Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Toward Sustainable Smart Manufacturing under Maritime Disruptions: Timing-Aware Optimization and State Feedback

Islam El-Nakib

islam.elnakib@aast.edu

Arab Academy for Science, Technology and Maritime Transport

Research Article

Keywords: state feedback, sustainable smart manufacturing, production continuity, circular material use, maritime disruption, rolling-horizon optimization

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11317187/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Toward Sustainable Smart Manufacturing under Maritime Disruptions: Timing-Aware Optimization and State Feedback

Islam El-Nakib

Arab Academy for Science, Technology and Maritime Transport, Alexandria, Egypt
Corresponding author: islam.elnakib@aast.edu

Abstract
Maritime disruptions threaten manufacturing when materials become usable after production deadlines or when recovery capacity is exhausted. This study develops and evaluates a timing-aware decision-support framework linking inbound bookings, protective inventory, reserved transport capacity and feasible production. Matched comparisons distinguish the value of state-responsive replanning, residual protection and production coordination. Ten policies are evaluated on 120 matched ten-week computational episodes, supplemented by 216 factorial and 24 multi-size episodes. Rolling updates reduce mean cost by 19.47% relative to retaining the initial plan; residual protection reduces cost by 3.53% relative to nominal planning. Advantages over persistence and production-first planning remain small and uncertain. A separate 26-week, author-confirmed empirical factory ledger records 22.87% on-time dispatch despite 95.34% fulfilment by cutoff and 96.49% regular-hour utilization. An analytical circularity extension links qualified secondary inputs, delayed recovery, carried work in progress and shared production capacity to physically supported customer dispatch and explicit environmental budgets. A worked example shows that maximizing recovery volume can reduce feasible finished output. The contribution is a decision-structure and evaluation framework, not a new learning algorithm. The findings support deadline-feasible resource decisions; environmental benefits and broader transferability require further comparative evaluation.

Keywords: state feedback; sustainable smart manufacturing; production continuity; circular material use; maritime disruption; rolling-horizon optimization

1. Introduction

Sustainable manufacturing concerns the way resources are transformed into useful output, rather than efficiency within an isolated machine or planning function. A systems perspective is therefore essential when a local improvement increases another resource burden or shifts it to a different lifecycle stage [1]. During an inbound disruption, a manufacturer may protect production through additional inventory, alternative transport, reserved capacity or recovery of secondary materials. These responses are not environmentally equivalent. They may also fail operationally when material becomes available after the production requirement it was intended to protect.

Consider a manufacturer with enough transport capacity over a planning horizon but insufficient usable components in its first shortage week. A later booking may support eventual recovery without preserving the original customer deadline. Timely arrivals also cannot restore service when another component or the production resource remains binding. Circular material availability is similarly insufficient: recovered inputs must meet qualification and timing requirements, and their processing must not displace the capacity needed for timely finished production. Cost, material volume and useful output therefore require separate evaluation.

Historical maritime evidence illustrates the timing mechanism. UN Trade and Development reports a Shenzhen–Rotterdam example of approximately 31 sailing days through Suez and 41 days around the Cape of Good Hope [2]. Arvis, Rastogi and Ulybina describe seven to ten additional days for typical Europe–Asia diversions [3]. These are published historical illustrations, not current service quotations or estimates for the factory examined here. Their relevance is that a route change can cross a production-use deadline, making eventual delivery and timely material availability different planning outcomes.

Artificial intelligence and optimization can support this decision, but combining prediction, simulation and planning is already established. Hybrid master-planning models, learning-assisted production control and learned-scenario inventory routing provide close architectural precedents [4, 5, 6]. Accordingly, the central question is not whether a learning model can be connected to an optimizer. It is which capability changes decisions usefully once arrival timing, complementary materials, reserve commitments and physical production are represented consistently. Forecast accuracy and downstream operating value must be evaluated separately.

This study develops and evaluates a timing-aware manufacturing decision-support framework that links inbound protection to material availability before production-use deadlines. Matched policy comparisons distinguish the incremental value of residual protection, production coordination and state-responsive replanning within a common execution model. An analytical extension incorporates qualified secondary inputs, delayed recovery and competition for production capacity, establishing a basis for evaluating circular production under explicit environmental and customer-service constraints.

The sustainability argument connects these decisions to resource conversion rather than equating resilience with environmental improvement. The operational evidence examines punctual output, protective-stock commitment and

recovery headroom. The analytical extension then distinguishes recoverable feedstock, qualified secondary material and actual consumption, consistent with the coordination of reverse flows and forward production [7, 8]. Its worked example shows how excessive recovery can reduce feasible output when processing competes for scarce production hours.

Five research questions organize the analysis. RQ1 asks how much the specified rolling update protocol contributes relative to an unchanged initial plan. RQ2 asks whether added forecasting and residual protection improve decision outcomes against simpler controls. RQ3 asks how the contribution of joint production–inbound planning varies across operating conditions and tested configurations. RQ4 asks what linked factory records reveal about punctual dispatch, accumulated backlog and recovery headroom. RQ5 asks how timing-qualified secondary materials and explicit environmental constraints should extend this operational architecture, and what further evidence is necessary to evaluate them.

Sections 2–5 position and specify the operational decision framework. Sections 6–7 report the computational comparisons and separate factory diagnostics. Section 8 develops the analytical circularity extension, including complete dispatch, recovery-pipeline and environmental-accounting balances. Section 9 discusses implications and limits, and Section 10 concludes. Supplementary Information S1 contains the prospective benchmarking and environmental evaluation protocol; its design quantities are not reported as completed experiments.

2. Related research and contribution positioning

2.1. Disruption protection and production conversion

The conditional value of inventory and backup capacity is a long-established issue in disruption management. Tomlin analyzes mitigation and contingency strategies with unreliable supply, a more expensive reliable alternative and capacity restrictions [9]. Supplier selection, inventory and uncertain replenishment are also jointly addressed in sourcing and network-design research [10, 11]. These studies establish the importance of considering protection alternatives together; the present work does not claim to introduce that general principle.

The focal decision boundary is narrower than strategic sourcing or network redesign. A manufacturer books inbound services, reserves eligible capacity, positions protective stock before the episode and revises aggregate production. Multiple complementary materials must become available in time for conversion through the production resource. The reconfigurable supply-chain perspective, aggregate production-resilience analysis and production–transportation collaboration research provide related mechanisms at different decision levels [12, 13, 14]. The present model makes the arrival-to-use connection explicit within a common optimization and execution environment.

Maritime disruption studies also connect transport to procurement, production, inventory and sales. Li and colleagues use system dynamics to examine these interactions during a Suez Canal blockage [15]. The present study instead optimizes bookings and production in a generic computational benchmark. It does not inherit industrial calibration merely by using the same maritime context. Similarly, terminal recovery and geopolitical reconfiguration studies address adjacent system boundaries rather than validating a manufacturer-level policy [16, 17].

2.2. Forecasting, learning and downstream decision value

Predictive and prescriptive analytics distinguish the quality of a prediction from the usefulness of a decision made with it [18]. Smart Predict-then-Optimize incorporates downstream optimization structure into prediction training, while the decision-focused learning literature considers this broader class of approaches [19, 20]. Here, prediction models are fitted using statistical losses; realized operational cost is used subsequently to select a protection multiplier. The implemented method is therefore forecast-informed optimization with decision-based validation, not end-to-end decision-focused learning.

Badakhshan and Ball combine simulation, machine learning and optimization for disrupted master planning. Assid, Gharbi and Pellerin develop learning-assisted production and replenishment control under irregular material supply. Jia, Schrotenboer and Chen combine learned scenarios and optimization for online inventory routing [4, 5, 6]. These are close precedents. The assembly of learning, uncertainty scenarios and rolling optimization is not, by itself, the novelty claim of this manuscript. The relevant distinction is the timing-qualified manufacturer decision boundary and the explicit interpretation of matched capability comparisons.

Simpler forecasting controls remain essential. Arora and colleagues address demand forecasting during and after disruption, supporting the broader need to compare statistical approaches alongside machine learning [21]. Demand and route-capacity shocks remain different prediction targets. Accordingly, the present comparisons include linear autoregression with an exogenous warning, persistence and historical means, as well as gradient boosting and random forests. Each predictor has its own calibration residuals and selected protection multiplier; comparisons across predictors therefore evaluate complete information pipelines rather than isolate prediction accuracy alone.

2.3. Recovery planning, resource competition and environmental trade-offs

Energy- and resource-efficient manufacturing requires consistent boundaries across processes and systems [1]. Closed-loop value recovery and circular resource strategies connect reverse flows to forward production rather than treating returned

volume as an outcome by itself [7, 8]. More specific production-planning models establish the mechanisms against which a circularity extension should be positioned.

Poles [22] uses system dynamics to examine production and inventory control with remanufacturing, including capacity planning, lead times and backlog. Flapper, Gayon and Lim [23] study stochastic manufacturing and remanufacturing on a single shared server, using a Markov decision formulation to investigate operating policies. These studies establish that recovery competes with forward production for resources; shared capacity is not a new principle introduced here. The present extension makes usable release and original customer deadlines explicit within a finite-horizon material-allocation model.

Reddy and Kumar [24] integrate capacity and inventory planning in a hybrid manufacturing–remanufacturing system under uncertain demand, returns and yield. Their work is a close precedent for evaluating recovery against available capacity, rather than presuming that all returned material can be used. In the present analytical extension, recovery yield and release lag are specified coefficients; the reported operational experiments do not test uncertain recovery yields. Independent yield scenarios and qualification outcomes therefore belong to the subsequent evaluation, not the existing evidence.

Lahmar and colleagues [25] formulate multi-objective production planning for a multistage manufacturing–remanufacturing system with grade-based product distinctions, minimizing cost and carbon emissions. Their model demonstrates why environmental performance requires its own objective or accounting rather than inference from operating cost. Here the extension specifies environmental ceilings alongside an original-due-cohort service floor and separates virgin acquisition, virgin consumption, recovery work in progress and waste. No numerical advantage over their method is claimed.

Hoffmann and colleagues [26] organize functional requirements and performance measures for remanufacturing planning and control, with an industrial application. This complements the mathematical studies by emphasizing the information needed to interpret variable returned-product conditions and process performance. Accordingly, qualification limits, gross-input coefficients, yields and release events must have identifiable sources before they are used in the proposed model.

These precedents establish the component mechanisms. The distinctive emphasis here is their connection to deadline-feasible manufacturer output and state-preserving policy evaluation: material may be recovered successfully yet arrive too late, displace finished-production capacity or remain unused at the endpoint. The extension formalizes that decision boundary; its environmental effectiveness still requires execution under documented inputs. The GHG Protocol Product Life Cycle Accounting and Reporting Standard [27] informs the inventory boundary and allocation discipline, without converting the operational benchmark into a completed lifecycle assessment.

2.4. Decision-structure contribution and evidence boundary

Finite-scenario feasibility, residual stress construction and rolling state feedback have substantial precedents. Robust optimization protects against a specified uncertainty representation; a small set of adverse trajectories does not create distribution-free or distributionally robust reliability guarantees [28, 29]. Related work on robust inventory routing, online inventory optimization and digital supply-chain twins further establishes the importance of stating the actual architecture and its information assumptions [30, 31, 32]. The current implementation has no live industrial data connection or state-estimation subsystem and is not described as a deployed digital twin.

Table 1 positions the closest research against the contribution actually examined. The comparison is a structured narrative review, not an exhaustive database-screened review, and it does not infer the absence of a feature from another study merely because that feature is not discussed here.

Table 1. Closest research and the comparison boundary

Research	Documented focus	Position of the present study
Tomlin [9]; Saputro et al. [10]; Wang et al. [11]	Disruption mitigation, sourcing and inventory decisions	Examines short-horizon arrival-to-production feasibility rather than claiming a new principle of joint protection
Badakhshan and Ball [4]	Hybrid modeling for disruption-sensitive master planning	Close architectural precedent; the hybrid combination is not claimed as new
Assid et al. [5]	Learning-assisted production/replenishment control	Relevant external benchmark candidate; not numerically reimplemented in the reported results
Jia et al. [6]	Learned scenarios and online inventory-routing optimization	Establishes a related prediction-to-scenario architecture; no numerical superiority is claimed

Research	Documented focus	Position of the present study
Duflou et al. [1]	Energy and resource efficiency across processes and systems	Motivates separate physical and environmental accounting
Guide and Van Wassenhove [7]; Bocken et al. [8]	Closed-loop value recovery and circular resource strategies	Motivates timing-qualified secondary inputs and recovery–production resource competition
Poles [22]	Remanufacturing lead times, capacity and inventory dynamics	Connects recovery-release timing to original production-use deadlines
Flapper et al. [23]	Stochastic manufacturing and remanufacturing with a shared server	Shared capacity is an established mechanism, not a new standalone claim
Reddy and Kumar [24]	Capacity/inventory decisions with uncertain demand, returns and yield	Motivates independent recovery-yield and capacity scenarios; not an implemented comparator
Lahmar et al. [25]	Grade-based hybrid production with cost and carbon objectives	Positions explicit environmental accounting; no numerical superiority claimed
Hoffmann et al. [26]	Requirements and performance measures for remanufacturing control	Connects qualification and process information to interpretable measures
Present study	Arrival-feasible protection, controlled capability comparisons and a circularity extension	Separates demonstrated operational effects, descriptive case diagnostics and analytically specified environmental mechanisms

The contribution concerns the connection between protective actions, material-use deadlines and feasible customer-serving output, together with matched evaluation of the planning capabilities involved. State-responsive replanning has the strongest demonstrated effect within the tested architecture. The analytical extension adds explicit recovery and service-accounting mechanisms; it does not supply an evaluated environmental intervention. Table 2 maps each claim to its supporting evidence.

Table 2. Contribution-to-evidence map

Contribution	Evidence in this article	Claim boundary
Timing-qualified material protection	Material-prefix condition; reserve diagnostic; common physical execution	Feasibility depends on the stated model and scenario set
Incremental value of planning capabilities	Ten policies; matched protection, coordination and feedback contrasts	Predictor comparisons change complete pipelines; effects are not an additive causal decomposition
Resource and service diagnostics	Protective-stock use; punctual-output ratios; 26-week factory-record summaries	Operational indicators, not measured energy, emissions or waste reductions
Circular-manufacturing extension	Secondary-material balances; recovery–production trade-off; environmental budgets	Analytical development; expanded policy experiments remain prospective

These distinctions turn the contribution into a set of testable claims. They also make the negative and inconclusive comparisons informative: a sophisticated forecasting pipeline that does not improve the relevant operating outcomes should not be preferred merely because it uses machine learning.

