International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

The Value of Early Delivery Alerts in Construction Scheduling with Limited Site Storage

Sop Mbou Constant and **Ologun Michael Oreoluwa**
[1] Wuhan Textile University, Wuhan, China

Abstract - Advance notice of a late delivery is valuable to a contractor only if a useful decision remains open when the notice arrives. On congested sites, acting on an imperfect alert can also backfire: advancing one load to protect the programme may cause two loads to reach a full laydown area together. This study examines how alert timing, classification error and limited site storage jointly determine the value of early delivery alerts after materials have been ordered. A finite-scenario stochastic programme is formulated with binding booking cutoffs, a storage check at receipt, paid off-site deferral and an objective combining expected cost with the conditional value at risk (CVaR) of completion delay. A two-activity instance yields a closed-form break-even fee for acting on an imperfect alert. Exact enumeration of 36 synthetic designs shows that early rather than late noisy alerts strictly improve the cost–risk objective in 15 designs and tie in 21, with gains governed by alert accuracy and storage capacity. A second experiment optimises zone-specific booking policies and evaluates them on 20,000 simulated project paths per setting, including correlated disruptions. With five-day zone releases, only the final zone can affect project delay; with four-day releases this separability fails, and policies calibrated under independence can underperform the late-alert benchmark when disruptions are shared. A certified mixed-integer benchmark shows that simple dispatch rules match the joint optimum in three of four networks and exceed it by €53.47 in the fourth. The findings are conditional on synthetic instances and require field validation.

Keywords: **construction logistics; early delivery alerts; just-in-time delivery; laydown capacity; project scheduling; stochastic programming; conditional value at risk**

1 INTRODUCTION

Material flow is one of the most persistent sources of disruption in construction execution. Deficient site material management has long been associated with lost labour productivity and idle crews (Thomas et al., 2005), and construction supply chains remain exposed to supplier, transport and coordination risks that propagate to the site programme (Baghalzadeh Shishehgarkhaneh et al., 2024). Just-in-time (JIT) and lean approaches respond by synchronising deliveries with work readiness and keeping on-site stock low (Akintoye, 1995; Tommelein, 1998). On urban and congested sites, low stock is often not a choice but a physical necessity, because laydown areas are small and crane time is shared (Said & El-Rayes, 2011; Nolz, 2021).

Digital tracking and supply chain visibility technologies now make it increasingly feasible to warn the site team that a delivery is likely to be late (Ergen et al., 2007; Dallasega et al., 2018). Such alerts are rarely perfect, however, and their usefulness is bounded in time. An arriving truck can create a planning conflict even when it carries material the project needs. If a contractor advances one delivery in response to a threatened delay and the warning is wrong, both loads may arrive together at a site that cannot hold them. Holding a truck off site preserves feasibility but adds cost and may occupy a later booking slot. Conversely, a warning received after the booking deadline leaves the site team with schedule changes alone. The value of an alert therefore depends on three interacting factors: when it arrives relative to the booking cutoff, how accurate it is, and whether the site can physically absorb the consequences of acting on it.

Existing research has produced sophisticated models that integrate project scheduling with material ordering under uncertainty (Parchami Afra et al., 2022; Tian et al., 2024; Hamlehvar et al., 2026), as well as empirical accounts of construction logistics coordination (Ekeskär & Rudberg, 2016; Thunberg & Fredriksson, 2018). Research on early warnings in project management, in turn, has focused on detecting warning signs and on the organisational barriers to responding to them (Haji-Kazemi et al., 2013, 2015). What remains under-examined is the *execution-stage* question that links these streams: once quantities and suppliers are committed, how much is an imperfect, time-limited delivery alert worth when storage is constrained, and under what conditions does acting on it improve rather than worsen schedule performance and delay risk?

This study addresses that gap through three research questions:

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

- **1.** How do alert classification errors and the availability of legal booking options change the break-even price of advancing a shipment?
- **2.** How does receiving information before rather than after the booking cutoff change completion delay and cost once competition for storage space and crane time is represented?
- **3.** Does a booking rule calibrated under independent delays retain its value when several shipments are disrupted together or the production rhythm changes?

First, it formulates the post-commitment delivery decision as a finite-scenario stochastic programme in which booking cutoffs, receipt-time storage checks and paid off-site deferral are explicit, and in which information timing is represented through non-anticipativity constraints. Second, it derives a closed-form break-even fee for acting on an imperfect alert and identifies a regime in which a more accurate alert reduces expected delay without reducing tail delay. Third, exact computational experiments isolate the value of alert lead time by holding the disruption process fixed while changing only when, and how accurately, it is observed. Fourth, the study shows analytically and numerically that schedule slack determines *where* in a project alerts are valuable, and that dependence between disruptions can reverse nominal gains. Practically, the results translate into a decision procedure that procurement teams, logistics coordinators and site managers can apply to booking systems on storage-constrained sites.

Section 2 reviews the literature, Section 3 develops the theoretical framework, Section 4 presents the methodology and Section 5 the results. Sections 6–9 discuss the findings, implications, conclusions and limitations.

2 LITERATURE REVIEW

2.1 Integrated Project Scheduling and Material Ordering

Operations research has increasingly treated scheduling and material supply as a single decision problem. Parchami Afra et al. (2022) systematically review formulations and solution methods for the integrated project scheduling and material ordering problem and show that most models determine baseline order and start times in advance. Tian et al. (2024) combine uncertain material supply with limited storage in a chance-constrained baseline scheduling model, while Zhang et al. (2025) integrate procurement and resource-constrained scheduling for prefabricated buildings using an activity-delay-based ordering method. Hamlehvar et al. (2026) address procurement resilience in a two-stage stochastic model with backup options, and Kouvelis et al. (2023) analyse how inventories, time buffers and supplier flexibility protect project supply chains from material shortages. Lu et al. (2018) further show that the performance of material allocation policies depends on characteristics of the project network.

These models are grounded in the broader project-scheduling literature. The resource-constrained project scheduling problem and its extensions provide the precedence and renewable-resource structure on which they rest (Hartmann & Briskorn, 2010), and standard benchmark libraries supply test instances (Kolisch & Sprecher, 1997). Herroelen and Leus (2005) distinguish reactive, proactive and stochastic approaches to scheduling under uncertainty. Most integrated models are proactive in this sense: uncertainty is absorbed by buffers or recourse planned at the outset rather than by decisions responding to information revealed during execution. Nolz (2021) is a notable exception, coordinating construction and urban deliveries hierarchically and updating delivery decisions as availability information arrives. Even there, however, the timing of information relative to a binding booking deadline and the cost of acting on a false signal are not isolated as objects of analysis.

2.2 Just-in-Time Delivery, Site Storage and Construction Logistics

A second stream examines how materials physically reach and are held on site. Early studies explored the transfer of JIT principles to building material management (Akintoye, 1995) and to the supply of precast concrete components (Pheng & Chuan, 2001). Polat and Arditi (2005) cautioned that the economic advantage of JIT over buffered supply depends on local supply reliability and cost conditions, which is an early indication that low-inventory delivery is not unconditionally beneficial. Thomas et al. (2005) articulated principles of site material management that connect material availability to labour productivity.

Optimisation research has addressed the storage dimension directly. Said and El-Rayes (2011) model the trade-off between procurement and on-site storage under limited space, and later integrate procurement, storage and site layout in an automated multi-objective logistics optimisation system (Said & El-Rayes, 2014). Bortolini et al. (2019) use 4D building information modelling to plan and control site logistics for engineer-to-order prefabricated systems. Organisational studies show that coordinated delivery booking is increasingly common in practice: Ekeskär and Rudberg (2016) document a third-party logistics arrangement that centralised delivery coordination in a large hospital project, and Janné and Fredriksson (2019) develop governance guidelines for construction logistics centres in urban development, reporting consolidation and planning benefits.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

These studies show that storage and delivery coordination matter but treat delivery dates as planned inputs rather than as decisions revised in response to imperfect, time-limited information.

