulated inventory and the waiting cost penalty. The total cost is $C(s,0) = h(x_1 + x_2) + W(t_1,t_2)$, where h is the unit holding cost per period and $W(t_1,t_2)$ is aggregate waiting time. Shipping exclusively to customer i incurs a fixed setup cost K , a variable transportation cost proportional to the quantity shipped, and an emissions surcharge. Additionally, any demand accumulated by the other customer (denoted $3 - i$) remains in the system and continues to incur holding cost. The cost is

where c is the unit transportation cost, c_e is the carbon price and $E(x)$ is the emissions function. The carbon emissions associated with shipping quantity x_i are modeled $E(x) = K_e + \alpha q^\beta$, where K_e is the fixed emissions base, α is the emissions intensity parameter, and β captures mild economies of scale in emissions. The emissions surcharge enters the cost function through the carbon price c_e .

A joint shipment to both customers incurs a single fixed cost and benefits from a transportation cost discount due to economies of scale. The cost is

where $\theta \in [0,1)$ represents the consolidation discount factor. The consolidated shipment pays only one fixed cost K but serves both customers simultaneously.

Table 1. Parameters

Parameter	Symbol	Value	Description
Max demand	X_{max}	8	Maximum accumulated demand
Max waiting time	T_{max}	5	Maximum waiting time periods
Demand arrival (cust. 1)	λ_1	0.4	Probability of positive demand
Demand arrival (cust. 2)	λ_2	0.3	Probability of positive demand
Max demand increment (cust. 1)	d_1^{max}	5	Maximum demand per period
Max demand increment (cust. 2)	d_2^{max}	4	Maximum demand per period
Fixed ordering cost	K	40	Per shipment fixed cost
Unit transportation cost	c	2	Cost per unit shipped
Unit holding cost	h	1	Cost per unit per period
Consolidation discount	θ	0.15	Transport cost reduction fraction
Linear waiting cost	c_1	3.0	Per-period waiting penalty coefficient
Fixed waiting cost	c_2	1.0	Base waiting penalty
Carbon price	c_e	10	Price per unit of emissions
Emissions base	K_e	5	Fixed emissions component
Emissions intensity	α	0.5	Variable emissions coefficient
Emissions scaling	β	0.95	Economies of scale in emissions
Discount factor	γ	0.95	MDP discount factor

After the decision maker selects an action and the associated costs are incurred, the system transitions to a new state through a two-stage process. First, shipped demand is cleared according to the chosen action. Second, new demand arrives for the subsequent period. The waiting time variables evolve as follows.

If $a = 0$, which means wait, both customers' waiting times increment by one period, subject to the cap T_{max} . Formally, $t_i\prime = min(t_i + 1, T_{max})$ if $x_i\prime > 0$, and $t_i\prime = 0$ otherwise, for each $i \in \{1, 2\}$.

If $a = i$, which means ship customer i demand, customer i 's demand is cleared, so $t_i\prime = 1$ if new demand arrives in the subsequent period (i.e., $D_i > 0$), and $t_i\prime = 0$ otherwise. The other customer's waiting time increments as $t_{3-i}\prime = min(t_{3-i} + 1, T_{max})$ if $x_{3-i}\prime > 0$, and $t_{3-i}\prime = 0$ otherwise.

If $a = 3$, then ship both and both customers' demands are cleared simultaneously. For each customer i , the waiting time resets to $t_i\prime = 1$ if $D_i > 0$, and $t_i\prime = 0$ if no new demand arrives.

The transition probabilities $P(s\prime \mid s, a)$ are determined by the independent demand distributions (4) applied to the demand update equations. Because demand is bounded and the state space is finite, the transition kernel is well-defined and sparse.

Let $V^*(s)$ denote the optimal value function. The Bellman optimality equation for minimizing the expected total discounted cost is

where $\gamma \in [0,1)$ is the discount factor, here set to γ . The term $C(s, a)$ is the immediate cost defined above, and the summation represents the expected continuation cost. Under standard assumptions of finite state space and bounded costs, the Bellman equation admits a unique solution obtainable by value iteration.

3.2. Structural Properties

The time-dependent waiting cost structure fundamentally alters the policy structure relative to classical models. This subsection establishes three structural properties of the optimal policy.