3. Decision setting, evidence layers and application boundary

3.1. Manufacturer-level planning problem

The implemented benchmark represents a manufacturer converting imported materials into product families through a shared production resource. At the beginning of each week, the policy observes usable raw and finished inventory, outstanding demand, already-dispatched arrivals, the current warning and the previous route-shock vector. It then books inbound quantities, activates services, purchases eligible reserve capacity and announces production. Realized transport capacity can curtail dispatch. Shipment arrival dates become known at dispatch and remain fixed thereafter in this implemented version.

Production consumes only material available in the current week. Initial protective stock is a pre-episode investment: its purchase and positioning costs are charged, but the upstream lead time required to establish it is outside the horizon. Production and material quantities are divisible, while route activation is binary. Machine sequencing, setups, minimum lots, substitution, explicit supplier failure, contractual booking freezes and post-dispatch arrival revisions are outside the executed core. These exclusions distinguish aggregate continuity planning from detailed shop-floor scheduling.

A simple illustration clarifies the timing requirement. Using the historical 31-day and 41-day maritime transit example, suppose a manufacturer needs a component on days 35–40, requires 100 components per day and has 250 usable components at the start of day 35. All factory quantities in this illustration are assumptions. The six-day requirement is 600 components, and the first deficient prefix ends on day 37. Another 350 components could cover the material requirement only if they arrive before the relevant deficient prefixes. A shipment arriving on day 41 cannot prevent that shortfall. If production is limited instead to 60 units per day, at most 360 units can be completed in the window; only 110 additional components are needed to support that maximum. Material volume and feasible output are therefore not interchangeable.

3.2. Separation of evidence

The study uses three evidence categories: computer-generated benchmark histories for controlled policy comparisons; a separate author-confirmed empirical factory ledger for linked operational diagnostics; and an analytical circularity formulation. Descriptive ratios in Sections 6–7 are derived from reported summary values and retain their rounding limitations. They do not increase the number of independent paths or provide episode-level uncertainty estimates. The environmental extension and broader benchmark design in Supplementary Information S1 have not been executed as additional policy experiments.

Table 3 identifies the experimental denominators. The main, factorial and multi-size banks contain 360 distinct final-evaluation paths, but they address different questions and must not be pooled as homogeneous replications. Validation adds 24 paths used only for protection selection.

Table 3. Reported computational banks and their purposes

Bank	Paths	Settings per path	Policy runs	Purpose
Validation	24	Six predictor/set variants; four multipliers	576	Protection selection
Main evaluation	120	Ten policies	1,200	Matched policy comparison
Factorial evaluation	216	Two policies	432	Operating-condition dependence
Multi-size extension	24	Three policies	72	Limited configuration transfer
Total, including validation	384	Separate banks	2,280	Not a pooled statistical sample

Fifteen additional initial-solve cases measure formulation size and computational time; they are not policy episodes. The prospective environmental and broader-structure design in Supplementary Information S1 is excluded from all completed-experiment counts.

3.3. Empirical factory-record layer

The empirical factory analysis uses IJPR__Research_Data.xlsx, containing linked operational records identified by the author as original, reviewed and validated. Validation covers transaction linkage and internal consistency, rather than independent authentication against the originating information system. Appendix B documents the legacy provenance labels, relative-time transformation and the distinction among operational observations, research calculations and model assumptions.

The ledger covers 26 relative weeks, two products, three materials, two routes and one production line. Product P01 uses 2 kg of M01 and 1 kg of M02, with 0.50 production hours per piece. Product P02 uses 1.5 kg of M01 and 1.2 kg of M03, with 0.65 hours per piece. Regular capacity is 110 hours per week, reduced by eight maintenance hours in weeks 6 and 18. Opening raw stocks are 450, 150 and 180 kg; opening finished stocks are 15 and 10 pieces. Opening overdue demand is 35 pieces, and three committed inbound quantities are 320, 100 and 110 kg.

The recorded baseline uses weekly order-up-to replenishment, production constrained by available materials and hours, and oldest-due-order-first dispatch. Recorded forecasts are treated as research calculations and costs as rate-based accruals. The

case is not silently substituted into the separate benchmark. It supplies evidence about timing, backlog and recovery capacity, not an implemented comparison between the proposed policy and factory practice.

4. Timing-aware operational formulation

4.1. State, notation and decision timing

At decision week $t \in \{0, \dots, T-1\}$, let $h = 0, \dots, H-1$ index the remaining horizon, where $H = T - t$. Materials, products, routes and planning scenarios are indexed by m, k, r, s, respectively. Table 4 defines the principal operational notation. A local horizon is resolved after each realized transition; future actions within a solve are common across scenarios.

Table 4. Principal notation for the executed operational model

Symbol	Definition
I_m^0, F_k^0, B_k^0	Observed raw inventory, finished inventory and backlog at the current solve
A_{hm}	Already-dispatched material scheduled to become usable in local week h
d_{hk}, a_{mk}, f_k, C	Known demand, material requirement per product, resource use and weekly capacity
$x_{hmr}, z_{hr}, e_{hr}, b_m$	New bookings, binary activation, purchased reserve and initial optional protection
$q_{hk}, F_{hk}, B_{hk}, \Delta_{hk}$	Planned production, finished inventory, backlog and absolute plan revision
$I_{shm}, K_{shr}, G_{sh}, \ell_{sur}$	Scenario raw stock, route capacity, shared fallback capacity and booking lead time
$\mathcal{R}_{alt}, \overline{e}_r, \overline{E}, \overline{b}_m$	Alternative services, route/total reserve limits and protective-stock bounds
S, γ, κ, π_k	Number of planning trajectories, protection multiplier, revision penalty and backlog penalty

Raw inventory varies by scenario because arrivals vary. Bookings, production, reserve and the other operating decisions are common across the scenarios in a given solve. This is an open-loop finite-scenario plan executed repeatedly with observed state feedback, not a multistage scenario tree with contingent future decisions.

4.2. Material and product balances

A booking dispatched in local week u becomes usable in week $u + \ell_{sur}$. With $I_{s,-1,m} = I_m^0$, the raw-material balance is

$$I_{shm} = I_{s,h-1,m} + A_{hm} + \mathbf{1}_{\{h=0\}} b_m + \sum_{u,r:\, u+\ell_{sur}=h} x_{umr} - \sum_k a_{mk}\, q_{hk}. \qquad (1)$$

All raw inventories are nonnegative. The initial protection decision satisfies $0 \le b_m \le \overline{b}_m$ at the first episode solve and is fixed to zero at later solves. The arrival indicators are precomputed, so the balance is linear. Finished-product balance and production capacity are

$$F_{hk} - B_{hk} = F_{h-1,k} - B_{h-1,k} + q_{hk} - d_{hk}, \qquad (2)$$

$$\sum_k f_k\, q_{hk} \le C. \qquad (3)$$

The pre-horizon finished and backlog states equal their observed values. Positive holding and backlog costs permit an optimum without simultaneous finished stock and backlog for the same product. The integrated policy can produce ahead or recover backlog when materials and capacity permit; it does not impose a fixed weekly demand cap on output.

4.3. Route capacity and contractual reserve accounting

For each scenario, bookings on an activated route are limited by its scenario capacity plus purchased reserve. Alternative routes also share a fallback pool:

$$\sum_m x_{hmr} \le K_{shr} z_{hr} + e_{hr}, \quad 0 \le e_{hr} \le \overline{e}_r z_{hr}, \quad \sum_r e_{hr} \le \overline{E}. \qquad (4)$$

$$\sum_{m} \sum_{r \in \mathcal{R}_{alt}} x_{hmr} \leq G_{sh} + \sum_{r \in \mathcal{R}_{alt}} e_{hr}\,, \quad \forall s, h. \qquad (5)$$

A reserved alternative-service unit increases both that route's usable limit and its contribution to the shared pool. Equations (4) and (5) constrain the same cargo; they do not authorize two independent uses of one reserve unit. Optimization and execution use this identical contractual interpretation. A reservation guaranteeing a route slot but not pooled access would require a different model.

Bookings arriving after the horizon in every planning scenario are disabled. A booking useful in one scenario can still arrive beyond the endpoint in another scenario or under realized conditions. Paid but unused terminal inventory and pipeline retain their costs without salvage credit in the reported experiments.

4.4. Economic objective and production-plan stability

Let $c_m^M, c_r^X, c_r^Z, c_r^E, c_k^Q$ denote material, freight, activation, reserve and production costs. Let c_m^I, c_k^F denote holding costs and c_m^B the additional positioning cost of protective stock. Booking and protection expenditure is

$$J_{book} = \sum_{h,m,r} (c_m^M + c_r^X)\, x_{hmr} + \sum_{h,r} (c_r^Z z_{hr} + c_r^E e_{hr}) + \sum_{m} (c_m^M + c_m^B)\, b_m. \qquad (6)$$

The implemented objective is

$$\min J_{book} + \sum_{h,k} \left(c_k^Q q_{hk} + c_k^F F_{hk} + \pi_k B_{hk} + \kappa \Delta_{hk}\right) + \frac{1}{S} \sum_{s,h,m} c_m^I\, I_{shm}. \qquad (7)$$

For the previously announced production plan q^{ref}, revision is represented by

$$\Delta_{hk} \geq q_{hk} - q_{hk}^{ref}, \qquad \Delta_{hk} \geq q_{hk}^{ref} - q_{hk}. \qquad (8)$$

Revision variables are fixed to zero at the initial solve. Route activations are binary; remaining variables are nonnegative and continuous. The factor $1/S$ averages the raw-inventory holding term across constructed scenarios. It is not an estimated probability. The model enforces feasibility across its finite scenario set but neither minimizes worst-case total cost nor establishes calibrated out-of-sample reliability. No additional terminal backlog penalty or terminal salvage value is included.

4.5. Physical execution and feedback

Only current bookings, reserves and production actions are executed. Actual route capacities follow the realized shock. Excess bookings are reduced proportionally across materials on each route; a second proportional reduction applies if alternative-service dispatch exceeds the realized shared pool including eligible reserves. Unshipped quantities are cancelled. Material and freight charges apply to actual dispatch, while activation and reserve fees remain payable.

When material availability prevents planned production, an allocation linear program maximizes weighted executed output subject to current material and resource limits:

$$\max \sum_{k} \pi_k\, q_k^{act}, \qquad 0 \leq q_k^{act} \leq q_{0k},$$
$$\sum_{k} a_{mk}\, q_k^{act} \leq I_m^{available}, \qquad \sum_{k} f_k\, q_k^{act} \leq C. \qquad (9)$$

Physical execution cannot exceed the announced target. Finished stock serves older backlog before current demand. Production, inventory and backlog costs use realized quantities and states; realized revision cost applies to the current target. Inventories, backlog and the immutable dispatched pipeline then become the next state. Replanning uses only information available at that decision point.

This common execution model is central to the comparisons. A policy is not credited with production merely because the optimizer requested it, or with a shipment merely because it was booked. The same accounting also identifies the activity data that an environmental extension must use: realized dispatch, actual processing and usable inventory, rather than nominal promises.

4.6. Timing and conversion diagnostics

Summing the raw balance through a prefix ending at h yields the necessary and sufficient condition for nonnegative raw inventory for a fixed plan and scenario:

$$\sum_{j=0}^{h}\sum_{k} a_{mk}\, q_{jk} \le I_m^0 + b_m + \sum_{j=0}^{h} A_{jm} + \sum_{u,r:\, u+\ell_{sur}\le h} x_{umr}\,. \qquad (10)$$

Necessity follows by cancellation of intermediate stocks. Conversely, defining inventory as cumulative available material minus cumulative consumption gives nonnegative inventory whenever the inequality holds at every prefix. Transport and factory constraints must still hold separately. A shipment arriving after a deficient prefix contributes nothing to that prefix. This is a conservation-based planning diagnostic, not a new general theorem of inventory theory.

For a single-product reduction with demand D, usable material A, requirement $a > 0$, output-capacity limit C' and timely material conversion $g \ge 0$ per reserve unit,

$$q(e) = \min\{D, C', (A + ge)/a\}. \qquad (11)$$

Reserve adds output only when material is binding and $g > 0$. If additional output has marginal value w and reserve costs c_e per unit, its local net value is $wg/a - c_e$ until another bound becomes active. Complementary materials and shared resources make this a local diagnostic rather than a universal reservation threshold. Section 8 extends the same logic to qualified secondary inputs and recovery capacity.

5. Forecasting, policy comparisons and computational design

5.1. Forecast and residual construction

The features are the current warning and lagged route shocks. Gradient boosting (GB) uses 80 depth-two trees, a learning rate of 0.05 and Huber loss. Random forests (RF) use 120 trees, maximum depth five and at least four observations per leaf. Linear autoregression with an exogenous warning (ARX) is fitted by least squares. Persistence uses the last observed route state, and the historical-mean model uses the training-period route means. Settings are fixed before final evaluation.

A 700-period computer-generated history supplies 450 training observations and 250 residual-calibration observations. Within-block lag construction gives 449 and 249 pairs, respectively. Beyond the current week, forecasts recursively use the preceding predicted route state and the training-period mean warning; future realized warnings are not supplied. A separate history provides 300 one-step predictions for accuracy assessment. Its population differs from the heterogeneous policy-test strata, so the forecast-error table is not interpreted as a causal regression of prediction error on decision cost.

For route-vector observation y_i and prediction $\hat{y}_i$, adverse residuals are $r_i = \max(0, y_i - \hat{y}_i)$ componentwise. Two k-means clusters are fitted, and the componentwise 90th residual percentile in each forms a prototype v_j. Fixed protection uses the entire calibration block. The warning-conditioned variant partitions warnings into tertiles, fits two prototypes within each group and selects the group using the current warning. The nominal and two adverse trajectories are

$$\mathcal{U}_t(\gamma) = \{(\hat{y}_h)_{h=0}^{H-1},\ (y_h^{(1)})_{h=0}^{H-1},\ (y_h^{(2)})_{h=0}^{H-1}\},$$
$$y_h^{(j)} = \mathrm{clip}_{[0,0.95]}\big(\hat{y}_h + \gamma\, 0.9^h v_j\big), \quad j = 1,2. \qquad (12)$$

The decay factor 0.9 is a modeling choice. Componentwise cluster percentiles do not establish 90% joint trajectory coverage. Warning conditioning changes residual selection, not online retraining or the globally selected multiplier. The nominal policy uses only the forecast trajectory.