2.3 Early Warnings and Information in Project Control

Early warning research in project management argues that problems are often signalled before they materialise. Haji-Kazemi et al. (2013) review approaches for detecting early warning signs, and Haji-Kazemi et al. (2015) show that organisational barriers frequently prevent project teams from responding effectively once signs are identified. In construction specifically, tracking technologies such as radio frequency identification and GPS enable the status of components to be monitored through the supply chain (Ergen et al., 2007), and Industry 4.0 technologies are seen as enablers of closer coordination between off-site supply and on-site production (Dallasega et al., 2018). Lean construction offers a production-control rationale for such visibility: shielding crews from upstream uncertainty requires that work be released only when its prerequisites, including materials, are assured (Ballard & Howell, 1998), and the integration of lean and digital information flows is a recognised research agenda (Sacks et al., 2010).

This literature is largely qualitative or technology-oriented: it shows that warnings can be generated and that responding is organisationally difficult, but it does not price the response as a function of lead time, false-alarm rate or the site's capacity to absorb a changed delivery.

2.4 Synthesis and Research Gap

Table 1 summarises the three streams and their principal limitations relative to the present problem. Collectively, they suggest that information updates, feasible recovery actions and physical storage should be treated as parts of one decision process, yet no study combines all three at the execution stage.

Table 1 Summary of Prior Research Streams Relevant to Early Delivery Alerts

Research stream	Representative studies	Main contribution	Limitation for the present problem
Integrated scheduling, material ordering and scheduling under uncertainty	Parchami Afra et al. (2022); Tian et al. (2024); Zhang et al. (2025); Hamlehvar et al. (2026); Kouvelis et al. (2023); Herroelen and Leus (2005); Nolz (2021)	Joint optimisation of orders, buffers and activity starts; proactive and reactive scheduling	Decisions largely fixed in a baseline; booking cutoffs and the cost of acting on false signals not isolated
JIT, site material management and storage optimisation	Akintoye (1995); Pheng and Chuan (2001); Polat and Arditi (2005); Thomas et al. (2005); Said and El-Rayes (2011, 2014); Bortolini et al. (2019)	Conditions for low-inventory supply; storage, layout and logistics planning	Delivery dates treated as planned inputs; no treatment of imperfect advance information
Construction logistics coordination	Ekeskär and Rudberg (2016); Thunberg and Fredriksson (2018); Janné and Fredriksson (2019)	Booking systems, third-party logistics, planning practice	Value of information not quantified
Early warning and tracking	Haji-Kazemi et al. (2013, 2015); Ergen et al. (2007); Dallasega et al. (2018)	Warning detection, barriers to response, enabling technologies	Warnings not priced against lead time, accuracy and storage

The research gap is therefore specific and defensible: there is no quantitative account of the value of an imperfect delivery alert at the execution stage that jointly represents (a) the timing of the alert relative to a binding booking cutoff, (b) false alarms and missed detections, (c) receipt-time storage limits with paid off-site deferral, and (d) tail delay risk alongside expected cost. The present study addresses this gap and focuses on a deliberately narrow but practically common decision: whether to pay to advance an already booked load.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

3 THEORETICAL FOUNDATION AND CONCEPTUAL FRAMEWORK

3.1 Just-in-Time and Lean Construction

JIT theory holds that material should arrive when it is needed, in the quantity needed, so that inventory does not conceal process problems (Akintoye, 1995). Lean construction extends this logic to production control: pull-driven supply matches deliveries to demonstrated work readiness (Tommelein, 1998), and shielding protects production from upstream variability by releasing only sound assignments (Ballard & Howell, 1998). Two implications follow. Under JIT the site has little stock to absorb a late delivery, so timely information substitutes for physical buffers. Yet JIT also raises the cost of *over*-reaction: when the laydown area is sized for one load, an unnecessary advance creates congestion, double handling or off-site holding. The net value of an alert is therefore the delay it prevents less the waste its false alarms create.

3.2 Construction Supply Chain Coordination

Construction supply chain management emphasises the interface between the supply chain and the site, where independent suppliers, carriers and trades must be coordinated around a temporary production system (Vrijhoef & Koskela, 2000). Many on-site disruptions originate in insufficient planning across this interface (Thunberg & Fredriksson, 2018), and coordination mechanisms such as delivery booking systems and logistics centres make the interface explicit (Ekeskär & Rudberg, 2016; Janné & Fredriksson, 2019). From this perspective, an early delivery alert is a coordination signal passed across the supply chain–site interface. Its value depends on whether the receiving party retains decision rights, in this case an open booking slot, at the moment the signal arrives.

3.3 Value of Information and Risk-Averse Decision Theory

Decision theory provides the formal link between information timing and value. Blackwell (1953) showed that a more informative experiment cannot make an optimising decision-maker worse off. In a multistage stochastic programme, non-anticipativity constraints restrict each decision to information observable when it is made (Birge & Louveaux, 2011; Shapiro et al., 2021). An alert received before a cutoff refines the information partition at the booking node, which weakly enlarges the feasible policy set. Because contractors are typically concerned with severe overruns as well as average performance, the objective adds a coherent tail risk measure (Artzner et al., 1999), namely the CVaR of completion delay, which can be linearised in scenario form (Rockafellar & Uryasev, 2000, 2002).

3.4 Conceptual Framework

Figure 1 integrates these perspectives. The early delivery alert, characterised by its timing relative to the booking cutoff and by its accuracy, enters three execution-stage coordination decisions: the booking decision, gate admission with possible off-site deferral, and activity dispatch on a shared crane. These decisions jointly determine expected cost and tail delay. The relationship is moderated by limited laydown capacity, the availability of a fallback slot, schedule slack between work zones and the dependence of disruptions across loads. RQ1 and RQ2 concern the path from alert attributes through the coordination decisions to performance; RQ3 concerns the moderating role of disruption dependence and schedule rhythm.

Table 2 contrasts the conventional, plan-driven approach to delivery scheduling with the alert-based coordination examined in this study. The comparison is conceptual and defines the two policy regimes compared in Section 5.

Table 2 Conventional Delivery Scheduling Compared With Alert-Based Delivery Coordination

Dimension	Conventional delivery scheduling	Alert-based delivery coordination
Information basis	Planned delivery dates; delays observed at the gate	Advance transport alerts with known sensitivity and false-positive rate
Decision timing	Bookings fixed at ordering or by a standing rule	Booking revised if the alert arrives before the cutoff τ
Response to a threatened delay	Reactive resequencing of site work after the delay occurs	Proactive advance of a substitute load to keep the crane productive
Treatment of storage	Capacity checked informally; conflicts resolved ad hoc	Capacity checked at receipt; conflicts priced through off-site deferral
Exposure to false alarms	None, as no action is taken on forecasts	Explicit: unnecessary advances incur fees, holding cost and congestion

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Dimension	Conventional delivery scheduling	Alert-based delivery coordination
Performance criterion	Typically on-time completion or mean cost	Expected cost plus weighted tail delay, $E[C] + \lambda \cdot CVaR$
Data requirements	Delivery schedule and site plan	Additionally, dated alerts, cutoff times, gate records and alert accuracy

Note. Authors' synthesis based on the decision model in Section 4.

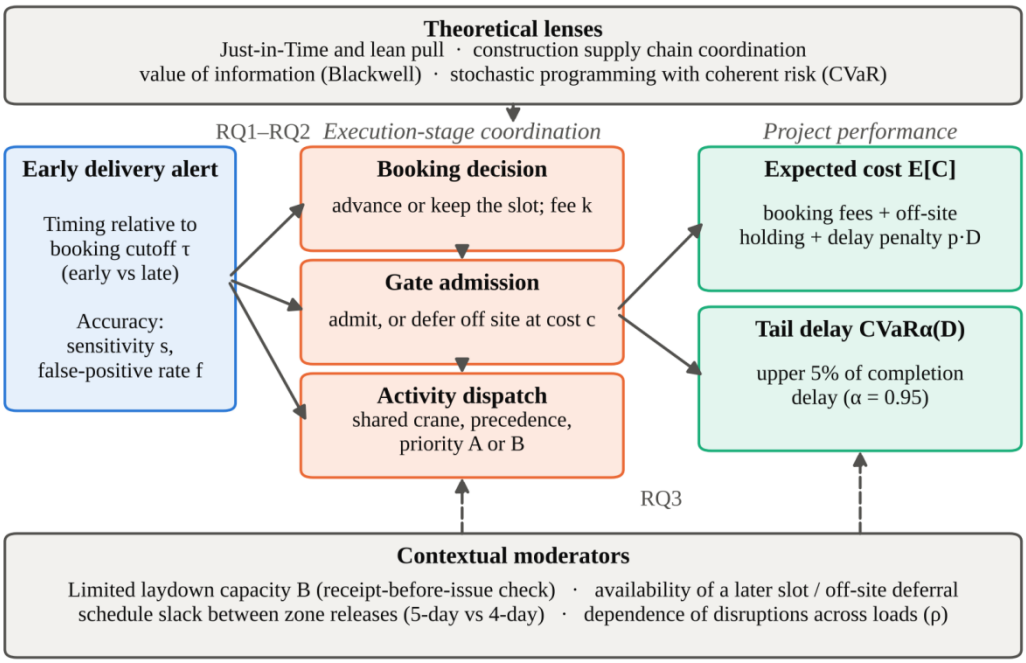

Figure 1 Conceptual Framework Linking Early Delivery Alerts, Execution-Stage Coordination, Limited Site Storage and Project Performance