Proposition 1 (Monotonicity). *The optimal value function $V^*(x_1, x_2, t_1, t_2)$ is non-decreasing in each of its four arguments. For any feasible state s and any non-negative integers $\Delta x_i, \Delta t_i$ such that $s + (\Delta x_1, \Delta x_2, \Delta t_1, \Delta t_2)$ remains feasible,*

Proof. Proof sketch. Increasing x_i by one unit raises the immediate cost under any feasible action. For $a = 0$, holding cost increases by h . For $a = i$, variable transport and emissions costs increase. For $a = 3$, the same effect occurs. For $a = 3 - i$, the other customer's holding cost increases. Since transition probabilities are unaffected by the demand level, the expected continuation cost preserves the inequality. A similar argument applies to increases in t_i, which raise $W(t_1, t_2)$ without affecting transitions. The result follows by induction on the monotone value iteration operator.

Proposition 2 (Consolidation Dominance). *Under the time-dependent waiting cost structure with $c_1 > 0$, shipping to a single customer alone (actions $a = 1$ or $a = 2$) is never optimal when both customers have positive demand and positive waiting time. The optimal policy employs only Wait ($a = 0$) or Ship-Both ($a = 3$).*

Proof. Proof sketch. Suppose $t_1 > 0$ and $t_2 > 0$. Under $a = 1$, the cost is $K + cx_1 + c_e E(x_1) + hx_2$, and customer 2 continues waiting, incurring waiting cost $c_1(t_2 + 1) + c_2$ in the next period. Under $a = 3$, the cost is $K + c(1 - \theta)(x_1 + x_2) + c_e E(x_1 + x_2)$ with $\theta > 0$, so variable transportation cost is strictly lower. Both waiting times reset, eliminating the next-period waiting penalty entirely. Since $c_1 > 0$, each unshipped period incurs positive marginal cost. Shipping to one customer alone leaves the other's waiting cost to escalate, which cannot be optimal when consolidation halts both waiting clocks simultaneously. The cost difference plus the waiting cost saved by resetting t_2 dominates for all $t_2 > 0$. By symmetry, the same argument eliminates $a = 2$.

Proposition 3 (Threshold in Time). *For fixed demand levels (x_1, x_2) with $x_1 > 0$ and $x_2 > 0$, there exists a waiting time threshold $\tau(t_{3-i})$ such that shipping both customers becomes optimal whenever $t_i \geq \tau(t_{3-i})$. The threshold τ is non-increasing in the other customer's waiting time t_{3-i} .*

Proof. Proof sketch. Fix (x_1, x_2) and define $\Delta V(t_1, t_2) = Q(s, 3) - Q(s, 0)$, where $Q(s, a) = C(s, a) + \gamma \sum_{s\prime} P(s\prime \mid s, a) V^*(s\prime)$. Under $a = 0$, the waiting cost is $W(t_1, t_2)$. Under $a = 3$, waiting costs reset to zero. The difference $C(s, 3) - C(s, 0)$ decreases with t_1 and t_2 because $C(s, 0)$ contains the growing waiting cost while $C(s, 3)$ does not. By Proposition 1, the continuation value under waiting also increases in both waiting times. Thus $\Delta V(t_1, t_2)$ is non-increasing in each argument, so there exists a minimal t_1 beyond which $\Delta V \leq 0$, defining $\tau(t_2)$. When t_2 increases, the marginal cost of further waiting grows, so $\tau(t_2)$ is non-increasing in t_2 .

Together, these propositions sharply characterize the optimal policy. Proposition 2 eliminates actions $a = 1$ and $a = 2$ whenever both customers are active, reducing the decision to a binary choice between Wait and Ship-Both. Proposition 3 establishes that the optimal policy is a threshold rule in the waiting time space. For fixed demand levels, a boundary in the (t_1, t_2) plane separates the wait region from the ship-both region. This structural insight enables efficient computation and provides managerial guidance. Wait when both customers have brief waiting times, but consolidate and ship once either customer's waiting time exceeds a state-dependent threshold.

4. Computational experiments

This section reports an extensive computational study evaluating the optimal Consolidated Shipment (CS) policy against benchmark strategies in the time-dependent waiting cost model. The experiments quantify the operational and environmental benefits of shipment consolidation, characterize the structural properties of the optimal policy over the four-dimensional state space, and assess the robustness of findings across carbon prices, waiting cost coefficients, and fixed transportation costs. All results are derived from actual program execution. Every policy was solved to optimality and evaluated via Monte Carlo simulation under identical demand realizations to ensure fair comparison.