5.2. Validation and comparators

Positive multipliers are selected from {0.5,0.75,1.0,1.5} by mean realized cost on 24 separate validation paths. The selected multiplier is 1.0 for fixed GB, warning-conditioned GB and ARX; 0.75 for RF; and 0.5 for persistence and historical means. The nominal policy is evaluated separately and is not selected and relabeled as a robust policy. Because residuals and multipliers differ across predictors, these comparisons do not hold statistical risk coverage constant.

Table 5 distinguishes matched ablations from complete-pipeline comparisons. GB-RO is the reference architecture, not an assumed best policy. Every policy shares the execution environment, initial state and cost accounting.

Table 5. Ten reported policies and the interpretation of their comparison

Policy	Forecast and protection	Architecture and comparison
GB-RO	GB; fixed residual prototypes	Integrated rolling reference
GB-Regime	GB; warning-tertile residual prototypes	Uncertainty-construction alternative
GB-Nominal	GB; nominal trajectory only	Matched protection ablation

Policy	Forecast and protection	Architecture and comparison
ARX-RO	ARX; own residuals and selected multiplier	Complete alternative pipeline
RF-RO	RF; own residuals and selected multiplier	Complete alternative pipeline
Persistence-RO	Last route state; own calibration	Simple information-pipeline control
Historical-RO	Training means; own calibration	No contextual warning prediction
Production-first	Same GB scenarios and multiplier	Production target set first; logistics then optimized
GB-Static	Same initial GB model and scenarios	Initial plan retained; matched feedback ablation
Perfect-information	Realized future route path	Non-implementable rolling diagnostic

The production-first comparator first optimizes production and finished-product inventory/backlog under factory capacity while assuming sufficient material. A second logistics-production solve bounds production by that target. It can respond to accumulated backlog at the next decision point. It is therefore stronger than a rule simply capping production at current demand, but remains an internal restriction rather than a replication of a leading published sequential algorithm. Perfect information is used within the same rolling architecture and is not a proved global lower bound on realized policy cost.

5.3. Benchmark parameters and disturbance process

The main benchmark has two products, two materials and three transport services over ten weeks. Weekly demand is $(12,8)$, material coefficients are $(2,1)$ and $(1,2)$ for the two products, resource requirements are $(1,1.4)$ and factory capacity is 28. Initial raw stock covers one week; finished stock, backlog and the dispatched pipeline start at zero. Table 6 states the remaining assumptions in normalized units.

Table 6. Principal benchmark assumptions

Quantity	Value
Route capacities; sensitivities	$(60,38,14)$; $(0.65,0.50,0.30)$
Base lead times; shock increments	$(1,2,1)$; $(2,2,0)$ with upward integer rounding
Shared alternative pool	$45\left[1-0.35\max_{r\in\mathcal{R}_{alt}} y_r\right]$, plus eligible reserve
Reserve limits; price	12 per route; 20 total per week; 1.2 per purchased unit
Material; freight costs	$(2,3)$; $(1,1.65,3.2)$
Activation; production costs	4 per active service-week; 1 per output unit
Raw; finished holding costs	0.1; 0.3 per unit-period
Backlog; revision penalties	$(24,14)$ per unit-period; 0.3 per changed output unit
Optional initial protection	Maximum $(64,56)$; material cost plus 1.5 positioning cost
Baseline raw inventory	$(32,28)$

Warnings follow a Beta(1.5,3.5) distribution. A clipped autoregressive latent state begins at 0.2, with Gaussian innovation standard deviation 0.11. Route shocks combine this state with loadings equally spaced from 1.0 to 0.45 and independent noise of standard deviation 0.075. Severity multiplies the route vector before clipping to $[0,0.95]$. Training persistence is 0.65 and warning loading is 0.35. During evaluation, warning loading equals one minus persistence. This preserves the untruncated latent mean, although clipping can alter the realized distribution. Each episode discards a 61-row prefix and uses its last shock as the initial observed lag.

Route capacity is $K_r(y) = \overline{K}_r(1-\alpha_r y_r)$ and lead time is $\ell_r^0 + \lceil \beta_r y_r \rceil$. Capacity remains positive throughout the specified shock range; complete corridor closure is not represented. The ceiling introduces discrete timing thresholds. Warning relationships and capacity coefficients are modeling assumptions, not estimates from the factory ledger.

5.4. Experimental scope and inference

The main bank crosses persistence $(0.45,0.65,0.85)$ with severity $(0.8,1.0,1.2,1.4)$, with ten independent paths in each of 12 strata. Validation uses two separate paths per stratum. The factorial bank crosses initial coverage $(0.5,1.0,2.0)$ weeks, capacity multipliers $(0.90,1.00,1.20)$ and severity $(0.8,1.1,1.4)$ at persistence 0.65. Eight replications in each of 27 cells yield 216 paths. All cells are retained, including unfavorable effects.

Five initial solves are evaluated at each of three sizes (K, M, R, T): $(2,2,3,10)$, $(5,4,4,12)$ and $(8,6,5,14)$. A separate policy extension compares GB-RO, ARX-RO and Production-first on twelve paths at each of the two larger sizes. Each size has one generated material structure; demand is known and sinusoidal with 12% amplitude. Forecasts are refitted for the route dimension, but protection multipliers transfer without retuning. Size, structure, horizon and parameters change together, so these checks cannot identify a pure system-size effect.

Realized cost includes optional protection, actual procurement and freight, activation, reserves, production, holding, backlog and executed revisions. End fill is one minus terminal backlog divided by total episode demand. On-time service counts current demand served after older backlog is cleared. The empirical upper-decile mean uses the twelve highest costs among 120 paths and is not a population tail-risk guarantee.

The main cost estimand equally weights the twelve strata. Paired differences use common realized paths. Ten thousand stratified bootstrap resamples resample ten paths within each stratum before recombining means. Pointwise 95% percentile intervals and Bonferroni-adjusted intervals for seven designated contrasts are reported. The source analysis also provides Holm-adjusted paired mean tests as secondary diagnostics. These are conditional Monte Carlo comparisons, not a preregistered confirmatory trial. Uncertainty from refitting the entire forecasting pipeline, retuning protection or transferring to another industry is not represented.

5.5. Implementation checks and reproducibility

The reported implementation uses SciPy's mixed-integer linear programming interface to HiGHS, a relative gap target of 10^{-6} and an eight-second policy-solve limit; initial-solve scaling permits 20 seconds. Incumbents are checked against bounds, linear constraints and integrality before execution. Thirteen regression checks cover prefix feasibility, arrival timing, zero capacity, deterministic static/rolling consistency, cost reconciliation and reserve pooling.

The source archive reports 21,951 main production–logistics MILP solves, all with status zero; maximum linear/bound violation is 3.71×10^{-10}, maximum integrality deviation is 9.47×10^{-12} and maximum reported relative gap is 9.82×10^{-7}. Target and execution allocation linear programs are not included in that MILP count. These are diagnostics of the reported operational implementation. Appendix A identifies its driver, seed specifications and the separately defined summary-level calculations.

6. Computational findings and resource-oriented interpretation

6.1. Forecast accuracy and operating performance

Table 7 reports independent one-step forecast accuracy under the training generator. All five predictors also enter the downstream policy comparison, so a stronger simple forecast is not omitted.

Table 7. Independent one-step route-shock accuracy across 300 predictions

Forecast	Mean MAE	Mean RMSE	Route 1 MAE	Route 2 MAE	Route 3 MAE
Linear ARX	0.0852	0.1088	0.1025	0.0830	0.0700
Random forest	0.0881	0.1130	0.1065	0.0861	0.0717
Gradient boosting	0.0885	0.1133	0.1073	0.0866	0.0715
Persistence	0.1054	0.1369	0.1247	0.1031	0.0884
Historical mean	0.1125	0.1428	0.1428	0.1136	0.0811

ARX has the lowest reported mean absolute error (MAE) and root mean squared error (RMSE). This does not by itself determine the policy ranking, because calibration, selected protection and operating-stratum distributions differ. The comparison supplies evidence against assuming that the machine-learning alternatives must be preferable.

Table 8 reports main evaluation means on the same 120 held-out paths. Cost is normalized, service is expressed as a percentage, and initial buffer is measured in the benchmark's aggregate material units. Standard deviations are retained in Appendix A.

Table 8. Main policy outcomes on the common 120-path evaluation bank

Policy	Mean cost	End fill (%)	On-time (%)	Initial buffer
GB-RO	2,398.53	96.96	96.94	118.87
ARX-RO	2,391.80	97.26	97.24	115.64
RF-RO	2,410.69	97.36	97.20	110.51
Persistence-RO	2,402.85	97.61	97.16	100.67
Historical-RO	2,362.19	96.56	96.51	120.00
GB-Regime	2,403.13	97.06	97.03	118.29
GB-Nominal	2,486.28	97.82	93.95	86.30
Production-first	2,399.76	96.93	96.87	118.87
GB-Static	2,978.47	91.94	83.08	118.87
Perfect-information	2,239.89	98.53	98.53	90.30

GB-RO is not the least-cost implementable policy. Historical-RO has mean cost 2,362.19 and on-time service 96.51%, compared with 2,398.53 and 96.94% for GB-RO. ARX-RO has lower mean cost and higher mean service than GB-RO, although its paired cost interval includes zero. The initial protective-stock decision is prominent: GB-RO purchases 118.87 of 120 allowed aggregate units, while GB-Nominal purchases 86.30 and Historical-RO purchases the entire allowance.

Fig. 1 displays all ten mean costs from a zero origin. It makes the larger static-plan penalty visible without magnifying the much smaller differences among rolling policies.

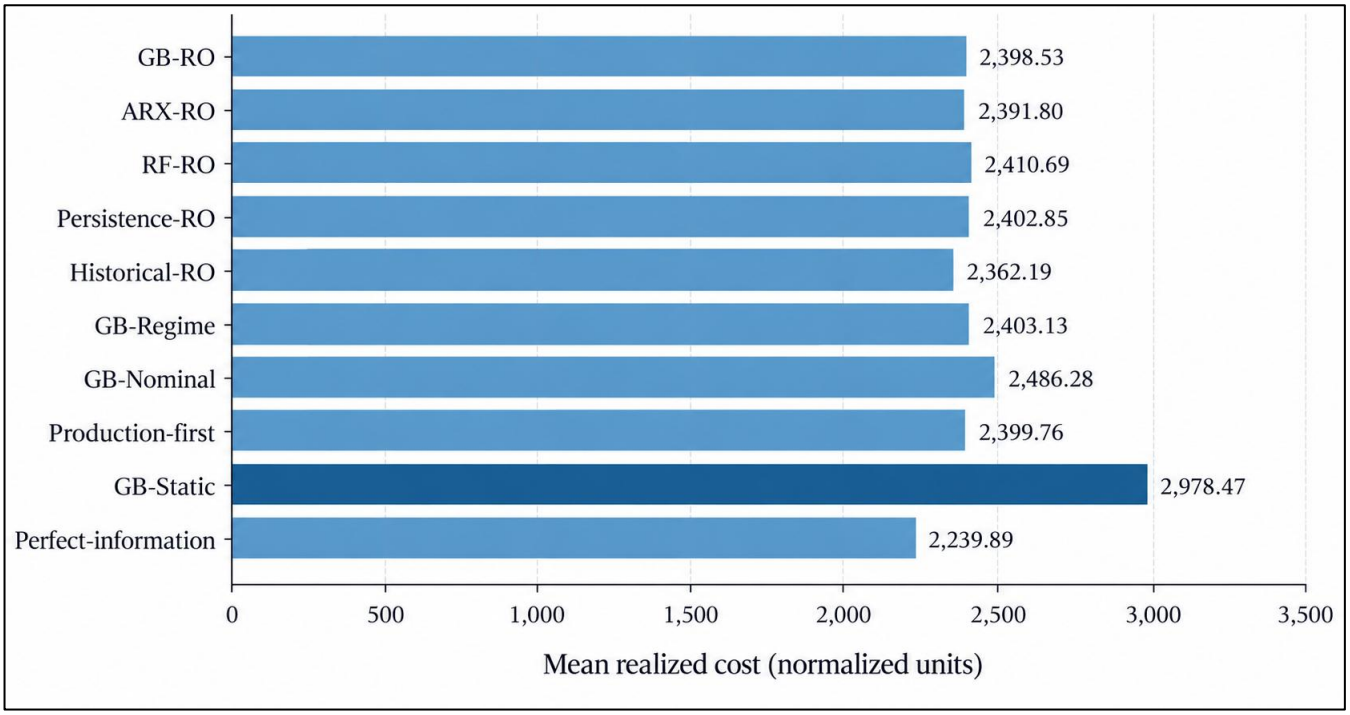


Fig. 1 Mean realized cost across the ten reported policy architectures

The figure does not identify a universally superior policy. Its perfect-information row is non-implementable, and service differences must accompany economic comparisons. All rolling policies have zero recorded material-forced shortfall relative to their current planned target, but they can still plan insufficient output and carry backlog. GB-Static instead averages 1.34 target-shortfall weeks per episode.

6.2. Matched capability contrasts

Table 9 uses GB-RO minus the comparator as the signed cost difference. Negative differences favor GB-RO. Intervals use the archived stratified bootstrap, not differences reconstructed from rounded policy means.

Table 9. Paired cost contrasts and pointwise 95% bootstrap intervals

Comparator	Cost difference	Pointwise 95% interval	GB cost saving (%)
GB-Nominal	-87.75	[-126.93, -50.56]	3.53
Persistence-RO	-4.33	[-22.93, 14.16]	0.18
ARX-RO	6.73	[-2.07, 15.78]	-0.28
RF-RO	-12.16	[-23.83, -0.95]	0.50
Historical-RO	36.33	[28.83, 43.85]	-1.54
Production-first	-1.23	[-2.68, 0.22]	0.05
GB-Static	-579.95	[-709.82, -454.87]	19.47

Rolling updates provide the clearest contrast: mean cost is 579.95 units lower than retention of the initial plan, corresponding to a 19.47% reduction, while on-time service is 13.86 percentage points higher. Both policies begin with the same model and residual construction. The comparison identifies the combined value of the specified update protocol, including newly observed operating information and repeated optimization. It does not separately identify inventory observation, forecast refreshing and production-plan revision.