4 RESEARCH METHODOLOGY

4.1 Research Design

The study adopts a quantitative modelling design that combines analytical derivation with exact computational experiments on synthetic instances. This design suits the research questions because the value of information is a counterfactual quantity: it requires comparing outcomes under the same disruption realisations with and without earlier information, which field observation cannot provide directly. The design proceeds in four steps. First, a general decision model is formulated (Sections 4.2–4.3). Second, a two-activity instance is solved in closed form to answer RQ1 (Section 4.4). Third, two computational experiments answer RQ2 and RQ3 (Sections 4.5–4.6). Fourth, a certified mixed-integer benchmark quantifies the loss caused by restricting attention to tractable dispatch rules (Section 4.7). Figure 2 summarises the operational logic represented in all experiments.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Figure 2 Process Flow Showing How an Early Delivery Alert Influences Booking, Gate Admission and Activity Dispatch

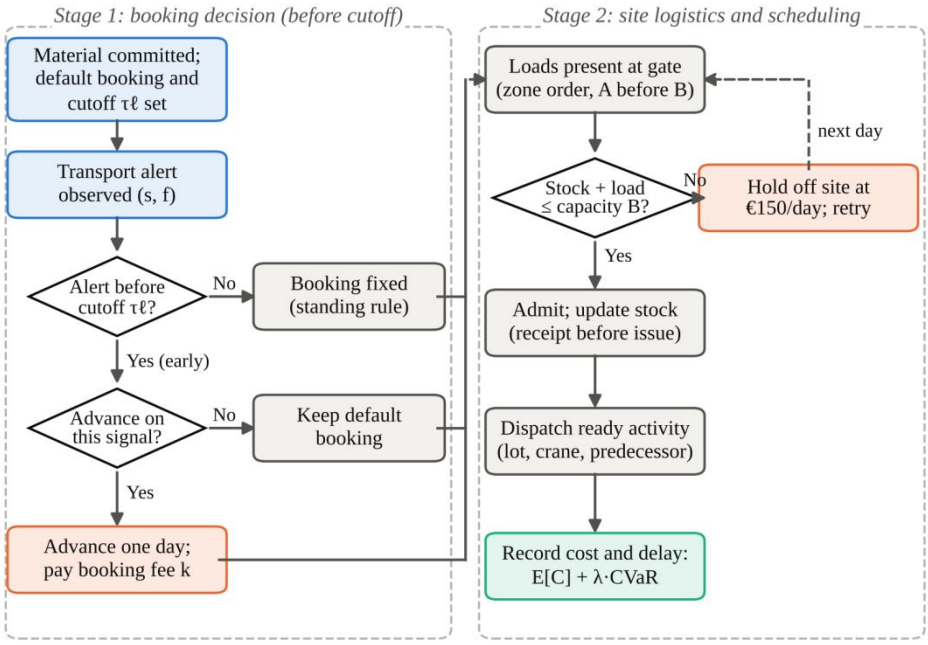

4.2 Decision Model

Events and information. A finite set of scenarios Ω has probabilities πω. Each scenario specifies the dates on which committed loads are ready to present at the site and the timestamps and contents of alerts. Project activities i have durations d_i, finish-to-start precedences and renewable resource requirements r_i^k. Load ℓ contains quantity $q\ell_m$ of material m. Each flexible load has a default booking, a set of admissible replacement slots and a cutoff τℓ; an alert observed at or before τℓ can affect that booking. Within a day, the site observes which loads are ready to enter, admits or defers them, checks storage, issues materials and starts eligible activities. An activity started on day t occupies its resources through day $t + d_i - 1$.

Booking and material balance. Let $b\ell u\omega$ select booking u for load ℓ, fixed after τℓ; let $G\ell tu\omega$ equal one if booking u makes load ℓ available at the gate on day t in scenario ω; let $a\ell t\omega$ denote admission and $Q\ell t\omega$ a presented load held off site at the end of day t. With initially empty queues,

$$\sum_u b\ell u\omega = 1,\ Q\ell t\omega = Q\ell,t-1,\omega + \sum_u G\ell tu\omega\, b\ell u\omega - a\ell t\omega,\ \sum_t a\ell t\omega = 1. \quad (1)$$

Booking and admission variables are binary and queues are nonnegative. Let $S_{mt}\omega$ be on-site stock immediately after admission, $I_{mt}\omega$ the stock after issue, a_{im} the material requirement of activity i, v_m the volume per unit and B the site capacity:

$$S_{mt}\omega = I_{m},t-1,\omega + \sum_\ell q\ell_m\, a\ell t\omega,\ \sum_m v_m\, S_{mt}\omega \le B,\ I_{mt}\omega = S_{mt}\omega - \sum_i a_{im}\, x_{it}\omega \ge 0. \quad (2)$$

The capacity constraint applies at receipt, before any issue: simultaneous admission of two loads is infeasible if it would exceed B, even if a task would consume one load later that day. The experiments assume full-lot issue and no direct unloading into installation.

Schedule and loss. For binary starts $x_{it}\omega$, $s_i\omega = \sum_t t\, x_{it}\omega$ with $\sum_t x_{it}\omega = 1$. Each precedence $i \to j$ imposes $s_j\omega \ge s_i\omega + d_i$, and for each resource k and day t, $\sum_i r_i^k \sum_a x_{ia}\omega \le R_{kt}$, where the inner sum runs over start days a with $a \le t < a + d_i$. With contractual last day H, completion delay is $D\omega = \max\{0, \max_i(s_i\omega + d_i - 1) - H\}$.

Objective. Scenario cost Cω comprises booking fees $\sum k\ell_u\, b\ell u\omega$, off-site holding charges $\sum c\ell_t\, Q\ell t\omega$ and the delay penalty $pD\omega$, where p is the contractual penalty per day. The objective, $E[C\omega] + \lambda\, CVaR\alpha(D\omega)$, is

$$\min E[C\omega] + \lambda\,[\eta + \sum_\omega \pi\omega\, \xi\omega\,/(1-\alpha)],\ \xi\omega \ge \max(0, D\omega - \eta), \quad (3)$$

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

where the bracket is the scenario representation of λ CVaRα(Dω) (Rockafellar & Uryasev, 2000), with η unrestricted, α = 0.95 and λ in euros per day of tail delay. With finite time and scenario sets and linear charges, the model is a mixed-integer linear programme.

4.3 Information Constraints and Feasibility Audit

Following the deterministic-equivalent approach (Birge & Louveaux, 2011; Shapiro et al., 2021), booking variables are identified across scenarios whose histories coincide at the cutoff: $b_{\ell u\omega} = b_{\ell u\omega'}$ whenever $H_\omega(\tau_\ell) = H_{\omega'}(\tau_\ell)$. Admission decisions share histories through the gate observation, and activity starts share histories through material receipt. Moving an alert across a cutoff changes only the information partition at the corresponding booking node; it does not change disruption outcomes, available bookings or physical gate capacities. This property is what allows the experiments to isolate the value of lead time.

Each load is presented once and admitted once. With $F_{\ell t}$ and $A_{\ell t}$ denoting cumulative presentations and admissions, $Q_{\ell t} = F_{\ell t} - A_{\ell t}$, $0 \le A_{\ell t} \le F_{\ell t} \le 1$ and $A_{\ell T} = F_{\ell T} = 1$; a repeat gate attempt by a held truck leaves $F_{\ell t}$ unchanged. A booking rule therefore cannot depend on the later event that confirms whether the alert was correct. The study uses one ex ante CVaR criterion throughout and makes no claim of recursive time consistency.