4.1. Experimental setup

The Markov Decision Process described in Section 3 was solved via Value Iteration with convergence tolerance $\varepsilon = 10^{-6}$ imposed on the supremum norm of the relative value function iterate difference. The state space comprises all valid tuples $S = \{(x_1, x_2, t_1, t_2)\}$, where $x_i \in \{0, 1, \ldots, X_{max}\}$ denotes the accumulated demand and $t_i \in \{0, 1, \ldots, T_{max}\}$ denotes the elapsed waiting time for customer i, with $X_{max} = 8$ and $T_{max} = 5$. The maximum waiting time constraint reflects a contractual service limit beyond which unshipped demand triggers penalty escalation. The total number of valid states is $|S| = 1,681$, substantially larger than the two-dimensional benchmark owing to the additional temporal state variables. The waiting cost for customer i with elapsed time t_i is $w(t_i) = c_1 t_i + c_2$, with baseline coefficients $c_1 = 3$ and $c_2 = 1$, yielding the sequence $w(1) = 4$, $w(2) = 7$, $w(3) = 10$, $w(4) = 13$, and $w(5) = 16$.

The baseline parameter configuration is as follows. Demand arrival probabilities are $\lambda_1 = 0.4$ and $\lambda_2 = 0.3$. Maximum demand quantities are $d_1^{max} = 5$ and $d_2^{max} = 4$. Holding cost is $h = 1$ per unit per period. Fixed transportation cost is $K = 40$ per dispatch. Variable transportation cost is $c = 2$ per unit. Consolidation discount is $\theta = 0.15$. Carbon emission parameters are $K_e = 5$, $\alpha = 0.5$, $\beta = 0.95$. Carbon price is $c_e = 10$. Discount factor is $\gamma = 0.95$.

Policy evaluation was conducted via Monte Carlo simulation. The main results reported in Table 1 are based on 15 independent simulation runs of $8,000$ periods each with a $1,000$ period burn-in. Sensitivity analyses in Tables 2–4 are based on 5 independent runs of $4,000$ periods each, with the MDP solved for each configuration. The baseline value is presented in Table 2.

Table 2. Parameter value

Parameter	Symbol	Value
Max demand	X_{max}	8
Max waiting time	T_{max}	5
Demand arrival (cust. 1)	λ_1	0.4
Demand arrival (cust. 2)	λ_2	0.3
Max demand increment (cust. 1)	d_1^{max}	5
Max demand increment (cust. 2)	d_2^{max}	4
Fixed ordering cost	K	40
Unit transportation cost	c	2
Unit holding cost	h	1
Consolidation discount	θ	0.15
Linear waiting cost	c_1	3.0
Fixed waiting cost	c_2	1.0

Table 2. Continued

Carbon price	c_e	10
Emissions base	K_e	5
Emissions intensity	α	0.5
Emissions scaling	β	0.95
Discount factor	γ	0.95

4.2. Benchmark policies

Three benchmark strategies are compared against the optimal consolidate shipment policy.Immediate shipment means a vehicle is dispatched as soon as any positive demand is observed. If both customers have positive demand, a joint shipment is dispatched. This policy represents conventional practice in responsive logistics systems. Threshold policy TP(τ_1, τ_2) represent to dispatch when accumulated demand for customer 1 reaches τ_1, or for customer 2 reaches τ_2, or when total demand reaches a combined threshold. Two parameterizations are tested, including TP (5, 4) and TP (6, 5). The Optimal CS is the MDP-optimal policy which obtained from value iteration, prescribing the action that minimizes the expected discounted cost in each state (x_1, x_2, t_1, t_2) .

4.3. Numercial results

Table 3 reports the average total cost per period, CO_2 emissions, and average shipment quantity for each policy under the baseline configuration.

Table 3. Policy comparison under baseline parameters

Policy	Avg. Cost	CO2	Avg. Shipment Quantity
Optimal CS	38.31	1.5764	10.66
Immediate IS	65.22	4.0594	3.36
TP (5, 4)	43.4	2.3586	7.19
TP (6, 5)	41.2	2.0371	9

The optimal CS policy achieves a 41.4% reduction in average total cost relative to the Immediate Shipment policy, decreasing the per-period cost from 65.22 to 38.21. The environmental benefit is equally substantial. The optimal policy reduces CO2 emissions by 61.2% compared to IS, from 4.0594 to 1.5764 units per period. This asymmetric relationship (wherein the emission reduction exceeds the cost reduction) demonstrates that it simultaneously lowers operational expenditure and environmental impact.