Residual protection around the GB forecast reduces mean cost by 3.53% relative to GB-Nominal. On-time service is 2.98 percentage points higher, but end fill is 0.86 points lower. Thus, timely fulfilment and recovery by the endpoint measure different outcomes. By contrast, the cost advantage over Persistence-RO is only 4.33 units, with an interval spanning zero. GB-RO is 6.73 units more expensive than ARX-RO on average, again with an interval spanning zero. Historical-RO costs less, with a different service outcome.

The coordination contrast is also small: mean cost is 1.23 units lower than Production-first, or approximately 0.051%, with a pointwise interval of [-2.68, 0.22]. After adjustment for seven cost comparisons, intervals exclude zero only for nominal planning, static planning and Historical-RO; the historical comparison favors the simpler policy on cost. The narrow pointwise advantage over RF-RO does not survive the reported multiplicity adjustment.

Fig. 2 shows the same matched contrasts with their pointwise intervals. Negative values indicate lower cost for GB-RO.

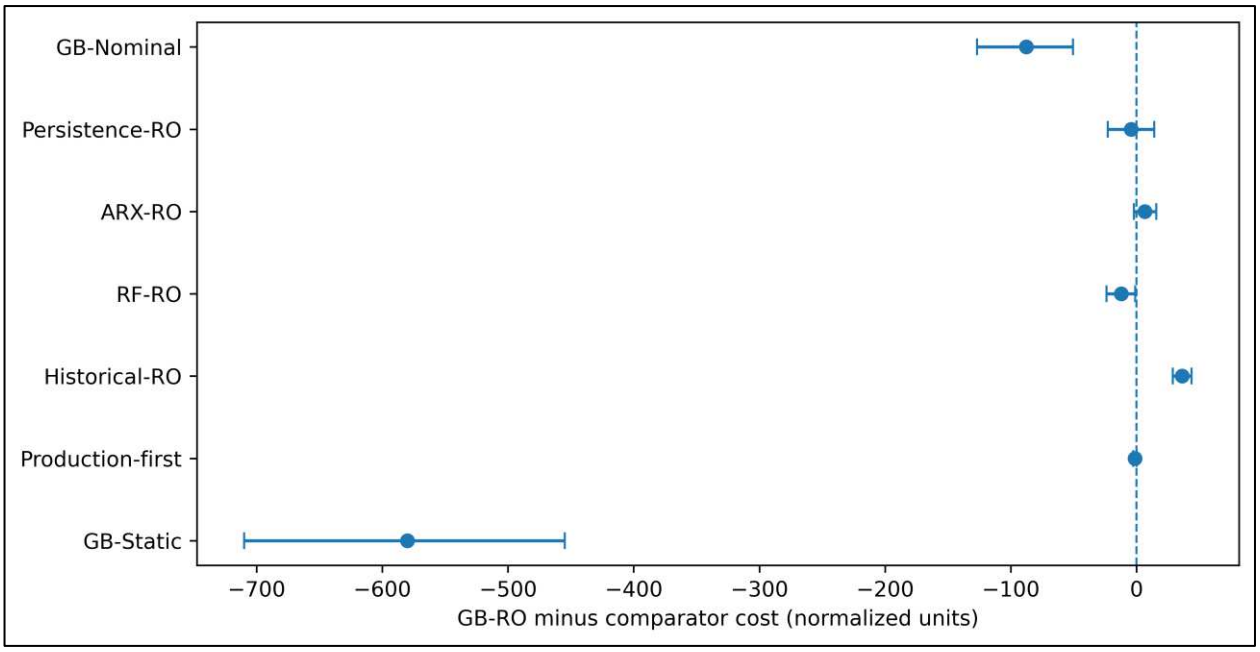


Fig. 2 Paired mean cost differences with pointwise 95% intervals

The figure supports a differentiated conclusion: state feedback and residual protection have demonstrated value within the tested architecture, while a general learning-related or coordination superiority claim is not supported. These effects must not be added as a causal decomposition of all possible planning improvements.

6.3. Operating conditions and configuration dependence

Across 216 factorial paths, the mean paired percentage advantage of integrated GB-RO over Production-first is approximately 0.079%. The 27 cell means range from -0.285% to 0.460%, with eight negative means. Only two pointwise cell intervals are wholly positive before adjustment for inspecting multiple cells. The largest cell mean occurs at 0.5 weeks of initial coverage, a capacity multiplier of 1.00 and severity 1.1, with interval [0.188, 0.756]%. This selected maximum is not a validated managerial threshold.

Fig. 3 includes all 27 cells, using the ordering reported in Appendix A. It preserves negative and near-zero effects rather than presenting only favorable conditions.

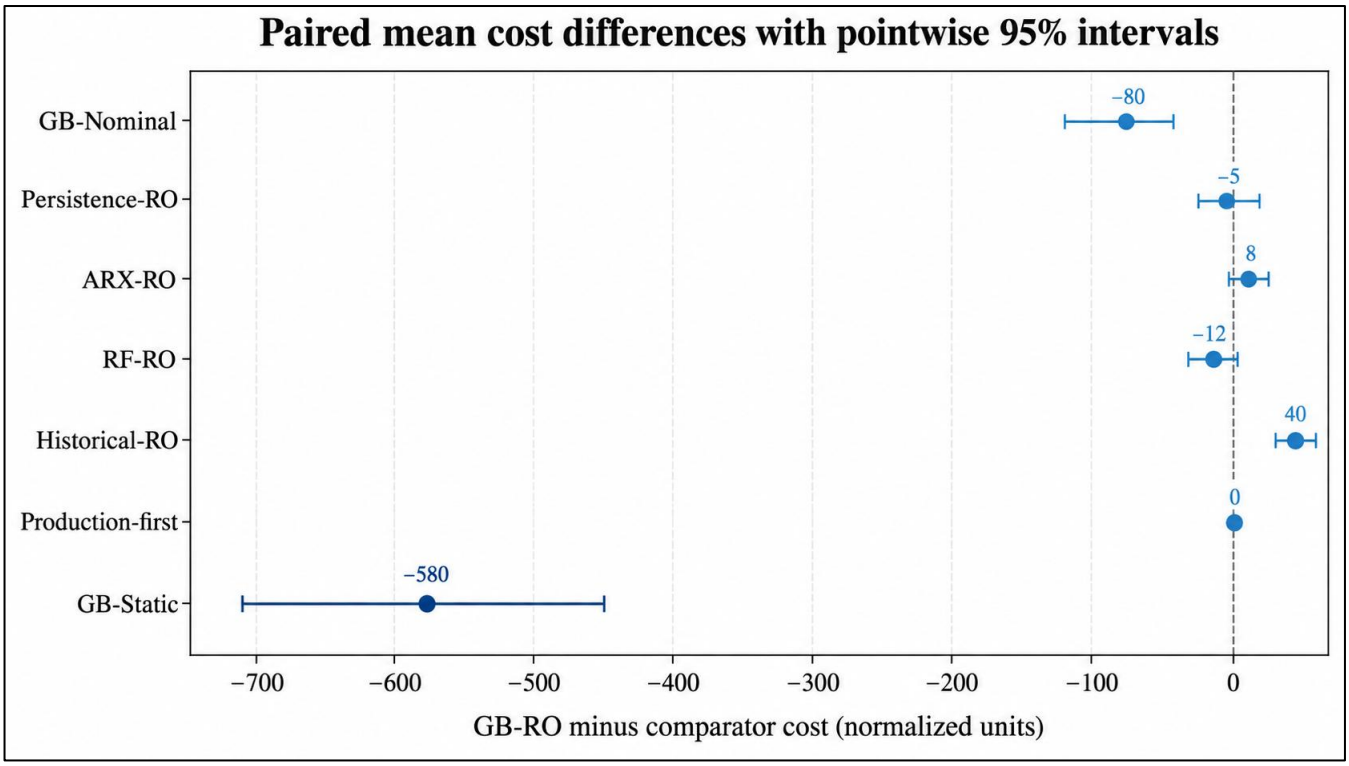


Fig. 3 All factorial cell means and pointwise intervals for the paired integration advantage

Both integrated and production-first policies can purchase protection and replan against accumulated backlog. The production-first target also respects factory capacity. Their similar responses under constant small-system demand make a small incremental coordination effect plausible, but this is a model-based interpretation, not identification of a unique causal mechanism.

The separate multi-size checks show greater differences in two particular configurations. In the medium configuration, GB-RO costs 13,142.04 versus 13,974.74 for Production-first, a 5.96% ratio-of-means difference. In the larger configuration, the corresponding means are 27,571.38 and 28,406.01, a 2.94% difference. Each has only twelve paths and one generated material structure. Structure, demand, horizon and parameters change alongside size, so these outcomes motivate broader testing rather than establish generalizability.

Table 10 distinguishes solver time from matrix-building-plus-solve time for five initial solves per size. Forecast fitting and complete rolling-policy execution are excluded from these timing figures.

Table 10. Initial-solve computational scale; K/M/R/T denotes products/materials/routes/weeks

Size (K/M/R/T)	Variables / constraints	Mean solver (s)	Maximum solver (s)	Mean build + solve (s)
2/2/3/10	262 / 250	0.0410	0.0520	0.0444
5/4/4/12	676 / 456	0.1277	0.1750	0.1358
8/6/5/14	1,266 / 714	0.1357	0.2786	0.1528

The reported maximum relative gap is zero for these fifteen initial solves. These modest models establish tractability only for the tested configurations; they do not demonstrate industrial-scale performance across rich integer scheduling structures.

6.4. Resource commitment and punctual-output diagnostics

A resource-oriented interpretation should not equate a purchased buffer with material efficiency. For the main benchmark, every episode has total demand of 200 aggregate output units. The reported mean on-time percentage therefore gives a mean timely quantity of $200\overline{S}/100$. Dividing mean cost by this quantity produces a descriptive ratio of means, not a mean of episode-specific unit costs. Buffer use can likewise be expressed as a percentage of the common 120-unit optional-stock cap.

Table 11 reports calculations from the rounded policy means. They are economic and material-commitment diagnostics, not new policy experiments or environmental indicators. Aggregate product and material units are appropriate only within this specified benchmark.

Table 11. Additional descriptive calculations from the reported policy means

Policy	Mean timely units	Cost / mean timely unit	Buffer-cap use (%)
GB-RO	193.88	12.37	99.06
ARX-RO	194.48	12.30	96.37
RF-RO	194.40	12.40	92.09
Persistence-RO	194.32	12.37	83.89
Historical-RO	193.02	12.24	100.00
GB-Regime	194.06	12.38	98.58
GB-Nominal	187.90	13.23	71.92
Production-first	193.74	12.39	99.06

GB-RO uses approximately 99.06% of the optional-stock cap, compared with 71.92% for GB-Nominal. Its lower cost and higher punctual service therefore do not establish lower material commitment. Its cost per mean timely unit is approximately 12.37, versus 13.23 for GB-Nominal and 17.93 for GB-Static. Historical-RO gives approximately 12.24 and ARX-RO 12.30. None of these ratios measures energy intensity, embodied emissions or waste prevention.

The supplementary post-hoc service-threshold screen in Appendix A.5 illustrates how the lowest-cost eligible mean changes with the service criterion. It is descriptive and is not used for validation selection or a deployment recommendation.

Fig. 4 displays mean cost and on-time service for the implementable rolling policies. The points are descriptive, not a proved efficient frontier.

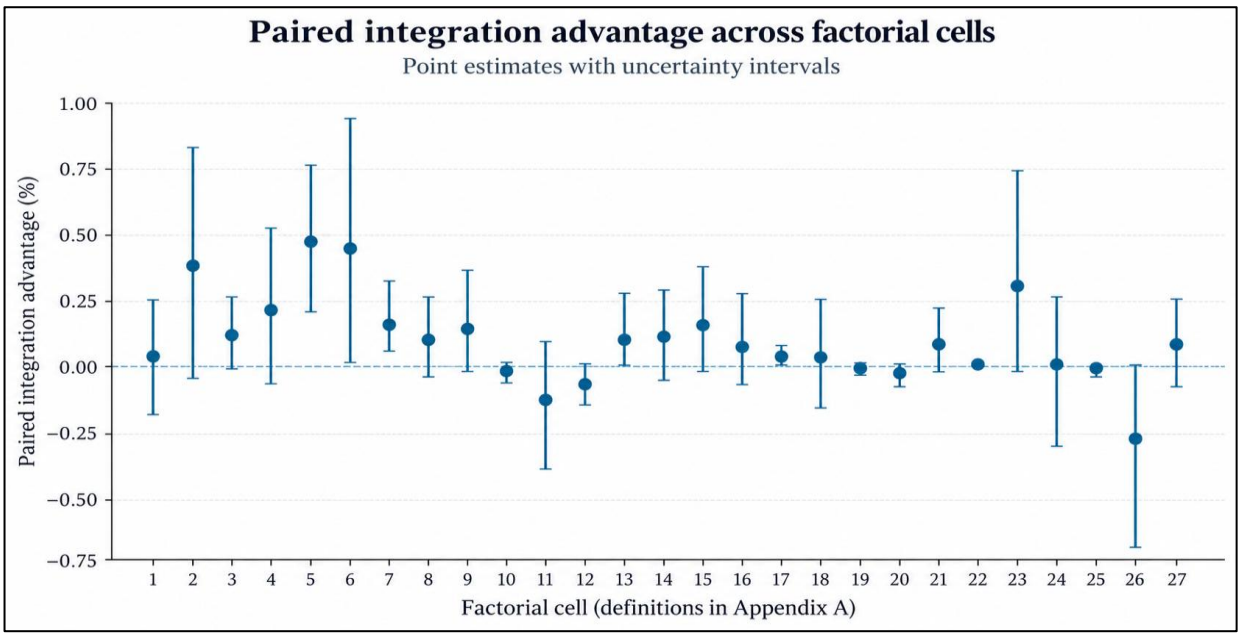


Fig. 4 Mean cost and on-time service for the implementable rolling policies

Cost and service need to be read jointly. A policy may spend less by accepting lateness, or spend more on protective stock without increasing punctual output materially. Supplementary Information S1 specifies how common-service comparison can be combined with explicit environmental accounting.