4.4 Analytical Instance

Two two-day activities share one crane. Activity A receives a 6 m³ load on day 3 (on time) or day 5 (delayed). The 6 m³ load for activity B is booked for day 4 and can be advanced to day 3 for fee k before the end of day 2. Storage capacity is 8 m³. If both loads present on day 3, A has priority and B can be held off site until day 4 at cost c. The contractual completion day is 6. Let q denote the probability that load A is delayed, s the alert sensitivity and f the false-positive probability.

4.5 Value of Early Alerts Under a Shared Booking Rule

The computational instances contain three or six work zones, each with two activities (A and B) and one dedicated 6 m³ load per activity. All activities use the same crane and take two days. Within each work stream an activity must follow the corresponding activity in the preceding zone; the streams can otherwise be interleaved. In zone g, numbered from zero, the normal arrival of A is day 1 + 5g and its disrupted arrival is two days later. B is normally booked one day after the normal A arrival and can be advanced by one day. The contractual last day is 5N − 1 for N zones. Loads presenting at the gate, including deferred loads, are processed in zone order and A before B; a load that would exceed capacity returns the next day at €150 per day held. Only after these admissions does the simulator start a ready activity from the earliest zone with at least one ready activity, where ready means the complete lot is on site, the crane is free and the predecessor, if any, is complete; the crane processes one activity at a time. Two dispatch variants break same-zone ties in favour of A or B. Table 3 lists the experimental factors.

Table 3 Experimental Factors and Fixed Parameters in Experiment 1

Element	Levels or value
Work zones N	3 or 6 (6 or 12 crane-dependent activities)
Site storage capacity B	8, 12 or 24 m³
Disruption probability per zone q	0.10 or 0.25 (independent across zones)
Alert quality (s, f)	Reference (0.80, 0.10); low (0.60, 0.20); high (0.95, 0.02)
Information regimes	Late noisy (after cutoff); early noisy (before cutoff); early perfect
Booking rules	Never, advance on positive alert, advance on negative alert, always; applied to all zones
Fixed costs	k = €100; p = €400/day; c = €150/day; λ = €100/day; α = 0.95

Early noisy alerts are observed before the booking cutoff; late noisy alerts arrive after the booking decision and are not used by the common dispatch rules, so the late regime can use only the two unconditional rules; early perfect alerts report the disruption state before the cutoff and serve as a benchmark. Each booking rule is crossed with the two dispatch priorities, and the best policy is selected per design and regime. All joint paths of disruption and alert are enumerated, up to $4^6 = 4{,}096$ paths, so the reported values are exact probability-weighted outcomes without Monte Carlo error.

4.6 Experiment 2: Zone-Specific Optimisation and Out-of-Sample Simulation

The second experiment allows the response to vary by zone, so that a critical late-stage delivery can be treated differently from an earlier delivery with slack. A zone rule maps a binary alert to an advance decision, giving four deterministic mappings per

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

zone. Six zones yield $4^6 = 4{,}096$ early-alert policies and $2^6 = 64$ unconditional late-alert policies, each crossed with the two dispatch priorities; the shared-rule benchmark evaluates the four identical-rule vectors. For each capacity, delay and cost are first tabulated by dispatch priority, disruption vector and booking vector, which requires $2 \times 2^6 \times 2^6 = 8{,}192$ physical schedules. Every candidate policy is then scored exactly on the joint distribution, and the optimum within each class follows from complete enumeration. CVaR is computed from the discrete delay histogram by accumulating probability mass from the largest delay downward until it reaches 0.05, with fractional weight at the boundary.

The calibrated policies are then fixed and evaluated on 20,000 independent project paths per test setting, paired across competing policies. Each zone has disruption probability 0.10. With probability ρ, the project receives one common Bernoulli disruption indicator shared by all zones; otherwise it receives independent indicators by zone, so the pairwise correlation between zones is ρ. Tests use $\rho \in \{0, 0.35, 0.65\}$. Nine combinations of storage capacity, alert quality and release spacing, each crossed with the three correlation levels, give 27 test settings. The reference schedule releases a zone every five days (deadline day 29 for six zones); a four-day schedule removes the nominal day of slack between zones (deadline day 24). The score equals mean cost plus €100 per day of empirical CVaR. Approximate 95% percentile intervals for paired score differences use 300 project-level bootstrap replications (path seed 20261119; bootstrap seed 20261120). These intervals measure simulation uncertainty only, not uncertainty in the assumed probabilities or site parameters. For four-day schedules, a dependence-aware rule is also calibrated by minimising the worst excess objective over the same fixed late policy across the three correlation levels.

4.7 Exact Joint Benchmark

The policy experiments fix the gate and dispatch logic. To quantify the loss caused by this restriction, the deterministic equivalent of the joint booking, admission and start-time problem is solved on four small synthetic networks (Table 4). Each has four activities and four dedicated 6 m³ loads; each of two zones has one disrupted and one flexible load, giving 16 joint scenarios with positive probability ($s = 0.80$, $f = 0.10$, $q = 0.10$). Both alerts are available at the common booking decision, and later decisions can use gate presentations observed to date. Off-site deferral costs €150 per load per day, the booking fee is €100, delay costs €400/day and the CVaR weight is €100/day within a 16-day horizon.

Table 4 Input Data for the Four Exact Benchmark Networks

Network	Edges	Durations (days)	Crane demands	Crane / storage	Due day / zone 2 release
Parallel streams	0→2, 1→3	2, 2, 2, 2	1, 1, 1, 1	1 / 8 m³	8 / 4
Fork and join	0→1, 0→2, 1→3, 2→3	1, 2, 3, 1	1, 1, 1, 1	2 / 12 m³	6 / 4
Cross precedence	1→2, 0→3, 2→3	2, 1, 3, 1	1, 1, 1, 1	1 / 12 m³	7 / 4
Unequal crane demand	0→2, 1→3	2, 3, 2, 2	2, 1, 1, 2	2 / 12 m³	8 / 5

Note. Activity indices are zero-based. The disrupted lot normally presents on day 1 in zone 1 and on the stated release day in zone 2; disruption adds two days. The flexible lot normally presents one day after the zone release and can advance one day.

With dedicated loads, cumulative admission must precede cumulative starts, and the receipt peak counts admissions up to and including day t while removing only lots consumed before day t:

$$\sum(u \le t)\, x_{ju} \le \sum(u \le t)\, a_{ju},\ 6\ \sum_j [\sum(u \le t)\, a_{ju} - \sum(u < t)\, x_{ju}] \le B. \quad (4)$$

Booking variables are shared across scenarios with the same alert pair; gate and start variables are shared across scenarios with the same alert pair and the same dated presentations of the disrupted loads. For comparison, the dispatch policy class allows each zone one of the four local booking rules, crossed with two crane priorities, giving 32 candidates per network. The problems are solved with SciPy 1.17.0 and the HiGHS mixed-integer solver at a requested relative gap of zero.

4.8 Verification and Reproducibility

Each simulated trajectory verifies the capacity invariant at receipt, single starts and termination of all activities, and an empty stock at completion. The two-activity instance is checked against all four analytical policy objectives, and scenario probabilities are checked to sum to one. The separability result in Section 5.3 is checked exhaustively against the multi-zone simulator. In the exact benchmark, every scenario path is reconstructed from the solver output and receipt times, stock, precedence and crane capacity are rechecked independently of the optimisation constraints.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

5 RESULTS AND ANALYSIS

5.1 Analytical Results: Break-Even Fee and the Role of False Alarms

Table 5 shows earliest completion in the analytical instance. Advancing B protects the deadline when A is late but requires off-site holding of B when A is on time.

Table 5 Earliest Completion Day in the Analytical Instance

Actual arrival of load A	Keep B on day 4	Advance B to day 3
Day 3 (on time)	Day 6	Day 6, with B held off site one day
Day 5 (delayed)	Day 7	Day 6

Under the rule that advances B on a positive alert, the probability of paying k is $a = qs + (1 - q)f$, the probability of paying c is $(1 - q)f$, and the probability of one day of delay is $q(1 - s)$; under a fixed day-4 booking the delay probability is q. For a one-day loss with probability u, the upper-tail CVaR is $L(u) = \min\{1, u/(1 - \alpha)\}$ days.