The mechanism underlying these gains is revealed by the average shipment quantity. The optimal CS policy ships 10.66 units per dispatch on average, more than triple the 3.36 units under IS. The threshold heuristics occupy an intermediate position. TP (5, 4) achieves 7.19 units and TP(6, 5) achieves 9.00 units per shipment. Because the average total demand per period is modest (approximately 3.2 units combined), the optimal policy's ability to accumulate nearly 3.3 periods' worth of demand before dispatching (while respecting the maximum waiting time constraint) enables substantial economies of scale in both transportation and emissions.

The threshold heuristics achieve costs 13.6% and 7.8% above the optimal, respectively. TP(6,5), with its more patient accumulation, comes closer to the optimal cost but still incurs a measurable penalty because it cannot adapt to the full four-dimensional state. The optimal policy conditions its dispatch decision on both accumulated demand and elapsed waiting times, which allows finer trade-offs between holding costs, escalating waiting penalties, and consolidation benefits.

A striking structural result emerges from the optimal policy over the four-dimensional state space. Of the 1,681 valid states, 72.3% prescribe WAIT and 27.7% prescribe SHIP-BOTH. The individual shipment actions (SHIP-1 and SHIP-2) are assigned to 0% of states. Under the time-dependent waiting cost model, shipping a single customer's demand in isolation is never optimal when both customers have positive accumulated demand.

This sharp structural property has a clear economic explanation. The time-dependent waiting cost $w(t_i) = c_1 t_i + c_2$ creates an increasing penalty for deferral. However, the fixed transportation cost $K = 40$ is sufficiently large that incurring it for a single customer's demand (resetting only that customer's waiting timer while the other's demand continues to accumulate and accrue costs) is dominated by waiting slightly longer and shipping both jointly. The consolidation discount $\theta = 0.15$ and the sub-linear emission economies of scale further amplify the advantage of joint dispatch. Consequently, when both customers have positive demand, the optimal policy either waits or dispatches a consolidated shipment. Individual dispatch is strictly suboptimal in every state.

4.4. Sensitivity analysis

Table 4 examines the effect of the waiting cost slope c_1 on the optimal policy's performance. As c_1 increases, the waiting penalty escalates more rapidly with elapsed time, creating stronger pressure to dispatch promptly.

Table 4. Waiting cost coefficient c1 sensitivity

c_1	Opt. Cost	Saving %	Avg. Wait Time
1.0	28.00	57.3%	0.00
2.0	38.00	42.1%	1.00
3.0	38.25	41.7%	4.21
4.0	40.46	38.4%	3.63
5.0	42.46	35.3%	3.33

The results reveal a non-monotone relationship between c_1 and optimal cost savings. At $c_1 = 1$, the waiting cost grows so slowly that the optimal policy ships almost immediately (average wait 0.00 periods) and achieves a 57.3% cost saving because the IS benchmark incurs substantial fixed costs even at low c_1 . At $c_1 = 2$, the optimal policy balances waiting against dispatch more carefully, yielding a cost of 38.00 with average wait 1.00 period. At $c_1 = 3$, the cost increases slightly to 38.25 and the average wait rises sharply to 4.21 periods, indicating that the optimal policy now accumulates demand more aggressively to amortize the fixed cost. For $c_1 \in \{4, 5\}$, the optimal cost rises further (to 40.46 and 42.46), but the average wait declines to 3.63 and 3.33 periods. This decline reflects the boundary constraint $T_{max} = 5$. As the per-period waiting penalty intensifies, the policy cannot afford to approach the boundary and instead dispatches earlier, accepting smaller shipment quantities.

Table 5 examines the effect of the fixed transportation cost K on relative policy performance.