6.5. Tail outcomes and uncertainty-set diagnostics

The empirical upper-decile mean cost is 2,587.78 for GB-RO, 3,236.77 for GB-Nominal and 4,944.80 for GB-Static. These descriptive tail results are consistent with protection and feedback reducing expensive episodes in the generator. They are neither a conditional-value-at-risk optimization result nor a certified tail-risk bound.

Warning-conditioned GB-Regime costs 2,403.13 versus 2,398.53 for fixed GB-RO. Its mean end-fill increment is 0.096 percentage points and its on-time increment is 0.091 points. The reported cost and service intervals include zero. A demonstrated gain from warning-conditioned prototypes is therefore not established.

The source's current-vector envelope exceedance rates are 14.58% for GB-RO, 13.83% for ARX-RO, 36.75% for Persistence-RO and 31.92% for Historical-RO. This diagnostic checks whether the realized current route vector exceeds the componentwise maximum of the planning trajectories. It is not joint trajectory coverage and does not prove feasibility for every vector below that envelope. Different exceedance rates further demonstrate why a common multiplier grid does not create equal risk protection.

7. Empirical factory diagnostics and the limits of resource proxies

7.1. Record reconstruction and service definitions

The source case analysis joins dispatches to orders by order_line_id and usable receipts to shipments by shipment_line_id. It carries opening stock through receipts, production, bill-of-material consumption and customer dispatch. It excludes 52 ETA rows containing only a source label, leaving 101 identified ETA events. Production is assigned the documented Friday 10:00 aggregation time because individual completion timestamps are not supplied. The reconstructed record counts and the full weekly summaries are retained in Appendix B.

Eligible demand consists of order lines with original due times inside the reporting window. Net quantity deducts cancellations known before the due date. On-time dispatch is capped at each line's net quantity and credits only dispatches at or before its original due time. The two opening overdue lines remain in backlog and physical reconciliation but are excluded from the new-demand service denominator. Fulfilment by the extraction cutoff includes late dispatches and is reported separately.

The source reconstruction reconciles 54 order balances, 135 inventory snapshots and 67 shipments, with a maximum inventory residual of 1.71×10^{-13} units. Recorded production respects regular capacity in all 26 weeks, and bookings respect recorded allowances in all 52 route-weeks. These are linkage and arithmetic checks. They do not independently authenticate originating records or estimate a counterfactual planning effect.

7.2. Punctual dispatch versus eventual fulfilment

Table 12 separates the case's customer-service measures, inbound availability and production-resource use. All are descriptive outcomes of the recorded baseline.

Table 12. Reported outcomes of the author-confirmed 26-week factory ledger

Measure	Value	Interpretation
Eligible in-window order lines	52	Two opening overdue lines excluded from this cohort
Net pieces due	5,064	Original due times retained; cancellations deducted
On-time dispatch	1,158; 22.87%	Punctual customer service
Late dispatch by cutoff	3,670 pieces	Fulfilled late, not credited on time
Fulfilled by cutoff	4,828; 95.34%	Timely and late fulfilment combined
Open quantity at cutoff	236 pieces	Retained in final backlog
Inbound usable release	13,480 of 14,120 kg	Another 640 kg remains in the committed pipeline

Measure	Value	Interpretation
Regular-hour use	2,744.05 of 2,844 hours	96.49% aggregate utilization

The gap between final fulfilment and punctual service is 72.47 percentage points. Product P01 has 600 on-time pieces out of 2,840 due, or 21.13%; P02 has 558 out of 2,224, or 25.09%. Recorded good output is 2,671 and 2,167 pieces, respectively. Opening finished stock and service of opening overdue orders explain why production totals do not equal fulfilment of the in-window cohort.

Fig. 5 retains all 26 weekly outcomes and the stated 95% service target. The route restriction is recorded in weeks 9–13, but the period boundaries are not randomized treatments and do not identify a causal disruption effect.

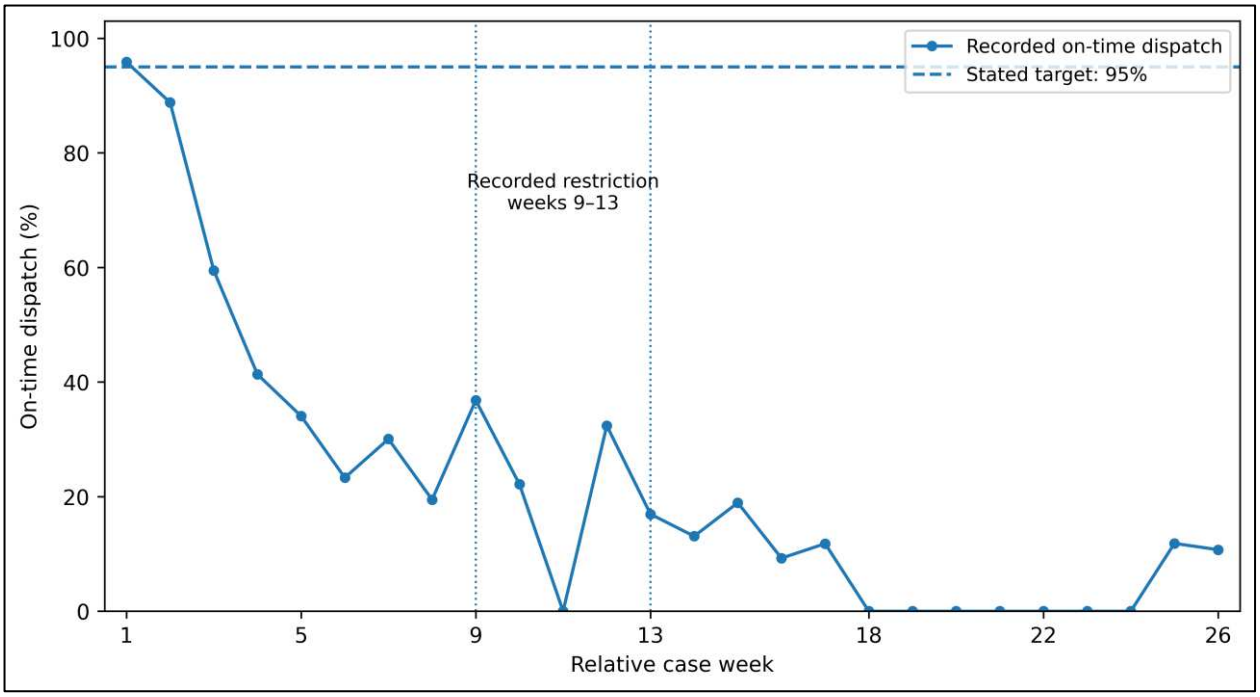


Fig. 5 Quantity-weighted weekly on-time dispatch and the stated service target

The continuing low service after the restriction period illustrates why restoring inbound availability need not restore punctuality when backlog and resource constraints carry forward. The case does not show that implementing GB-RO would reproduce its computational cost saving.

7.3. Recovery headroom and economic accounting

Table 13 groups the continuing case into the source's chronological segments. Each segment inherits the previous state; the groups are not independent replications, fresh holdouts or separately executed policy experiments.

Table 13. Chronological service and realized regular-capacity headroom

Weeks	Net due	On time	Service (%)	Hours used / available	Unused hours
1–14	2,744	1,035	37.72	1,432.90 / 1,532	99.10
15–20	1,281	84	6.56	651.60 / 652	0.40
21–26	1,039	39	3.75	659.55 / 660	0.45

Additional calculations from the full weekly table give 99.95 unused regular hours over the entire case, but only 0.85 unused hours in the last twelve weeks. Eight weeks have zero on-time dispatch. Closing overdue demand reaches 364 pieces in week 20 and ends at 236, compared with an opening overdue quantity of 35. These values describe the recorded schedules; unused time is not an estimate of optimized recovery capacity.

Within-window accruals total SAR 538,932.05: material SAR 244,960.00; production SAR 219,524.00; freight SAR 32,860.00; backlog SAR 35,180.00; holding SAR 3,268.05; and reserve SAR 3,140.00. The source analysis checks charges against supplied rates and excludes pre-window shipment charges. The total is not independently verified cash expenditure, and the normalized computational saving is not applied to it.

7.4. What the records support for sustainable manufacturing

The ledger supports an operational diagnosis: materials, backlog age and conversion headroom should be reviewed together. High utilization is compatible with poor punctual service, and additional material may not produce timely output when recovery capacity is exhausted. These findings motivate measuring the resource consequences of protection rather than treating recovery volume as a benefit automatically.

The case does not contain the measurements needed to estimate machine energy profiles, route-specific emissions, processing scrap, product end-of-life returns, recovery yields or secondary-material qualification. Regular hours cannot be converted into electricity using an undocumented factor, and stock balances cannot establish material waste without quality and disposal records. The environmental and circularity extension therefore remains analytically specified rather than retrospectively imposed on this ledger.

8. Analytical extension for timing-qualified circular manufacturing

8.1. Decision boundary and modeling assumptions

The extension connects secondary-material recovery to the same production-use deadlines as inbound protection. It distinguishes usable virgin material, qualified secondary material, recoverable feedstock and recovery work in progress. Its purpose is a mathematically specified decision structure; the reported policy experiments in Sections 5–6 evaluate the operational core, not this extension. The broader evaluation is specified separately in Supplementary Information S1.

All material quantities in this section use a common mass unit, such as kilograms, with material coefficients and transport limits converted consistently. Periods retain the local index $h = 0, \dots, H-1$ and absolute time $t + h$. Bookings, production, material allocations, recovery, disposal and customer dispatch are common across scenarios within a solve; material inventories reflect scenario-dependent arrivals. Thus, the extension retains open-loop finite-scenario planning with repeated state feedback, rather than introducing scenario-contingent decisions with perfect foresight. Recovery yields and release delays are specified coefficients within each solve. Varying them across separate sensitivity settings is not an estimate of their probability distributions.

8.2. Material qualification, recovery and carried work in progress

Let u^V_{hmk} and u^S_{hmk} be virgin and secondary mass allocated to product k, and let $\theta_{mk} \in [0,1]$ be the permitted secondary share. With a_{mk} denoting gross material input per finished unit, allocation and qualification satisfy

$$u^V_{hmk} + u^S_{hmk} = a_{mk} q_{hk}, \qquad 0 \le u^S_{hmk} \le \theta_{mk} a_{mk} q_{hk}. \qquad (13)$$

Set $\theta_{mk} = 0$ for ineligible material–product combinations. Let I^V_{shm}, I^S_{shm} and W_{shm} denote closing virgin, qualified-secondary and unprocessed recoverable stocks. External usable secondary receipts and recoverable feed receipts are exogenous quantities A^S_{shm} and A^W_{shm}. They must originate from identified external flows, not duplicate internally generated scrap. Virgin balance retains the original booking and reserve structure:

$$I^V_{shm} = I^V_{s,h-1,m} + A_{hm} + \mathbf{1}_{\{h=0\}} b_m + \sum_{u,r:\, u+\ell_{sur}=h} x_{umr} - \sum_k u^V_{hmk}. \qquad (14)$$

Recovery input r_{hm} consumes available feedstock and processing capacity in period h. Its qualified fraction $0 < \eta_m \le 1$ is released after $L^R_m \ge 1$ integer periods. All active processing is charged in the initiation period; the subsequent delay represents completion/qualification waiting without additional modeled machine occupancy. Multiperiod processing would require an occupancy profile instead. Let P^0_{jm} be qualified mass from recovery initiated before the current solve and scheduled for release in local period j. This opening pipeline includes every future release, including those beyond the current optimization window. The secondary-material balance is

$$I^S_{shm} = I^S_{s,h-1,m} + A^S_{shm} + P^0_{hm} + \sum_{u:\, 0 \le u < H,\ u+L^R_m = h} \eta_m r_{um} - \sum_k u^S_{hmk}. \qquad (15)$$

For each product–material pair, σ_{mk} is the recoverable scrap fraction and τ_{mk} the nonrecoverable process-loss fraction, with $0 \le \sigma_{mk} + \tau_{mk} \le 1$. Material retained in the finished output is the remaining fraction of gross input. Let w_{hm} be disposal directly from the recoverable pool. Its balance and input-availability constraint are

$$W_{shm} = W_{s,h-1,m} + A^W_{shm} + \sum_k \sigma_{mk} a_{mk} q_{hk} - r_{hm} - w_{hm},$$
$$r_{hm} + w_{hm} \le W_{s,h-1,m} + A^W_{shm}, \qquad \forall s, h, m. \qquad (16)$$

The second line prevents new production scrap from being recovered or disposed of before the end of its generation period. Nonrecoverable production loss, rejected recovery output and direct disposal are recorded once in the waste stream

$$d_{hm}^{W} = w_{hm} + \sum_{k} \tau_{mk}\, a_{mk} q_{hk} + (1 - \eta_m) r_{hm}. \qquad (17)$$

Recovery rejects are assumed to enter waste treatment in the initiation period. A facility that holds rejects before treatment needs a separate reject-stock balance. Gross-input coefficients must be reconciled with retained product mass, recoverable scrap and nonrecoverable loss; introducing scrap must not create material in addition to the original mass input.

Let Q_{hm} be qualified recovery work in progress awaiting usable release. The opening stocks in (14)–(16) are observed usable quantities, with P_{0m}^{0} excluded until it is released in period zero. Work in progress is tracked by

$$\begin{aligned} Q_{-1,m} &= \sum_{j \geq 0} P_{jm}^{0}, \\ Q_{hm} &= Q_{h-1,m} + \eta_m r_{hm} - P_{hm}^{0} - \sum_{u:\, 0 \leq u < H,\; u + L_m^R = h} \eta_m\, r_{um}. \qquad (18) \end{aligned}$$

All stocks, allocations, processing and disposal quantities are nonnegative. Pending recovery is therefore distinct from usable secondary stock, and its terminal quantity remains visible. After actual period-zero execution, the pipeline is shifted, not re-created from a new plan:

$$P_{jm}^{0,t+1} = P_{j+1,m}^{0,t} + \mathbf{1}_{\{j+1 = L_m^R\}} \eta_m r_{0m}^{\mathrm{act},t}, \qquad j \geq 0. \qquad (19)$$

The update records only recovery actually initiated. Planned future recoveries do not become commitments. Together, (15), (18) and (19) prevent duplicate credit for opening work in progress, loss of releases beyond the current window and premature reuse of recovered material.