Proposition 1. If $a > 0$ and deferral to day 4 is guaranteed, the positive-alert rule has a strictly smaller objective than the fixed day-4 booking if and only if

$$k < [pqs + \lambda\{L(q) - L(q(1 - s))\} - c(1 - q)f] / a. \quad (5)$$

Proof. The positive-alert rule has expected cost $ak + c(1 - q)f + pq(1 - s)$ and risk term $\lambda L(q(1 - s))$; the fixed booking has objective $pq + \lambda L(q)$. Subtraction and division by a give the threshold, with a tie at equality.

To compare all deterministic rules, let y_0 and y_1 be the advance decisions after a negative and positive alert. The probabilities of advancing a truly late load (u) and of advancing an on-time load (v), and the objective, are

$$u = q[(1 - s)y_0 + sy_1], \; v = (1 - q)[(1 - f)y_0 + fy_1], \quad (6a)$$

$$J(y_0, y_1) = (u + v)k + vc + p(q - u) + \lambda L(q - u). \quad (6b)$$

A fee is paid on every advance; on-time branches with an advance need the fallback at cost c; and a one-day loss occurs only when the load is late and was not advanced. Table 6 lists the resulting event probabilities.

Table 6 Event Probabilities for the Four Deterministic Booking Rules

Rule	Decisions (y_0, y_1)	u	v	Probability of delay
Never	(0, 0)	0	0	q
Positive alert	(0, 1)	qs	$(1 - q)f$	$q(1 - s)$
Negative alert	(1, 0)	$q(1 - s)$	$(1 - q)(1 - f)$	qs
Always	(1, 1)	q	$1 - q$	0

With $q = 0.10$, $s = 0.80$, $f = 0.10$, $p = €400$/day, $c = €150$ and $\lambda = €100$/day, the four objectives as functions of k are 140, $0.17k + 61.50$, $0.83k + 253.50$ and $k + 135$ euros, in the order of Table 6. Figure 3 plots this lower envelope. Advancing on a positive alert is uniquely optimal below $k = €461.76$, keeping the booking is uniquely optimal above it, and the two tie at the boundary. At $k = €100$, the objectives for never advancing, following a positive alert, following a negative alert and always advancing are €140.00, €78.50, €336.50 and €235.00. The €78.50 decomposes into booking fees of €17.00 (action probability 0.17), fallback cost of €13.50 (deferral probability 0.09), expected delay damages of €8.00 (delay probability 0.02) and weighted tail risk of €40.00 (tail delay 0.02/0.05 = 0.40 days). This global ordering is specific to the parameterisation.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Figure 3 Objective of the Four Booking Rules as a Function of the Advance Booking Fee in the Analytical Instance

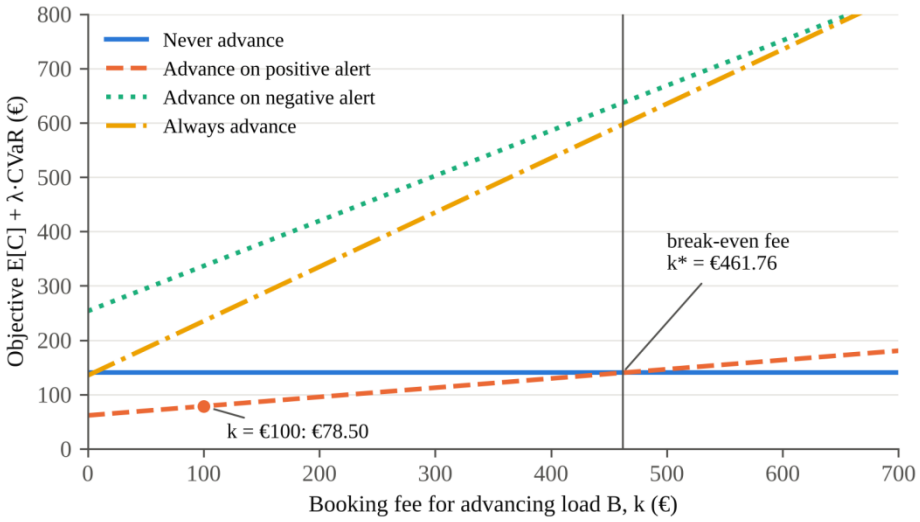

Three further analytical results clarify the mechanisms. First, *a better alert does not always reduce tail risk.* Let $\varepsilon = 1 - \alpha$. If $q(1 - s) \geq \varepsilon$, the risk improvement $L(q) - L(q(1 - s))$ is zero even though the alert reduces mean delay; if $q \geq \varepsilon > q(1 - s)$, the improvement is $1 - q(1 - s)/\varepsilon$ days; and if $q < \varepsilon$, it is qs/ε days. When baseline delays already fill the tail, sensitivity must exceed $1 - \varepsilon/q$ before CVaR begins to fall; at $q = 0.10$ and $\alpha = 0.95$ this requires sensitivity above 0.50. Second, *false alarms are priced through the fallback.* An increment δf in the false-positive probability raises the expected fee and deferral cost of the positive-alert rule by $(1 - q)(k + c)\delta f$ without changing its delay distribution, provided the day-4 fallback preserves on-time completion. If deferral instead causes a work stoppage, a false alert contributes to both cost and tail risk. Third, *storage determines feasibility, not merely cost.* Under the receipt-before-issue convention, simultaneous admission of volumes v_A and v_B is feasible precisely when $B \geq v_A + v_B$. If $\max(v_A, v_B) \leq B < v_A + v_B$ and neither load can be deferred, any rule that advances on an on-time branch with positive probability is infeasible; with $f > 0$ this disqualifies advance-on-positive entirely.

Finally, the value of earlier information is weakly positive by construction. If the early history contains all information available to the late regime and ignoring information is permitted, every feasible late policy can be copied into the early regime with unchanged pathwise cost and delay, so the minimum early objective is weakly smaller. Strict improvement requires a policy that uses the earlier distinction to change a feasible decision profitably. This inclusion argument, a pathwise analogue of Blackwell (1953), does not order objectives for policies fitted under one distribution and tested under another.

5.2 Value of Early Alerts Under a Shared Rule

Across the 36 designs, early noisy alerts reduced the objective relative to late noisy alerts in 15 cases and tied in 21, with reductions ranging from €0.00 to €94.25 per project. Table 7 reports the reference alert quality at disruption probability 0.10.

Table 7 Expected Cost Plus Weighted Tail Delay (€) for the Best Policy in Each Information Regime

Zones	Storage (m³)	Late noisy	Early noisy	Early perfect	Early noisy rule
3	8	140.00	139.50	30.00	positive
3	12	140.00	99.00	30.00	positive
3	24	140.00	99.00	30.00	positive
6	8	140.00	140.00	60.00	never
6	12	140.00	140.00	60.00	never
6	24	140.00	140.00	60.00	never

Because the early-noisy policy family contains both unconditional late-regime policies, a tie indicates that the best alert-based rule does no better than the best unconditional rule. Alert quality changes the result materially: the low-quality setting ($s = 0.60$, $f = 0.20$) improved the objective in none of its 12 designs, the reference setting in three and the high-quality setting ($s = 0.95$, $f = 0.02$) in all 12. Averaged over the equally weighted design grid, the early-alert gain was €18.58 at 8 m³ and €25.67 at either 12 or 24 m³; the equality at the two larger capacities reflects the load sizes and arrival pattern of the experiment. In this

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

experiment, storage therefore acted as a threshold rather than a continuous moderator: once two loads could be held together, additional space added no further value. The largest gain occurs with three zones, 12 m³ of storage, disruption probability 0.25, sensitivity 0.95 and false-positive probability 0.02, where the early noisy regime selects the positive-alert rule (Table 8).

Table 8 ***Outcomes for the Design With the Largest Early Noisy Alert Gain***

Regime	Mean delay (days)	CVaR (days)	Mean cost (€)	Objective (€)
Late noisy	0.25	1.00	100.00	200.00
Early noisy	0.01	0.25	80.75	105.75
Early perfect	0.00	0.00	75.00	75.00

Here the early noisy alert recovers most of the gap to the perfect-information benchmark, mainly by reducing tail delay from 1.00 to 0.25 days. The single-zone break-even threshold of Section 5.1 is not applied separately to each zone, because holds can propagate between zones and completion depends on the interaction of gate admission, precedence and crane availability.