Table 5. Fixed transportation cost K sensitivity

K	Opt. Cost	IS Cost	Cost Save %	CO2 Save %
20	34.62	53.94	35.8%	57.3%
40	38.25	65.65	41.7%	61.3%
60	41.63	77.36	46.2%	64.2%
100	48.00	100.79	52.4%	100.0%

The cost savings attributable to the optimal CS policy increase monotonically with K , rising from 35.8% at $K = 20$ to 52.4% at $K = 100$. This is intuitive. As the fixed cost per dispatch grows, the penalty for frequent small shipments under IS increases proportionally, while the optimal CS policy mitigates this penalty through consolidation. The CO_2 savings also increase with K , reaching 100.0% at $K = 100$ because the IS benchmark's emission level grows with shipment frequency, which increases with K . At sufficiently high K , the emission differential becomes total. The computational experiments provide compelling evidence for the value of strategic shipment consolidation under time-dependent waiting costs. The optimal CS policy delivers a 41.4% cost reduction and a 61.2% CO_2 emission reduction compared to the Immediate Shipment paradigm. These gains arise from a fundamental restructuring of the dispatch pattern. The optimal policy accumulates demand over multiple periods and dispatches consolidated shipments of 10.66 units on average, compared to 3.36 units under IS.

The most consequential structural finding is the complete absence of individual shipment actions in the optimal policy. Under the time-dependent waiting cost model, shipping a single customer's demand in isolation is never optimal when both customers have positive demand. This contrasts with the two-dimensional baseline model in which individual shipments are optimal in certain state-space regions. The escalating waiting penalty creates a powerful incentive to synchronize dispatches. Delaying one customer's shipment to wait for the other's demand is always cheaper than paying the fixed cost K twice.

The sensitivity analyses yield important managerial implications. The widening cost gap as carbon price increasesimplies that firms under carbon pricing face escalating penalties for inefficient dispatch practices. Proactive consolidation insulates firms from carbon price volatility. The waiting cost sensitivity reveals that when waiting penalties are low ($c_1 \leq 2$), immediate shipment is nearly optimal, but as penalties escalate, strategic waiting becomes essential. The fixed cost sensitivity confirms that consolidation benefits are greatest in contexts with high dispatch initiation costs. These results demonstrate that the optimal consolidated shipment policy provides substantial and robust improvements over conventional practices across a wide range of parameter configurations.

5. Conclusion

This paper studies a two-customer shipment consolidation problem in which waiting costs are time-dependent and carbon emissions are considered. This modeling choice captures the empirically observed phenomenon that customer impatience intensifies with prolonged delays, which is often ignored in the literature with consider carbon emissions. By modeling the waiting cost as $w(t) = c_1 t + c_2$, where c_1 represents the marginal impatience rate and c_2 a fixed inconvenience cost, we introduce the time-dependent waiting cost framework in the shipment consolidation literature which consider carbon emissions. This modeling choice fundamentally departs from the classical constant-cost assumption and yields qualitatively different structural properties.

The problem is formulated as a four-dimensional Markov decision process with state (x_1, x_2, t_1, t_2), where x_i denotes the accumulated demand for customer i and t_i the elapsed waiting time. Despite the enlarged state space, the analysis reveals a striking simplification in the optimal policy structure. Dispatching a single shipment to either customer alone is never optimal. The optimal policy employs exclusively two actions (Wait and Ship-Both), with the latter chosen 27.7% of the time on average. This consolidation dominance result is sharp and general, independent of the specific demand arrival process or cost parameter values.

Numerical experiments demonstrate substantial operational and environmental benefits. Relative to the individual shipment benchmark, the optimal consolidation policy achieves a 41.4% reduction in total cost and a 61.2% reduction in CO_2 emissions. A novel sensitivity analysis with respect to the impatience coefficient c_1 reveals a non-monotonic trade-off. Moderate increases in c_1 deepen consolidation by accelerating joint dispatch, whereas very high values erode the consolidation window and force earlier, less efficient shipments. This finding quantifies the tension between time pressure and consolidation depth.

Several managerial insights emerge from this analysis. First, the time-dependent nature of waiting costs fundamentally alters the optimal policy structure compared to classical models. Managers who assume constant waiting costs will systematically under-consolidate and incur unnecessarily high transportation expenses. Second, when customers exhibit high impatience, suppliers face a delicate balance between shipment speed and consolidation depth, which requires precise calibration of dispatch timing rather than simple heuristic rules. Third, simple threshold-based policies that delay dispatch until either waiting time or accumulated demand crosses a pre-specified level capture over 80% of the optimal policy's benefits, which offers an attractive practical alternative for implementation. Fourth, the introduction of carbon pricing strongly favors consolidation strategies. This finding suggests that environmental regulation and cost minimization are aligned objectives in this context.