8.3. Shared capacity and the conditional value of recovery

Let f_m^R be processing hours per recovery-feed unit and let C_h follow the production calendar. When recovery and production share the constrained resource, replace the original production-capacity constraint with

$$\sum_{k} f_k\, q_{hk} + \sum_{m} f_m^R\, r_{hm} \leq C_h. \qquad (20)$$

Capacity sharing in manufacturing–remanufacturing systems has established precedents [22, 23]. Here it is connected to qualified material availability before an original customer deadline. Additional recovery can remove a material shortage while consuming the hours needed to convert that material into finished output.

Consider a single-window diagnostic in which all recovery counted becomes usable before the production deadline, secondary eligibility is nonbinding and the within-window timing is feasible. With available material A, demand D, hours C, production time $f^P > 0$, recovery time $f^R > 0$, requirement $a > 0$, yield $\eta > 0$ and feedstock U, feasible output is

$$q(r) = \min\left\{D, \frac{C - f^R r}{f^P}, \frac{A + \eta r}{a}\right\}, \quad 0 \leq r \leq \min\{U, C/f^R\}. \qquad (21)$$

When material initially binds and demand does not bind at the crossing, the material and capacity branches intersect at

$$r^{\mathrm{cross}} = \frac{aC - f^P A}{a f^R + f^P \eta}. \qquad (22)$$

Below the crossing, the material-limited branch increases at η/a; above it, the capacity-limited branch decreases at f^R/f^P. A binding demand cap can create a plateau, while an infeasible crossing requires consideration of the relevant endpoint. With additional timing or qualification restrictions, (21) is only an upper bound until those restrictions are enforced. It is not a general economic recovery policy.

Table 14 illustrates the mechanism using explicit assumptions: $D = 100$, $C = 80$ hours, $f^P = 1$ hour/unit, $f^R = 0.2$ hour/feed unit, $A = 60$, $a = 1$, $\eta = 0.8$ and $U = 50$. The values are analytical assumptions, not estimates from the factory ledger.

Table 14. Analytical illustration of recovery competing with finished production

Recovery feed	Qualified recovery output	Recovery hours	Feasible finished output	Unused usable material
0	0	0	60	0
10	8	2	68	0
20	16	4	76	0
40	32	8	72	20
50	40	10	70	30

Processing 20 feed units permits 76 finished units; processing all 50 permits only 70 because of the shared-capacity burden. Qualified recovery output, material consumed and terminal unused material must therefore be distinguished. The example establishes a production mechanism, not a measured environmental improvement.

8.4. Physically supported dispatch and original due-cohort service

A service floor must be linked to actual finished-product availability. Let $\mathcal{O}_t$ contain known order lines, including carried overdue orders and completed lines retained for service accounting. Each order has product $k(i)$, fixed net quantity D_i, earliest permitted dispatch time ρ_i and original deadline δ_i, with $\rho_i \le \delta_i$. Order knowledge and dispatch eligibility are different: an order can be known before ρ_i. Unrevealed future orders are not supplied to a planning run. Net quantities incorporate only the stated eligible cancellations; due dates and historical service credits are not reset during replanning.

Let c_i^t and $c_i^{OT,t}$ be cumulative actual total and on-time dispatch before the solve, with $0 \le c_i^{OT,t} \le c_i^t \le D_i$. Let o_{hi} be planned dispatch and R_{hi} remaining order quantity. Then

$$R_{-1,i} = D_i - c_i^t, \qquad R_{hi} = R_{h-1,i} - o_{hi} \ge 0, \qquad o_{hi} \ge 0. \qquad (23)$$

Dispatch is prohibited before eligibility, and every unit dispatched must be available as the correct finished product. Finished stock and due backlog satisfy

$$\begin{aligned} o_{hi} &= 0 \quad \text{if } t + h < \rho_i, \\ F_{hk} &= F_{h-1,k} + q_{hk} - \sum_{i:\, k(i)=k} o_{hi} \ge 0, \\ B_{hk} &= \sum_{i:\, k(i)=k,\ \delta_i \le t+h} R_{hi}, \qquad F_{-1,k} = F_k^0. \end{aligned} \qquad (24)$$

These equations replace the aggregate finished-inventory/backlog balance in the extension. They distinguish unfulfilled future orders from overdue demand and prevent a service claim unsupported by production or opening finished stock. Setting $\rho_i = \delta_i$ disallows early shipment where that is the relevant operating convention.

To retain the baseline's oldest-due-order-first rule, define $i \prec j$ as earlier priority for the same product, breaking equal deadlines by a fixed order identifier. For eligible higher-priority orders, the following binary constraints are sufficient:

$$\begin{aligned} o_{hj} &\le R_{-1,j}\, \zeta_{hj}, \qquad \zeta_{hj} \in \{0,1\}, \\ R_{hi} &\le R_{-1,i}\left(1 - \zeta_{hj}\right), \quad i \prec j,\ \rho_i \le t + h. \end{aligned} \qquad (25)$$

A positive dispatch to a lower-priority order requires earlier eligible orders to be cleared in that period. Omitting (25) permits a different allocation policy and must be applied consistently to all compared methods, rather than changing the dispatch rule only for the proposed policy. The added binaries are not included in the core model-size or timing results.

Let $\mathcal{C}$ be the fixed evaluation cohort with original deadlines inside the declared reporting window, all covered by the current fixed terminal date. Retain completed lines in its denominator and exclude opening overdue lines outside that window. For $\sum_{i \in \mathcal{C}} D_i > 0$, planned on-time service is

$$S^{OT,t} = \frac{\sum_{i \in \mathcal{C}} \left(c_i^{OT,t} + \sum_{h:\, \rho_i \le t+h \le \delta_i} o_{hi}\right)}{\sum_{i \in \mathcal{C}} D_i}. \qquad (26)$$

An expired deadline contributes only its previously earned on-time credit. With common production and dispatch decisions, this planning ratio does not require a scenario subscript. A fixed-cohort guarantee applies only to orders known to the solve; prospective evaluation with later order revelation must retain the declared cohort-construction rule and report violations rather than provide future order information to the policy. If the cohort is empty, service is reported as not applicable and its floor is omitted.

At the next planning date, actual—not planned—dispatch updates both credits:

$$c_i^{t+1} = c_i^t + o_{0i}^{\mathrm{act},t},$$
$$c_i^{OT,t+1} = c_i^{OT,t} + \mathbf{1}_{\{\rho_i \le t \le \delta_i\}} o_{0i}^{\mathrm{act},t}. \qquad (27)$$

Opening finished stock, remaining orders and due backlog are updated from the same executed transactions. Together with (19), this preserves the physical and service state across successive planning runs.

8.5. Explicit economic, energy and emissions accounting

Let E_h^{base} and E_h^{IT} be documented background and allocated computing electricity in kWh per period. Let ϵ_k^P be incremental kWh per finished unit, ϵ_m^R kWh per recovery-feed unit and $\epsilon_m^V, \epsilon_m^S, \epsilon_m^W, \epsilon_m^Q$ storage electricity per kg-period for the four material states. Period and remaining-horizon electricity are

$$E_{sh} = E_h^{base} + E_h^{IT} + \sum_k \epsilon_k^P q_{hk} + \sum_m \epsilon_m^R r_{hm} + \sum_m \left(\epsilon_m^V I_{shm}^V + \epsilon_m^S I_{shm}^S + \epsilon_m^W W_{shm} + \epsilon_m^Q Q_{hm}\right),$$
$$E_s^{rem,t} = \sum_{h=0}^{H-1} E_{sh}. \qquad (28)$$

Loads must be disjoint: an inclusive machine-energy coefficient cannot be added to the same background electricity again. Storage terms require an appropriate physical or allocated basis. Computing electricity includes a declared allocation of training, forecasting and optimization; it is not assumed to be zero. An unchanged background load is not a saving when output decreases.

Separate virgin acquisition from virgin consumption. Remaining-horizon acquisitions $P^{V,t}$, virgin input consumed $V^{rem,t}$, secondary input consumed $U^{S,t}$ and waste treated $D^{W,rem,t}$ are

$$P^{V,t} = \sum_{h,m,r} x_{hmr} + \sum_m b_m, V^{rem,t} = \sum_{h,m,k} u_{hmk}^V,$$
$$U^{S,t} = \sum_{h,m,k} u_{hmk}^S, D^{W,rem,t} = \sum_{h,m} d_{hm}^W. \qquad (29)$$

Pre-existing virgin stocks and committed shipments are already-acquired resources and are not purchased again in $P^{V,t}$. Their acquisition burdens remain in the declared opening/past ledger. Secondary input consumed is not automatically secondary content in punctual output: such a claim additionally requires product/batch tracing and treatment of process losses.

For costs, let $h_m^V, h_m^S, h_m^W, h_m^Q$ be inventory charges per kg-period. Define the scenario material-holding charge

$$H_s^{mat} = \sum_{h,m} \left(h_m^V I_{shm}^V + h_m^S I_{shm}^S + h_m^W W_{shm} + h_m^Q Q_{hm}\right). \qquad (30)$$

Retain J_{book} from (6) for new virgin bookings, activation, reserve and initial optional protection. Let $c_k^{P,ne}$ and $c_m^{R,ne}$ be production and recovery costs excluding electricity priced separately; c_m^W is treatment cost per kg of waste; p_h^E is cost per kWh. Let $J^{fix,t}$ itemize fixed remaining costs of precontracted secondary/feedstock flows and other commitments, and $c^{IT,ne,t}$ non-electric digital-system costs. A complete remaining-cost objective is

$$J^{circ,t} = J_{\mathrm{book}} + J^{fix,t} + c^{IT,ne,t} + \sum_{h,k} \left(c_k^{P,ne} q_{hk} + c_k^F F_{hk} + \pi_k B_{hk} + \kappa \Delta_{hk}\right) + \sum_{h,m} \left(c_m^{R,ne} r_{hm} + c_m^W d_{hm}^W\right) + \frac{1}{S} \sum_s \left(H_s^{mat} + \sum_h p_h^E E_{sh}\right). \qquad (31)$$

The scenario average is a modeling convention, not a calibrated probability expectation. Costs already incurred are sunk for the current optimization but retained in episode totals. When an existing production or recovery rate already includes energy, either remove that component before using (31) or omit the corresponding additional energy charge. Recomputed remaining-plan objectives are not summed as realized episode costs.

Define acquisition factors g_m^V in kg CO2e/kg virgin material; freight factors g_{hmr}^X in kg CO2e/kg shipped on the stated route and distance; protective-stock positioning factors g_m^b in kg CO2e/kg; electricity factors g_h^E in kg CO2e/kWh; non-electric production factors $g_k^{P,ne}$ per finished unit; non-electric recovery factors $g_m^{R,ne}$ per kg feed; and treatment factors g_m^W per kg waste. For a declared short-horizon acquisition-and-operations boundary,

$$G_s^{rem,t} = G_s^{fix,t} + \sum_{h,m,r} (g_m^V + g_{hmr}^X)\, x_{hmr} + \sum_m (g_m^V + g_m^b)\, b_m + \sum_h g_h^E\, E_{sh} + \sum_{h,k} g_k^{P,ne}\, q_{hk} + \sum_{h,m} \left(g_m^{R,ne} r_{hm} + g_m^W d_{hm}^W\right). \qquad (32)$$

Here $G_s^{fix,t}$ is an itemized fixed-event ledger for precontracted external secondary/feedstock acquisition and transport, and other fixed remaining burdens not yet attributed to the past ledger. It is not a fitted residual. Secondary materials receive the burden specified by the chosen allocation convention; they are not assumed burden-free. Acquisition and freight are attributed once at the defined purchase/dispatch event, including material still unused at the endpoint. A reserve slot alone creates no automatic freight saving or additional carrier emission. Its contractual fee and any documented physical carrier consequences are different quantities.

These factors require documented units, geography, reference period, allocation convention and uncertainty. The accounting boundary follows lifecycle-inventory discipline [27] but does not include all product-use and end-of-life stages. Consequently, (32) is a bounded decision inventory, not a complete lifecycle assessment. For the narrow case in which ηr qualified units displace an equal virgin mass at unchanged output and service, recovery is beneficial on the affected inventory terms only when its full incremental burden is below the displaced virgin burden. Transport, yield, processing and terminal-stock differences must also enter that comparison.

8.6. Service-constrained environmental budgets and execution boundary

Use an epsilon-constraint formulation rather than combining incomparable physical quantities with unexplained weights. For a fixed reporting horizon, let $G^{past,t}$, $E^{past,t}$, $V^{past,t}$ and $D^{W,past,t}$ be burdens and resource use already attributed to that horizon. Environmental ceilings remain fixed across replanning:

$$\begin{aligned} &\min \quad J^{circ,t} \\ &\text{subject to} \quad G^{past,t} + G_s^{rem,t} \le \overline{G}, \qquad \forall s, \\ &E^{past,t} + E_s^{rem,t} \le \overline{E}^{energy}, \qquad \forall s, \\ &V^{past,t} + V^{rem,t} \le \overline{V}, \\ &D^{W,past,t} + D^{W,rem,t} \le \overline{W}, \qquad S^{OT,t} \ge \overline{S}. \end{aligned} \qquad (33)$$

The within-solve model comprises route constraints (4)–(5), initial-protection bounds, plan-deviation constraints (8), the physical balances and capacity limits (13)–(18) and (20), order constraints (23)–(26), and accounting definitions (28)–(32). Stock, capacity and dispatch balances replace their operational-core counterparts. Equations (19) and (27) update the state between solves; the single-window diagnostics (21)–(22) are not constraints in the multi-product optimizer. The ceilings are respectively kg CO2e, kWh, kg virgin input and kg waste; $\overline{S}$ is a fraction. A separate ceiling on acquisition $P^{V,t}$ can be imposed where procurement rather than consumption is the intended resource constraint. Positive holding and treatment costs do not substitute for these physical definitions.