5.3 Experiment 2: Zone-Specific Policies, Schedule Slack and Disruption Dependence

Table 9 reports out-of-sample objectives for five-day releases under reference alert quality. Across all 27 test settings, 19 paired intervals for the zone-specific policy's gain over the optimised late regime lie wholly above zero and five lie wholly below zero; point estimates range from −€63.37 to €114.31.

Table 9 Out-of-Sample Objective (€) and Paired 95% Interval for the Zone-Specific Policy Gain

Storage (m³)	Correlation ρ	Late	Shared rule	By zone	Gain interval
8	0.00	139.34	139.34	77.48	57.60 to 65.91
8	0.35	138.74	138.74	74.40	60.16 to 68.34
8	0.65	138.76	138.76	75.81	58.19 to 67.05
12	0.00	100.00	139.34	63.73	31.47 to 40.85
12	0.35	100.00	138.74	60.75	34.63 to 43.68
12	0.65	100.00	138.76	62.08	32.61 to 42.63
24	0.00	100.00	139.34	63.73	31.47 to 40.85
24	0.35	100.00	138.74	60.75	34.63 to 43.68
24	0.65	100.00	138.76	62.08	32.61 to 42.63

Note. Five-day zone releases, reference alert quality. A positive gain denotes a lower objective than the late comparator. Intervals are from 300 paired bootstrap replications of 20,000 paths.

Table 10 shows cost and tail delay at the two extreme simulated gains. Rule vectors are ordered from the first to the sixth zone, where 0 = never advance, 1 = advance on a positive alert, 2 = advance on a negative alert and 3 = always advance.

Table 10 Cost and Tail Delay at the Two Extreme Simulated Gains

Release, capacity, sensitivity, ρ	Policy	Rule vector	Mean cost (€)	CVaR (days)
4 d, 8 m³, 0.80, 0.65	Late optimised	000000	153.01	1.000
4 d, 8 m³, 0.80, 0.65	Zone-specific alert rules	111110	216.38	1.000
5 d, 8 m³, 0.95, 0.65	Late optimised	000000	38.76	1.000
5 d, 8 m³, 0.95, 0.65	Zone-specific alert rules	000001	15.74	0.087

Figure 4 shows how schedule slack determines where alerts are useful. At 8 m³ and reference quality, exact calibration selects rule vector 000001 for five-day releases and 111110 for four-day releases. With one day of planned separation, local lateness in early zones is absorbed before the next release, and only the last zone lacks a subsequent buffer. With a four-day rhythm, the selected policy instead uses alerts in the earlier zones and retains the default booking in the last.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Figure 4

Paired Out-of-Sample Gain of Zone-Specific Alert Rules Over the Late-Alert Policy Under Reference Alert Quality

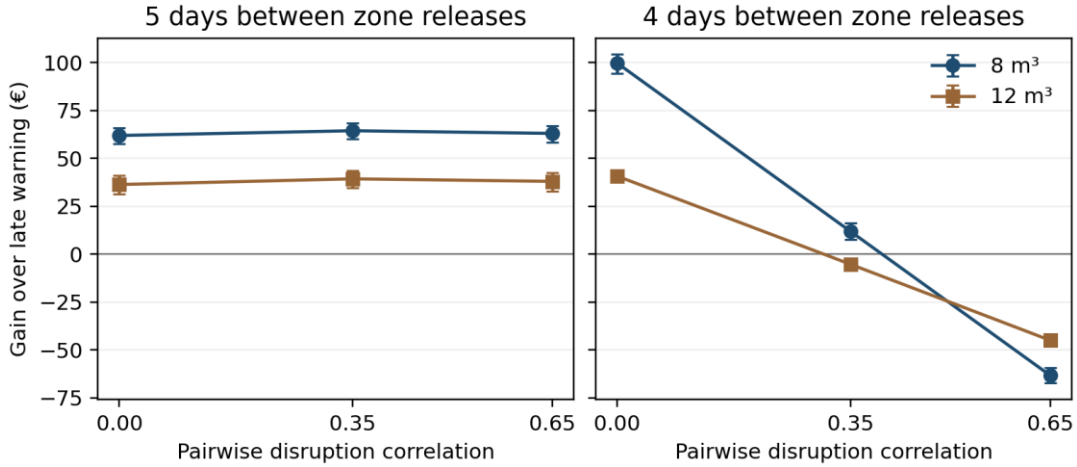

Proposition 2. For the two-stream benchmark with five-day zone releases, one shared crane, two two-day activities per zone, admission priority A before B, next-day deferral and capacity of at least one lot, the six-zone trajectory separates into one-zone trajectories, and project completion delay equals the delay of the last zone, provided no other work competes for the crane or stock and the earliest-ready dispatch rule is used.

Proof. Let r be the first possible delivery day in a zone. If A arrives at r, B is available by r + 1, including the one-day fallback if needed, and the zone finishes by r + 3. If A arrives at r + 2 and B advances to r, B runs on r and r + 1 and A on r + 2 and r + 3. If A arrives at r + 2 and B keeps r + 1, B runs on r + 1 and r + 2 and A on r + 3 and r + 4. In every branch all lots are consumed and the crane is free before the next zone's first decision on r + 5. The first zone starts with no stock or active tasks, and induction yields the claim.

$$D = D^N,\ C = \sum_g (k\, a^g + c\, z^g) + p D^N. \quad (7)$$

Here a^g and z^g indicate an advance and an off-site deferral in zone g. Equation 7 implies that advancing in any zone before N cannot reduce project delay, so with $k > 0$ and $c \geq 0$ never advancing those zones is weakly optimal and any rule causing an advance there with positive probability is strictly more expensive for the same last-zone decision. This explains the 000001 solution without requiring independent disruptions. An exhaustive check confirmed equal delay and cost between Equation 7 and the simulator on all 24,576 combinations of 64 disruption vectors, 64 booking vectors, three capacities and two crane priorities. When releases occur every four days, the induction fails and delays propagate through the shared crane.

Table 11 examines the consequence for robustness. Nominal policies calibrated under independence are attractive at $\rho = 0$ but lose their advantage, and at 8 m³ incur a loss of €63.37, when disruptions are shared. In all three four-day storage settings, the dependence-aware criterion selects 000000, the same decision as the late benchmark, so its simulated gain and interval are exactly zero.

Table 11

Four-Day Schedules Under Nominal and Dependence-Aware Calibration

Storage (m³)	ρ	Nominal rule	Robust rule	Nominal gain (€)	Robust gain (€)	Robust interval (€)
8	0.00	111110	000000	99.58	0.00	0.00 to 0.00
8	0.35	111110	000000	11.83	0.00	0.00 to 0.00
8	0.65	111110	000000	−63.37	0.00	0.00 to 0.00
12	0.00	111111	000000	40.63	0.00	0.00 to 0.00
12	0.35	111111	000000	−5.34	0.00	0.00 to 0.00
12	0.65	111111	000000	−45.04	0.00	0.00 to 0.00
24	0.00	111111	000000	40.63	0.00	0.00 to 0.00
24	0.35	111111	000000	−5.34	0.00	0.00 to 0.00
24	0.65	111111	000000	−45.04	0.00	0.00 to 0.00

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Because the robust rule was calibrated on three correlation levels only, the continuous interval between them was also verified for the selected policies.

Proposition 3. Let fixed policies induce nonnegative integer delays bounded by M days and let ε be the upper-tail mass. Under the disruption mixture of Section 4.6, $T^j(\rho) = \Pr\rho(D \geq j)$ and $\bar{C}(\rho) = E\rho[C]$ are affine in ρ, and

$$J(\rho) = \bar{C}(\rho) + \lambda \sum_{j=1}^{M} \min\{1, T^j(\rho)/\varepsilon\}. \quad (8)$$

Proof. Sort delay outcomes from largest to smallest and collect mass ε, allowing fractional mass at the boundary (Rockafellar & Uryasev, 2002). Each integer delay d is the sum of d indicators of exceeding successive integer levels, and at level j the collected mass is $\min\{\varepsilon, T^j\}$. Summing over levels gives Equation 8.

J(ρ) is therefore continuous and piecewise affine, with slope changes only where $T^j(\rho) = \varepsilon$. On each segment between the crossing points of two policies, their difference is affine, so its maximum occurs at an endpoint. Table 12 reports the resulting worst excess over $0 \leq \rho \leq 0.65$.