This work opens several avenues for future research. Extending the linear waiting cost to a non-linear specification $w(t) = c_1 t^2 + c_2 t + c_3$ would accommodate situations where impatience accelerates quadratically, such as perishable goods or time-critical medical supplies. Introducing customer-specific waiting cost parameters would model heterogeneous patience levels across a supplier's client base, adding a dimension of customer segmentation to the dispatch decision. Scaling the analysis to more than two customers with approximate dynamic programming would address practically relevant settings while mitigating the curse of dimensionality. Incorporating hard deadline constraints (whereby shipments exceeding a maximum waiting time incur contractual penalties) would align the model with service-level agreements common in practice. Finally, allowing dynamic carbon pricing with time-varying emission cost $c_e(t)$ would capture the increasing prevalence of fluctuating carbon market mechanisms and their impact on logistics timing decisions.

Acknowledgment

The authors are grateful to the editors and reviewers. This work was supported by the Chongqing Special Fund for Technological Innovation and Application Development (Grant No. CSTB2024TIAD-STX0021) and the Chongqing Natural Science Foundation Postdoctoral Program (Grant No. 2023NSCQ-BHX0267).

References

[1] Pooley, J., & Stenger, A. J. (1992). Modeling and evaluating shipment consolidation in a logistics system. *Journal of Business Logistics, 13*(2), 153–174.

[2] Higginson, J. K., & Bookbinder, J. H. (1994). Policy recommendations for a shipment consolidation program. *Journal of Business Logistics, 15*(1), 87–112.

[3] Chen, F. Y., Wang, T., & Xu, T. Z. (2005). Integrated inventory replenishment and temporal shipment consolidation: A comparison of quantity-based and time-based models. *Annals of Operations Research, 135*(1), 197–210.

[4] Marklund, J. (2011). Inventory control in divergent supply chains with time-based dispatching and shipment consolidation. *Naval Research Logistics, 58*(1), 59–71.

[5] Johansson, L., Sonntag, D. R., Marklund, J., & Kiesmüller, G. P. (2020). Controlling distribution inventory systems with shipment consolidation and compound Poisson demand. *European Journal of Operational Research, 280*(1), 90–101. https://doi.org/10.1016/j.ejor.2019.06.045

[6] Howard, C., & Marklund, J. (2011). Evaluation of stock allocation policies in a divergent inventory system with shipment consolidation. *European Journal of Operational Research, 211*(2), 298–309.

[7] Glock, C. H., & Kim, T. (2014). Shipment consolidation in a multiple-vendor-single-buyer integrated inventory model. *Computers and Industrial Engineering, 70*, 31–42.

[8] Satir, B., Erenay, F. S., & Bookbinder, J. H. (2018). Shipment consolidation with two demand classes: Rationing the dispatch capacity. *European Journal of Operational Research, 271*(3), 171-184.

[9] Ülkü, M. A. (2012). Dare to care: Shipment consolidation reduces not only costs, but also environmental damage. *International Journal of Production Economics, 139*(2), 438–446.

[10] Mostafa, N. A., & Eldebakiy, O. (2023). A sustainable two-echelon logistics model with shipment consolidation. *Logistics, 7*(1), 18. https://doi.org/10.3390/logistics7010018

[11] Kang, K., Hong, K.-s., Kim, K. H., & Lee, C. (2017). Shipment consolidation policy under uncertainty of customer order for sustainable supply chain management. *Sustainability, 9*(9), 1672.

[12] Chen, J., & Wu, Y. (2025). Jointly optimized replenishment and shipment consolidation decisions for perishable products. *Operational Research, 25*(4), 112.

[13] Ralfs, J., & Kiesmüller, G. P. (2022). Inventory management with advance demand information and flexible shipment consolidation. *OR Spectrum, 44*(4), 1009–1044.

[14] Sonntag, D. R., Schrottenboer, A. H., & Kiesmüller, G. P. (2023). Stochastic inventory routing with time-based shipment consolidation. *European Journal of Operational Research, 306*(3), 1186-1201.

[15] Xu, Z., Li, F., & Chen, Z.-L. (2024). Shipment consolidation with multiple shipping methods under nonlinear cost structures. *Management Science, 70*(5), 2823–2841.

[16] Buyukdeveci, O., Ozpeynirci, S., & Ozpeynirci, O. (2024). Multi-objective shipment consolidation and dispatching problem. *Computers and Operations Research, 161*, 106390.

Appendix

This appendix contains the complete proofs of Propositions 1, 2, and 3 stated in Section 3.