Each realized period adds only newly attributable events to the past ledger. For any recorded quantity $X \in \{G, E, V, D^W\}$,

$$X^{past,t+1} = X^{past,t} + X_0^{act,t}. \qquad (34)$$

The opening allocation and event boundary must be declared consistently, including pre-positioned protection and opening recovery work in progress. Already-counted commitments cannot re-enter $X_0^{act,t}$. Fixed background burdens may affect feasibility even when they do not change the preferred decision. Infeasible environmental ceilings or an irrecoverable service deficit must be reported, not repaired by resetting past burdens, removing overdue orders or silently relaxing the target.

For given coefficients the extension is a mixed-integer linear formulation. Environmental-budget grids can be evaluated with explicit feasibility and nondominance checks. Disabling secondary/recovery flows, process losses, environmental ceilings and separate energy charges gives the corresponding operational specialization under matching arrival, dispatch-priority and cost conventions. This formulation-level relationship is not a claim of an already tested extended solver.

Only current feasible actions can be executed; actual material releases, recoveries and dispatches determine the next state and outcome ledger. Out-of-sample supplies can invalidate a planned service level even when the finite planning scenarios were feasible. Implementation therefore requires conservation, eligibility, priority, commitment and accounting tests,

followed by common-service policy evaluation. The formulation specifies what must be evaluated without adding unexecuted environmental or circularity outcomes to the reported evidence.

9. Discussion, implementation and limitations

9.1. Implications for sustainable smart manufacturing

The central operating mechanism is useful conversion before a deadline. A reservation can increase nominal transport availability while arriving too late; a timely component can remain unproductive when a complementary input is absent; a recovered material can be qualified yet consume scarce capacity needed for finished output. These mechanisms explain why resilience interventions and circularity targets should be evaluated at the production-system level rather than through isolated volume measures.

The computational evidence supports investment in maintaining and updating the physical state before assuming that a more elaborate learning method is necessary. State feedback has the largest demonstrated effect among the matched contrasts. Residual protection improves cost relative to nominal planning, but increases initial protective-stock use and changes service outcomes. Comparisons with persistence, ARX and historical means make the cost of additional forecasting complexity a genuine decision question, not a predetermined conclusion. Reviews of smart supply-chain optimization and resilience likewise provide a broad context for adaptive decision support without proving that every sophisticated architecture is preferable [33, 34].

Circular material availability must be qualified, timely and compatible with production capacity. The worked recovery example shows that processing more feedstock can leave additional secondary material unused while reducing finished output. Equations (13)–(34) connect that mechanism to carried recovery commitments, actual finished-stock availability, original due-cohort service and environmental budgets that cannot be reset at each planning run. The methodological implication is to evaluate useful material conversion, not recovery volume alone.

9.2. Managerial workflow and implementation responsibilities

A planner should first identify the earliest material-deficient prefix using usable stock and confirmed arrivals. Each alternative booking or recovery action can then be checked against that prefix, the other required materials and the available production hours. The factory ledger shows why backlog age and punctual service must accompany gross receipts and utilization. A dashboard displaying 95.34% final fulfilment alone would obscure the recorded 22.87% on-time outcome.

Implementation can begin with structured exports rather than a full digital twin. Production planning should reconcile usable stock, allocations and backlog; manufacturing engineering should maintain bills of materials, recovery qualifications and capacity calendars; logistics should preserve booking commitments and timestamped ETA versions; sustainability specialists should document physical boundaries and environmental factors; and analytics staff should retain simple predictors while monitoring drift and validation-service performance.

A shadow-operation stage can compare recommendations with actual planner decisions without issuing bookings automatically. Human-approved execution should follow only after discrepancies in material availability, commitment terms and accounting are understood. Neither the time required for deployment nor the resulting factory saving is estimated here. Data sharing between functions supports an implementation pathway, but no labor, health, partnership or other social-sustainability effect is inferred from the present results.

9.3. Limits of inference

The operational comparisons use computer-generated histories and normalized costs. The author-confirmed empirical ledger provides separate descriptive evidence after internal reconciliation, with its provenance and time transformation documented in Appendix B. It does not authenticate the originating system, calibrate the computational generator or identify a field-implemented policy saving. Summary-level ratios retain the information limits of the reported tables.

The main system is small and uses constant known demand. The larger checks contain one structure and twelve paths per configuration. Complete route closure, post-dispatch arrival changes, booking freezes, richer production technologies and recovery activities are outside the executed core. Initial protective stock is positioned before the episode without modeling its acquisition lead time, and the shrinking horizon has no salvage value or extended recovery tail. These assumptions can materially influence protection choices.

The residual construction is heuristic, and statistical inference conditions on one fitted history and validation selection. Envelope exceedance diagnostics do not establish coverage. Most comparators are internal ablations rather than harmonized implementations of published methods; larger-system checks contain one structure per size. The prospective design in Supplementary Information S1 identifies how these limitations should be tested but does not increase the completed sample.

The analytical extension fixes yields and recovery delays within each solve, treats external secondary/feedstock receipts as exogenous, and aggregates active recovery processing into its initiation period. Stochastic quality outcomes, multiperiod

processing and detailed equipment scheduling require further specification. No completed environmental inventory, emissions reduction or circular-policy performance result is reported. Documented activity data and factors, together with matched service and output comparisons, are needed to evaluate those effects.

9.4. Focused comparative research agenda

The next evaluation should first implement one closely related published procedure with a source-to-code adaptation map and common information, cost and execution rules. It should then test independent material structures at each fixed size, refit forecasting pipelines across independent histories, and vary booking commitments and arrival-information quality. Finally, environmental and circularity variants should be selected on separate validation data and assessed on common held-out paths at predefined service requirements. Supplementary Information S1 preserves the detailed prospective design and reporting rules; execution of those comparisons remains outstanding.

10. Conclusion

This study develops timing-aware decision support for manufacturing exposed to maritime material disruptions. Its implemented core links bookings, protective stock, reserved capacity and feasible production, while matched comparisons distinguish residual protection, production coordination and repeated state updating. The strongest demonstrated result is the value of the specified feedback protocol: mean cost is 19.47% lower than retaining the initial plan. Residual protection reduces cost by 3.53% relative to nominal planning, while advantages over persistence and production-first planning remain small and uncertain.

A separate 26-week, author-confirmed empirical factory ledger records 22.87% on-time dispatch despite 95.34% fulfilment by cutoff and 96.49% regular-hour utilization. Its service and capacity diagnostics show why eventual recovery and punctual output must be distinguished. Protective-stock commitment and recovery headroom remain operational measures, not substitutes for energy, emissions or waste accounting.

The analytical extension specifies qualified secondary inputs, delayed release, carried recovery work in progress, physically supported dispatch and original due-cohort service. Explicit economic and environmental ledgers distinguish acquisitions, consumption, waste and past commitments. The recovery illustration shows that maximizing processed returns can reduce feasible finished output when recovery consumes the binding resource. The central proposition is to protect deadline-feasible, resource-accounted output rather than nominal availability alone. External benchmarking, broader structural tests and comparable-service environmental evaluation remain necessary to establish the performance of the extended model.

Declarations

Funding declaration

This research did not receive any specific grant from funding agencies in the public, commercial or not-for-profit sectors.

Competing interests. The author declares no competing interests.

Author contributions. Islam El-Nakib: conceptualization, methodology, investigation, formal analysis, visualization, writing—original draft, and writing—review and editing. The sole author is accountable for the research and manuscript.

Ethics and operational-data governance. The study reports computational experiments and coded factory operational records; it does not report interviews, surveys or participant interventions. Operational records are reported in coded or aggregated form. External access to restricted records is subject to confidentiality safeguards and the necessary data-provider and institutional authorization. No independent source-system authentication or ethics approval is claimed on the basis of arithmetic consistency checks.

Clinical trial number: not applicable.

Consent to Participate declaration: not applicable.

Consent to Publish declaration: not applicable.

Use of generative AI and AI-assisted technologies. Grammarly and Google Gemini Pro were used for editorial assistance, supported manuscript restructuring, literature searching, formulation drafting, summary-level calculations, figure preparation and document formatting. AI assistance is not independent validation of the operational records or execution of the reported policy experiments. The author is responsible for verification, interpretation and final approval of the methods, evidence, references and manuscript.

Data Availability Statement

The computational code, model outputs, episode-level results and supporting analytical materials for the reported operational experiments are available from the corresponding author upon reasonable request at islam.elnakib@aast.edu. The empirical case uses coded operational records; raw records are not publicly deposited because access is subject to confidentiality and data-provider or institutional authorization. Requests for restricted materials will be considered under the necessary authorization and confidentiality safeguards; de-identified or aggregated materials may be shared where permissible. The analytical circularity formulation is provided in Section 8 and the prospective evaluation protocol in Supplementary Information S1. No completed environmental-extension results archive is claimed.

References

1. Duflou JR, Sutherland JW, Dornfeld D, Herrmann C, Jeswiet J, Kara S, Hauschild M, Kellens K. Towards energy and resource efficient manufacturing: a processes and systems approach. CIRP Annals. 2012;61(2):587–609. https://doi.org/10.1016/j.cirp.2012.05.002

2. UN Trade and Development (UNCTAD). Review of Maritime Transport 2024: Navigating Maritime Chokepoints. Geneva: United Nations; 2024. https://unctad.org/publication/review-maritime-transport-2024

3. Arvis JF, Rastogi C, Ulybina D. Will a prolonged rerouting of ships from Suez trigger a new supply chain crisis? World Bank Blogs. 19 January 2024. https://blogs.worldbank.org/en/trade/will-prolonged-rerouting-ships-suez-trigger-new-supply-chain-crisis

4. Badakhshan E, Ball P. Deploying hybrid modelling to support the development of a digital twin for supply chain master planning under disruptions. International Journal of Production Research. 2024;62(10):3606–3637. https://doi.org/10.1080/00207543.2023.2244604

5. Assid M, Gharbi A, Pellerin R. Machine learning-based dynamic production planning and control in unreliable manufacturing systems with supply disruptions. International Journal of Production Research. 2025;63(23):9229–9247. https://doi.org/10.1080/00207543.2025.2536195

6. Jia M, Schrotenboer AH, Chen F. Scenario Predict-then-Optimize for Data-Driven Online Inventory Routing. Transportation Science. 2025;59(5):1032–1056. https://doi.org/10.1287/trsc.2024.0613

7. Guide VDR Jr, Van Wassenhove LN. OR FORUM—The evolution of closed-loop supply chain research. Operations Research. 2009;57(1):10–18. https://doi.org/10.1287/opre.1080.0628

8. Bocken NMP, de Pauw I, Bakker C, van der Grinten B. Product design and business model strategies for a circular economy. Journal of Industrial and Production Engineering. 2016;33(5):308–320. https://doi.org/10.1080/21681015.2016.1172124

9. Tomlin B. On the value of mitigation and contingency strategies for managing supply chain disruption risks. Management Science. 2006;52(5):639–657. https://doi.org/10.1287/mnsc.1060.0515

10. Saputro TE, Figueira G, Almada-Lobo B. Integrating supplier selection with inventory management under supply disruptions. International Journal of Production Research. 2021;59(11):3304–3322. https://doi.org/10.1080/00207543.2020.1866223

11. Wang M, Amiri-Aref M, Klibi W, Babai MZ. Global multi-sourcing network design with inventory planning under uncertainty. International Journal of Production Research. 2024;62(21):7660–7686. https://doi.org/10.1080/00207543.2024.2323068

12. Dolgui A, Ivanov D, Sokolov B. Reconfigurable supply chain: the X-network. International Journal of Production Research. 2020;58(13):4138–4163. https://doi.org/10.1080/00207543.2020.1774679

13. Garrido A, Pongutá F, García-Reyes H. Zero-inventory plans, constant workforce, or hybrid approach? Analysing pure production strategies for enhancing factory resilience with demand variability. International Journal of Production Research. 2025;63(10):3589–3607. https://doi.org/10.1080/00207543.2024.2425771

14. Bo H, He S, He Z. Inventory-kit modelling and delay-level construction for assembled products with production-transportation collaboration in two-echelon suppliers. International Journal of Production Research. 2026. Advance online publication. https://doi.org/10.1080/00207543.2026.2732142

15. Li D, Jiao J, Wang S, Zhou G, Liu D. System dynamics simulation for supply chain performance under maritime transport disruptions: a case study of Suez Canal blockage. International Transactions in Operational Research. 2026;33(4):2190–2218. https://doi.org/10.1111/itor.70097

16. Niebles-Atencio F, Lalla-Ruiz E, Voß S. Recovery strategies for disrupted seaside planning in container terminals: a systematic review. International Journal of Production Research. 2026. Advance online publication. https://doi.org/10.1080/00207543.2026.2719841

17. Liu X, Ho W, Lam HKS, Talluri S. Reconfiguring supply chains amid geopolitical risks: a systematic review of emerging topics. International Journal of Production Research. 2026. Advance online publication. https://doi.org/10.1080/00207543.2026.2704029

18. Bertsimas D, Kallus N. From predictive to prescriptive analytics. Management Science. 2020;66(3):1025–1044. https://doi.org/10.1287/mnsc.2018.3253

19. Elmachtoub AN, Grigas P. Smart "Predict, then Optimize". Management Science. 2022;68(1):9–26. https://doi.org/10.1287/mnsc.2020.3922

20. Mandi J, Kotary J, Berden S, Mulamba M, Bucarey V, Guns T, Fioretto F. Decision-focused learning: foundations, state of the art, benchmark and future opportunities. Journal of Artificial Intelligence Research. 2024;80:1623–1701. https://doi.org/10.1613/jair.1.15320

21. Arora R, Sachs AL, Svetunkov I, Muth M, Boylan JE. Forecasting demand during and after supply chain disruptions using a shock smoother ETS. International Journal of Production Research. 2025. Advance online publication. https://doi.org/10.1080/00207543.2025.2577158

22. Poles R. System Dynamics modelling of a production and inventory system for remanufacturing to evaluate system improvement strategies. International Journal of Production Economics. 2013;144(1):189–199. https://doi.org/10.1016/j.ijpe.2013.02.003

23. Flapper SD, Gayon JP, Lim LL. On the optimal control of manufacturing and remanufacturing activities with a single shared server. European Journal of Operational Research. 2014;234(1):86–98. https://doi.org/10.1016/j.ejor.2013.10.049

24. Reddy KN, Kumar A. Capacity investment and inventory planning for a hybrid manufacturing–remanufacturing system in the circular economy. International Journal of Production Research. 2021;59(8):2450–2478. https://doi.org/10.1080/00207543.2020.1734681