Table 12

Maximum Modelled Excess Relative to the Fixed Late Policy Over $0 \leq \rho \leq 0.65$

Storage (m³)	Early policy	Candidate points	Worst ρ	Excess (€)
8	Zone-specific alert rules	2	0.6500	64.70
8	Robust relative to late	2	0.0000	0.00
12	Zone-specific alert rules	2	0.6500	45.57
12	Robust relative to late	2	0.0000	0.00
24	Zone-specific alert rules	2	0.6500	45.57
24	Robust relative to late	2	0.0000	0.00

5.4 Exact Joint Benchmark

Table 13 compares the best dispatch-rule policy with the certified joint optimum. For each instance, the recovered feasible objective and the solver lower bound agree within €0.00001, and all 32 dispatch policies finish within the horizon.

Table 13 Best Dispatch-Rule Policy Compared With the Certified Joint Optimum (€)

Network	Joint objective	Solver lower bound	Best rule	Rule gap
Parallel streams	219.20	219.20	219.20	0.00
Fork and join	780.00	780.00	780.00	0.00
Cross precedence	316.00	316.00	316.00	0.00
Unequal crane demand	187.53	187.53	241.00	53.47

The best dispatch policy ties the joint solution in the first three networks. With unequal crane demands, its objective of €241.00 exceeds the joint optimum of €187.53 by €53.47 (28.51%): expected cost is higher by €20.67, and CVaR is higher by 0.328 days, accounting for the remaining €32.80. Because the joint model relaxes the booking, admission and dispatch restrictions simultaneously, the gap cannot be attributed to any single decision. These are certified optima for four finite synthetic instances; the agreement in three networks supports the rule family there, and the unequal-demand case is a counterexample to its general optimality.

6 DISCUSSION

Proposition 1 and Figure 3 show that the price a contractor can justify for advancing a load rises with the delay penalty and the tail-risk reduction, and falls with the expected cost of false alarms, $(1 - q)(k + c)f$. This formalises the JIT intuition that over-reaction is costly on lean sites (Akintoye, 1995; Polat & Arditi, 2005) and gives it an operational form: the false-alarm cost is borne through the fallback slot. More fundamentally, the storage analysis shows that a fallback changes the *feasible* policy set, not only its cost. Without a deferral option, a site whose laydown area cannot hold two loads cannot act on an imperfect alert at all. This extends storage-aware procurement models (Said & El-Rayes, 2011, 2014), which treat space as a cost or capacity parameter, by showing that space also determines whether information is actionable. The CVaR result adds a less intuitive insight: when baseline delays already fill the tail, a more accurate alert can reduce average penalties while leaving exposure to the worst

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

outcomes unchanged. Contractors whose primary concern is catastrophic overrun should therefore evaluate alert systems against tail metrics, not only against mean lateness.

Experiment 1 confirms that earlier information is weakly valuable by construction but strictly valuable only under specific conditions: in 15 of 36 designs, and in none of the designs with low alert quality. This is consistent with value-of-information theory (Blackwell, 1953) but qualifies the presumption that greater supply chain visibility is inherently beneficial (Dallasega et al., 2018): visibility creates value only when a decision right remains open, the signal is accurate enough to overcome false-alarm costs, and the site can physically absorb the response. The threshold effect of storage, with identical average gains at 12 and 24 m³, suggests that the critical design question for site logistics is whether two consecutive loads can coexist, rather than how much total space is available.

Experiment 2 yields the study's most managerially significant finding. Proposition 2 shows that, with one day of slack between zone releases, delay cannot propagate and only the final zone's booking matters. Removing that slack reverses the location of valuable alerts, and policies calibrated under independence can then underperform the late benchmark by up to €63.37 when disruptions are shared. This connects to the observation that material allocation policies depend on network characteristics (Lu et al., 2018) and to the buffer-versus-flexibility trade-off identified by Kouvelis et al. (2023), but adds a specific mechanism: release slack acts as a firewall that decouples zones, and its removal transmits delay through the shared crane. It also complements the risk literature (Baghalzadeh Shishehgarkhaneh et al., 2024): correlated supply failures can turn a locally rational response into a project-level loss. The dependence-aware rule avoids that loss only by forgoing alerts altogether, which suggests that robustness and information value may be in tension when the correlation structure is uncertain.

Finally, the exact benchmark qualifies the policy experiments. Simple, implementable rules were optimal in three of four networks, which supports their practical use, but the 28.51% gap under unequal crane demands shows that heterogeneous resource requirements can make joint optimisation of bookings, admissions and starts worthwhile. This echoes calls to integrate supply chain planning with site production planning rather than treating them sequentially (Thunberg & Fredriksson, 2018; Vrijhoef & Koskela, 2000).

7 PRACTICAL AND MANAGERIAL IMPLICATIONS

The findings yield five recommendations for construction managers, site engineers, procurement teams, logistics coordinators and contractors on storage-constrained sites.

Record alert timing against booking cutoffs. Procurement and logistics teams should log when each alert arrives relative to the latest time the slot can be changed. An alert received after the cutoff has no booking value however accurate it is, so earlier transmission may be worth more than greater accuracy.

Price the response before acting. Before paying to advance a load, the booking fee should be compared with the expected penalty and tail delay avoided net of expected false-alarm costs, as in Equation 5. Logistics coordinators can maintain simple break-even thresholds per supplier based on the observed sensitivity and false-alarm rates of each supplier's alerts.

Secure a fallback before relying on alerts. Site engineers should confirm that a deferral option, such as an off-site holding area, a later slot or a consolidation centre (Janné & Fredriksson, 2019), exists whenever the laydown area cannot hold two consecutive loads. Without it, acting on imperfect alerts may be physically infeasible, and the alerts lose their value.

Target alerts by schedule position, not supplier average. Where zones are separated by slack, attention should focus on the final or least-buffered work areas; where the programme is compressed, earlier zones matter more.

Account for shared disruptions. Contractors should collect joint event records, including disruptions affecting several loads on the same day, and test booking rules under correlated scenarios before adopting them. Rules tuned to independent lateness may increase cost when weather, plant breakdowns or regional transport problems delay several loads together.

8 CONCLUSION AND RECOMMENDATIONS

This study quantified the value of receiving a delivery alert before a binding booking decision on a construction site with limited laydown space. The analytical result links the fee a contractor can justify to alert reliability, the off-site deferral charge and the reduction in contractual and tail-delay exposure, and shows that storage capacity governs whether acting on an alert is feasible at all. Exact experiments on two-stream networks with a shared crane show that early alerts change the most economical rule only when alert quality, storage capacity and decision timing align; that schedule slack determines where in the project alerts are valuable; and that dependence between disruptions can turn nominal gains into losses. A certified benchmark indicates that simple dispatch rules are often, though not always, optimal.

Three recommendations follow. Contractors should integrate alert timing and accuracy into delivery booking systems and evaluate responses with explicit break-even calculations. Site logistics planning should guarantee a fallback for storage-

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

constrained deliveries before relying on advance information. Booking rules should be stress-tested against correlated disruptions and changes in production rhythm.

9 LIMITATIONS AND FUTURE RESEARCH

The findings are subject to several limitations. First, all evidence derives from synthetic networks and stated probability models; no field data on alert accuracy, delivery lateness or gate operations were available, so the numerical values should be read as conditional on the assumed parameters rather than as empirical estimates. Second, the networks are stylised: a single crane, deterministic two-day durations, dedicated full-lot loads, a regular release pattern and a next-day return option for held trucks. Shared supplier failures, correlated alarms, restricted carrier capacity and congestion at transfer areas could change both feasible actions and measured benefits. Third, apart from the common-shock tests in Experiment 2, disruptions are independent across zones, and the common-shock mechanism is a stylised representation of dependence rather than an empirical claim about its frequency. Fourth, the policy experiments optimise within declared rule families, and the global optimum of the unrestricted formulation is computed only for four small networks. Fifth, the study uses a single ex ante CVaR criterion and does not address time-consistent risk measures under rolling reoptimisation.

These limitations suggest five research directions. Empirical validation should use dated booking changes, alert timestamps and accuracy, gate records and progress data from projects with booking systems or logistics centres to estimate sensitivity, false-alarm rates and disruption correlation. Larger instances with heterogeneous resources, stochastic durations and multiple cranes require decomposition or approximate dynamic programming, given that the exact training problem grows exponentially with the number of zones. Distributionally robust formulations could optimise over plausible correlation ranges rather than evaluating fixed policies, and could clarify the tension between robustness and information value. Endogenous acquisition of alerts, in which the contractor chooses how much to invest in tracking, would link this work to technology adoption decisions. Finally, organisational studies could examine whether the barriers to responding to early warnings identified by Haji-Kazemi et al. (2015) apply to delivery alerts and how booking authority should be allocated among procurement, logistics and site teams.