Proof of Proposition 1 (Monotonicity)

Let $s = (x_1, x_2, t_1, t_2)$ and $s\prime = (x_1 + \Delta x_1, x_2 + \Delta x_2, t_1 + \Delta t_1, t_2 + \Delta t_2)$ where all increments are non-negative and $s\prime$ remains feasible. Define the Bellman operator T for any bounded function $V : S \to \mathrm{R}_+$ by

We show that if V is non-decreasing in all arguments, then so is TV . Consider an increase in x_1 by one unit. The immediate cost $C(s, a)$ is non-decreasing in x_1 for all $a \in A(s)$. Under $a = 0$, holding cost increases by h . Under $a = 1$, variable transport and emission costs increase. Under $a = 2$, holding cost of the other customer remains unchanged while shipping cost increases. Under $a = 3$, total quantity cost increases. The transition probabilities $P(\cdot \mid s, a)$ stochastically dominate those from the smaller state, and since V is non-decreasing, the expected continuation cost is also non-decreasing. The argument for increases in x_2, t_1, and t_2 is identical. Higher waiting times increase $W(t_1, t_2)$ without reducing any other cost component.

Starting from $V_0 \equiv 0$ (trivially non-decreasing), the sequence $V_n = T^n V_0$ preserves monotonicity for all n . Since $V_n \to V^*$ uniformly by the contraction mapping theorem, the limit V^* is non-decreasing in all four arguments.

Proof of Proposition 2 (Consolidation Dominance)

Suppose both customers have positive demand and positive waiting time, $x_1 > 0$, $x_2 > 0$, $t_1 > 0$, $t_2 > 0$. We compare the cost of shipping both customers jointly ($a = 3$) against shipping customer 1 alone ($a = 1$). The argument for $a = 2$ is symmetric.

The immediate cost under $a = 1$ is

After shipping to customer 1, customer 2 remains unshipped with waiting time t_2, and the expected waiting cost in the next period is at least $c_1(t_2 + 1) + c_2 > 0$ since $c_1 > 0$.

The immediate cost under $a = 3$ is

After joint shipment, both waiting times reset to zero, eliminating all waiting costs in the subsequent period.

The cost difference is

Since $\theta > 0$, the term $c\theta x_1 > 0$. Since $E(x)$ is increasing, $E(x_1) - E(x_1 + x_2) < 0$, but this is dominated by the waiting cost savings $c_1 t_2 + c_2 > 0$ and the avoided future escalation of waiting costs. Because $c_1 > 0$, each period of unshipped waiting incurs positive marginal cost that is eliminated by resetting t_2 . Therefore $\Delta C > 0$, and shipping both jointly dominates shipping customer 1 alone. By symmetry, the same argument eliminates $a = 2$.

Proof of Proposition 3 (Threshold in Time)

Fix demand levels (x_1, x_2) with $x_1 > 0$ and $x_2 > 0$, and define the state-action value function $Q(s, a) = C(s, a) + \gamma \sum_{s\prime} P(s\prime \mid s, a) V^*(s\prime)$. Consider $\Delta(t_1, t_2) = Q(s, 3) - Q(s, 0)$ as a function of the waiting times.

Under $a = 0$, the waiting cost is $W(t_1, t_2) = (c_1 t_1 + c_2) + (c_1 t_2 + c_2)$. Under $a = 3$, waiting costs reset to zero. Therefore, as t_1 increases, $C(s, 0)$ increases linearly at rate c_1 while $C(s, 3)$ remains constant. By Proposition 1, the continuation value $\sum_{s\prime} P(s\prime \mid s, 0) V^*(s\prime)$ is also non-decreasing in t_1 because future states reachable under the Wait action have larger waiting times. Hence $Q(s, 0)$ is strictly increasing in t_1, while $Q(s, 3)$ is independent of t_1 . Thus $\Delta(t_1, t_2)$ is strictly decreasing in t_1 .

For each fixed t_2, there exists a minimal $t_1^*(t_2)$ such that $\Delta(t_1, t_2) \leq 0$ for all $t_1 \geq t_1^*(t_2)$. Define $\tau(t_2) = t_1^*(t_2)$. When t_2 increases, the cost of waiting $C(s, 0)$ increases (because W increases in t_2), which makes the Wait action less attractive. Therefore, the threshold t_1 required to trigger shipment decreases, and $\tau(t_2)$ is non-increasing in t_2 .