REFERENCES

[1] Akintoye, A. (1995). Just-in-time application and implementation for building material management. *Construction Management and Economics, 13*(2), 105–113. https://doi.org/10.1080/01446199500000013

[2] Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance, 9*(3), 203–228. https://doi.org/10.1111/1467-9965.00068

[3] Baghalzadeh Shishehgarkhaneh, M., Moehler, R. C., Fang, Y., Aboutorab, H., & Hijazi, A. A. (2024). Construction supply chain risk management. *Automation in Construction, 162*, Article 105396. https://doi.org/10.1016/j.autcon.2024.105396

[4] Ballard, G., & Howell, G. (1998). Shielding production: Essential step in production control. *Journal of Construction Engineering and Management, 124*(1), 11–17. https://doi.org/10.1061/(ASCE)0733-9364(1998)124:1(11)

[5] Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming* (2nd ed.). Springer. https://doi.org/10.1007/978-1-4614-0237-4

[6] Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics, 24*(2), 265–272. https://doi.org/10.1214/aoms/1177729032

[7] Bortolini, R., Formoso, C. T., & Viana, D. D. (2019). Site logistics planning and control for engineer-to-order prefabricated building systems using BIM 4D modeling. *Automation in Construction, 98*, 248–264. https://doi.org/10.1016/j.autcon.2018.11.031

[8] Dallasega, P., Rauch, E., & Linder, C. (2018). Industry 4.0 as an enabler of proximity for construction supply chains: A systematic literature review. *Computers in Industry, 99*, 205–225. https://doi.org/10.1016/j.compind.2018.03.039

[9] Ekeskär, A., & Rudberg, M. (2016). Third-party logistics in construction: The case of a large hospital project. *Construction Management and Economics, 34*(3), 174–191. https://doi.org/10.1080/01446193.2016.1186809

[10] Ergen, E., Akinci, B., & Sacks, R. (2007). Tracking and locating components in a precast storage yard utilizing radio frequency identification technology and GPS. *Automation in Construction, 16*(3), 354–367. https://doi.org/10.1016/j.autcon.2006.07.004

[11] Haji-Kazemi, S., Andersen, B., & Klakegg, O. J. (2015). Barriers against effective responses to early warning signs in projects. *International Journal of Project Management, 33*(5), 1068–1083. https://doi.org/10.1016/j.ijproman.2015.01.002

[12] Haji-Kazemi, S., Andersen, B., & Krane, H. P. (2013). A review on possible approaches for detecting early warning signs in projects. *Project Management Journal, 44*(5), 55–69. https://doi.org/10.1002/pmj.21360

[13] Hamlehvar, T., Aazami, A., Kebriyaii, O., & Heidari, A. (2026). Project scheduling problem integrated with resilient material procurement: A two-stage stochastic optimization approach. *Journal of the Operational Research Society*. Advance online publication. https://doi.org/10.1080/01605682.2026.2649385

[14] Hartmann, S., & Briskorn, D. (2010). A survey of variants and extensions of the resource-constrained project scheduling problem. *European Journal of Operational Research, 207*(1), 1–14. https://doi.org/10.1016/j.ejor.2009.11.005

[15] Herroelen, W., & Leus, R. (2005). Project scheduling under uncertainty: Survey and research potentials. *European Journal of Operational Research, 165*(2), 289–306. https://doi.org/10.1016/j.ejor.2004.04.002

[16] Janné, M., & Fredriksson, A. (2019). Construction logistics governing guidelines in urban development projects. *Construction Innovation, 19*(1), 89–109. https://doi.org/10.1108/CI-03-2018-0024

[17] Kolisch, R., & Sprecher, A. (1997). PSPLIB – A project scheduling problem library. *European Journal of Operational Research, 96*(1), 205–216. https://doi.org/10.1016/S0377-2217(96)00170-1

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

[18] Kouvelis, P., Chen, X., & Xia, Y. (2023). Managing material shortages in project supply chains: Inventories, time buffers, and supplier flexibility. *Production and Operations Management, 32*(11), 3717–3735. https://doi.org/10.1111/poms.14059
[19] Lu, H., Wang, H., Xie, Y., & Wang, X. (2018). Study on construction material allocation policies: A simulation optimization method. *Automation in Construction, 90*, 201–212. https://doi.org/10.1016/j.autcon.2018.02.012
[20] Nolz, P. C. (2021). Optimizing construction schedules and material deliveries in city logistics: A case study from the building industry. *Flexible Services and Manufacturing Journal, 33*(3), 846–878. https://doi.org/10.1007/s10696-020-09391-7
[21] Parchami Afra, A., Kheirkhah, A., & Ahadi, H. (2022). Systematic literature review of integrated project scheduling and material ordering problem: Formulations and solution methods. *Computers & Industrial Engineering, 173*, Article 108711. https://doi.org/10.1016/j.cie.2022.108711
[22] Pheng, L. S., & Chuan, C. J. (2001). Just-in-time management of precast concrete components. *Journal of Construction Engineering and Management, 127*(6), 494–501. https://doi.org/10.1061/(ASCE)0733-9364(2001)127:6(494)
[23] Polat, G., & Arditi, D. (2005). The JIT materials management system in developing countries. *Construction Management and Economics, 23*(7), 697–712. https://doi.org/10.1080/01446190500041388
[24] Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk, 2*(3), 21–41. https://doi.org/10.21314/JOR.2000.038
[25] Rockafellar, R. T., & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance, 26*(7), 1443–1471. https://doi.org/10.1016/S0378-4266(02)00271-6
[26] Sacks, R., Koskela, L., Dave, B. A., & Owen, R. (2010). Interaction of lean and building information modeling in construction. *Journal of Construction Engineering and Management, 136*(9), 968–980. https://doi.org/10.1061/(ASCE)CO.1943-7862.0000203
[27] Said, H., & El-Rayes, K. (2011). Optimizing material procurement and storage on construction sites. *Journal of Construction Engineering and Management, 137*(6), 421–431. https://doi.org/10.1061/(ASCE)CO.1943-7862.0000307
[28] Said, H., & El-Rayes, K. (2014). Automated multi-objective construction logistics optimization system. *Automation in Construction, 43*, 110–122. https://doi.org/10.1016/j.autcon.2014.03.017
[29] Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2021). *Lectures on stochastic programming: Modeling and theory* (3rd ed.). Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9781611976595
[30] Thomas, H. R., Riley, D. R., & Messner, J. I. (2005). Fundamental principles of site material management. *Journal of Construction Engineering and Management, 131*(7), 808–815. https://doi.org/10.1061/(ASCE)0733-9364(2005)131:7(808)
[31] Thunberg, M., & Fredriksson, A. (2018). Bringing planning back into the picture – How can supply chain planning aid in dealing with supply chain-related problems in construction? *Construction Management and Economics, 36*(8), 425–442. https://doi.org/10.1080/01446193.2017.1394579
[32] Tian, B., Zhang, J., Demeulemeester, E., & Liu, H. (2024). A chance-constrained optimization approach integrating project scheduling and material ordering to manage the uncertain material supply. *Computers & Operations Research, 166*, Article 106624. https://doi.org/10.1016/j.cor.2024.106624
[33] Tommelein, I. D. (1998). Pull-driven scheduling for pipe-spool installation: Simulation of lean construction technique. *Journal of Construction Engineering and Management, 124*(4), 279–288. https://doi.org/10.1061/(ASCE)0733-9364(1998)124:4(279)
[34] Vrijhoef, R., & Koskela, L. (2000). The four roles of supply chain management in construction. *European Journal of Purchasing & Supply Management, 6*(3–4), 169–178. https://doi.org/10.1016/S0969-7012(00)00013-7
[35] Zhang, L., Lin, X., & Wang, M. (2025). Material ordering based on activity delay in multi-project scheduling of prefabricated buildings. *KSCE Journal of Civil Engineering, 29*(4), Article 100062. https://doi.org/10.1016/j.kscej.2024.100062

(This work is licensed under a Creative Commons Attribution 4.0 International License.)