25. Lahmar H, Dahane M, Mouss KN, Haoues M. Multi-objective sustainable production planning for a hybrid multi-stage manufacturing-remanufacturing system with grade-based classification of recovered and remanufactured products. Journal of Intelligent Manufacturing. 2025;36:1385–1407. https://doi.org/10.1007/s10845-023-02308-9

26. Hoffmann M, Krini A, Mueller A, Knorn S. Remanufacturing production planning and control: Conceptual framework for requirement definition. Journal of Remanufacturing. 2025;15:97–126. https://doi.org/10.1007/s13243-025-00149-8

27. World Resources Institute, World Business Council for Sustainable Development. Product Life Cycle Accounting and Reporting Standard. Greenhouse Gas Protocol; 2011. https://ghgprotocol.org/product-standard

28. Ben-Tal A, El Ghaoui L, Nemirovski A. Robust Optimization. Princeton: Princeton University Press; 2009.

29. Bertsimas D, Sim M. The price of robustness. Operations Research. 2004;52(1):35–53. https://doi.org/10.1287/opre.1030.0065

30. Man X, Chang D, Zheng P, Han W. A logistics ratio-optimised inventory routing problem with lateral transhipment using mixed-integer robust fractional programming. International Journal of Production Research. 2025;63(24):10381–10412. https://doi.org/10.1080/00207543.2025.2551240

31. Lejarza F, Venkatesan S, Baldea M. Rolling horizon product quality estimation and online optimisation for supply chain management of perishable inventory. International Journal of Production Research. 2025;63(10):3709–3732. https://doi.org/10.1080/00207543.2024.2427891

32. Ivanov D, Dolgui A. A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. Production Planning & Control. 2021;32(9):775–788. https://doi.org/10.1080/09537287.2020.1768450

33. Vaezinejad S, Pishvaee MS, Schniederjans D. A critical review of optimisation models in smart supply chains: insights and future directions. International Journal of Production Research. 2026;64(18):8021–8054. https://doi.org/10.1080/00207543.2026.2645688

34. Ruel S, Zsidisin GA, Blackhurst J. A taxonomy of supply chain resilience. International Journal of Production Research. 2026;64(18):7997–8020. https://doi.org/10.1080/00207543.2026.2645249

Appendix A. Reported computational detail and summary-level calculations

A.1. Computational driver and seed specification

The reported operational experiment is identified by run_v2.py (experimental driver), analyse_v2.py (tables and intervals), test_model_v2.py (regression checks) and plot_v2.py (figures). Reproduction requires the matching model version, parameter manifest, episode outputs and dependency environment. The descriptive calculations below use the reported summary values; the analytical recovery illustration is separate from the policy experiment banks.

The reported training/calibration history uses seed 20261001; independent accuracy assessment uses 410001. Numbering the twelve persistence/severity cells from zero in lexicographic order, validation seeds are $110000 + 100c + r$ for $r = 0{,}1$, and main seeds are $210000 + 100c + r$ for $r = 0, \ldots, 9$. Factorial seeds are $310000 + 100c + r$ for eight replications per cell. Initial scaling uses 510000–510004 within each size. Medium and larger policy extensions use 610000–610011 and 620000–620011. Material structures use seed 77. Stratified bootstrap resampling uses 20261001 and factorial intervals use 2929.

The reported optimizer gap target is the fraction 10^{-6}, equivalent to 0.0001%. Library versions, hardware details, model checksums and episode manifests identify the matching operational archive. Reproduction from rounded summaries alone cannot recover those inputs. The prospective banks in Supplementary Information S1 require separate manifests and are excluded from the completed sample counts.

A.2. Dispersion and adjusted comparisons

Table 15 retains the main cost standard deviations. Their denominator is the common 120-path bank, not the number of individual weekly solves.

Table 15. Reported standard deviation of realized episode cost

Policy	Cost standard deviation
GB-RO	81.61
ARX-RO	54.16
RF-RO	96.34
Persistence-RO	116.24
Historical-RO	72.21
GB-Regime	84.46
GB-Nominal	326.49
Production-first	84.77
GB-Static	907.52
Perfect-information	52.27

The higher dispersion for nominal and static planning is consistent with the reported descriptive tail differences, but a standard deviation does not establish a probabilistic reliability guarantee.

Table 16 retains the multiplicity-adjusted bootstrap intervals from the source study. Rounding of marginal policy means can make their subtraction differ by 0.01 from a paired difference retained at source precision.

Table 16. Source-reported paired cost differences and Bonferroni-adjusted intervals

Comparator	Paired difference	Adjusted interval
GB-Nominal	-87.75	[-142.91, -36.34]
Persistence-RO	-4.33	[-31.31, 19.91]
ARX-RO	6.73	[-5.82, 19.30]
RF-RO	-12.16	[-28.31, 3.20]

Comparator	Paired difference	Adjusted interval
Historical-RO	36.33	[26.09, 47.00]
Production-first	-1.23	[-3.26, 0.77]
GB-Static	-579.95	[-771.69, -409.04]

Negative differences favor GB-RO. The intervals use the archived episode-level bootstrap. The historical-mean comparison favors Historical-RO on cost; its lower punctual service is reported separately.

A.3. Complete factorial cells

Table 17 gives every reported factorial cell. Cell numbers match Fig. 3. Gains are mean paired percentage advantages, not ratios of pooled means; each cell has eight independent paths.

Table 17. Complete reported factorial results

Cell	Coverage (weeks)	Capacity multiplier	Severity	Mean gain (%)	Pointwise interval (%)
1	0.5	0.90	0.8	0.024	[-0.1900, 0.2310]
2	0.5	0.90	1.1	0.365	[-0.0600, 0.8300]
3	0.5	0.90	1.4	0.100	[-0.0100, 0.2420]
4	0.5	1.00	0.8	0.193	[-0.0780, 0.5070]
5	0.5	1.00	1.1	0.460	[0.1880, 0.7560]
6	0.5	1.00	1.4	0.430	[-0.0180, 0.9420]
7	0.5	1.20	0.8	0.138	[0.0410, 0.3050]
8	0.5	1.20	1.1	0.083	[-0.0560, 0.2400]
9	0.5	1.20	1.4	0.121	[-0.0210, 0.3440]
10	1.0	0.90	0.8	-0.025	[-0.0750, 0.0000]
11	1.0	0.90	1.1	-0.137	[-0.4020, 0.0770]
12	1.0	0.90	1.4	-0.080	[-0.1560, -0.0200]
13	1.0	1.00	0.8	0.085	[0.0000, 0.2540]
14	1.0	1.00	1.1	0.092	[-0.0640, 0.2640]
15	1.0	1.00	1.4	0.135	[-0.0280, 0.3570]
16	1.0	1.20	0.8	0.058	[-0.0790, 0.2520]
17	1.0	1.20	1.1	0.022	[-0.0010, 0.0590]
18	1.0	1.20	1.4	0.017	[-0.1670, 0.2300]
19	2.0	0.90	0.8	-0.016	[-0.0430, -0.0002]
20	2.0	0.90	1.1	-0.035	[-0.0850, 0.0010]
21	2.0	0.90	1.4	0.067	[-0.0230, 0.1990]
22	2.0	1.00	0.8	0.000	[0.0000, 0.0000]
23	2.0	1.00	1.1	0.282	[-0.0280, 0.7360]

Cell	Coverage (weeks)	Capacity multiplier	Severity	Mean gain (%)	Pointwise interval (%)
24	2.0	1.00	1.4	-0.005	[-0.3140, 0.2410]
25	2.0	1.20	0.8	-0.015	[-0.0460, 0.0000]
26	2.0	1.20	1.1	-0.285	[-0.7290, 0.0005]
27	2.0	1.20	1.4	0.067	[-0.0880, 0.2330]

The mean across all 216 paths is approximately 0.079%; eight cell means are negative. Only two pointwise intervals are wholly positive. The largest observed cell mean remains an exploratory result rather than an externally validated decision threshold.

A.4. Added arithmetic and analytical illustration

The punctual-output ratio uses reported mean cost divided by $200\overline{S}/100$, where $\overline{S}$ is the main-table on-time percentage. Optional buffer utilization is $100\overline{b}/120$. Both calculations use rounded source means and are reported to two decimal places. The post-hoc service screen selects the lowest reported mean cost among the eight rolling policies satisfying the displayed mean-service threshold; it excludes static and perfect-information policies and performs no new parameter tuning or statistical inference.

The weekly case calculations sum the 26 listed available and used hours, count zero on-time observations and inspect the listed closing backlog. They reproduce 5,064 due pieces, 1,158 on-time pieces, 2,744.05 used hours and 2,844 available hours. Added descriptive values are 99.95 unused hours in total, 0.85 in weeks 15–26, eight zero-service weeks and maximum closing backlog of 364 in week 20. They are summary-table checks, not a repeat of the transaction-level audit.

The recovery illustration substitutes the stated assumptions into (21). For recovery feed of 0, 10, 20, 40 and 50 units, feasible output is 60, 68, 76, 72 and 70, respectively. The crossing from (22) is 20 feed units. This arithmetic is independent of the original generated policy episodes and does not introduce a new environmental result.

A.5. Descriptive service-threshold screen

A further descriptive screen illustrates sensitivity to the stated service criterion. Table 18 retains only implementable rolling policies whose reported mean on-time service is at least the displayed threshold and identifies the lowest reported mean cost among them. This is a post-hoc screen of test means, not validation-constrained selection, a reliability certificate or a deployment recommendation.

Table 18. Post-hoc service-threshold screen of the eight reported rolling-policy means

Threshold (%)	Eligible means	Lowest-cost eligible policy	Mean cost	On-time (%)
95.0	7	Historical-RO	2,362.19	96.51
96.5	7	Historical-RO	2,362.19	96.51
97.0	4	ARX-RO	2,391.80	97.24
97.2	2	ARX-RO	2,391.80	97.24
98.0	0	None	Not applicable	Not applicable

The identity of the lowest-cost eligible mean changes when the threshold increases from 96.5% to 97.0%. At 98%, none of the reported implementable rolling-policy means qualifies. This does not show that the target is infeasible or that no differently tuned policy could meet it. A fair prospective comparison must select settings on separate validation data and evaluate them without retuning on held-out paths.

Appendix B. Case-record provenance, complete weekly outcomes and source registers

B.1. Provenance and relative time

The empirical source is IJPR__Research_Data.xlsx. The author identifies its records as original, reviewed and validated operational data. Internal validation concerns transaction linkage, identifiers, units, ranges and cross-sheet balances. Legacy projection/generation labels and associated template metadata are retained in the archived source for traceability and treated according to the author's provenance correction. The analytical dataset distinguishes observed fields, derived research variables and scenario assumptions. This attribution is not independent authentication against the originating information system.

A uniform 273-day shift of operational timestamps is documented in the preparation record. Relative-week reporting preserves order-to-dispatch intervals, usable-release timing and weekly service without treating shifted calendar labels as the original observation period. The purpose of the shift, originating-system extraction dates and original calendar coverage are not independently established here. Source-system verification and external disclosure remain separate from internal analytical validation.

Forecasts are treated as recorded research calculations: the source audit reports that all 52 demand forecasts reproduce the rounded mean of up to four previously due gross order quantities, including opening overdue orders. Cost rates and accruals retain their stated research basis rather than being described as authenticated cash payments. The original sheet name "Projected charges" is a locator, not an adopted classification of every underlying transaction.

B.2. Coverage and reconstruction conventions

Table 19 reports case coverage and inclusion rules used in the documented reconstruction. The counts describe linked operational records; they are not counts of independent policy experiments.

Table 19. Source-reported case coverage

Record group	Count	Inclusion or interpretation
Master items; core BOM links	5; 4	Product/material definitions
Orders	54	52 in-window lines plus two opening overdue lines
Deliveries	80	Partial dispatch events retained
Inventory	135	Opening and weekly usable-stock snapshots
Shipments; receipts	67; 78	Commitments, physical receipts and usable release
Complete ETA events	101	52 source-label-only rows excluded
Production; capacity	52; 26	Two products and one resource calendar over 26 weeks
Cost rates; charge rows	16; 367	Window cutoff applied to accruals
Planning decisions	133	Bookings, reserve and production actions
Recorded disruptions	1	Route restriction in weeks 9–13
Forecasts and signals	104	52 demand forecasts and 52 route warnings
Route observations	208	Offered/realized normal and reserve allowances

The source reports 27 scoped case checks: 26 pass and one flags the 52 excluded incomplete ETA rows. Recorded production uses the Friday aggregation convention; material is credited at usable release, not merely at booking or ETA revision. Physical tolerances are 10^{-7} base units and SAR 0.011 per rounded charge.

B.3. Complete weekly summaries

Table 20 retains every source-reported week, including zero-service weeks. Closing backlog includes all overdue quantities remaining at the cutoff, not only the current week's cohort.

Table 20. Complete reported weekly case outcomes

Week	Net due	On time	Service (%)	Closing overdue	Hours used / available
1	193	185	95.85	8	109.65 / 110
2	205	182	88.78	23	109.55 / 110
3	227	135	59.47	92	87.25 / 110
4	225	93	41.33	150	104.80 / 110
5	229	78	34.06	159	110.00 / 110
6	206	48	23.30	208	101.90 / 102
7	193	58	30.05	181	110.00 / 110
8	190	37	19.47	202	109.85 / 110
9	182	67	36.81	164	110.00 / 110
10	176	39	22.16	171	109.85 / 110
11	168	0	0.00	205	67.00 / 110
12	176	57	32.39	212	109.85 / 110
13	183	31	16.94	175	110.00 / 110
14	191	25	13.09	238	83.20 / 110
15	206	39	18.93	224	110.00 / 110
16	216	20	9.26	271	109.85 / 110
17	212	25	11.79	263	110.00 / 110
18	236	0	0.00	342	101.90 / 102
19	219	0	0.00	341	110.00 / 110
20	192	0	0.00	364	109.85 / 110
21	181	0	0.00	325	110.00 / 110
22	172	0	0.00	328	109.85 / 110
23	168	0	0.00	276	110.00 / 110
24	172	0	0.00	279	109.85 / 110
25	169	20	11.83	228	110.00 / 110
26	177	19	10.73	236	109.85 / 110

The overall service ratio uses summed quantities rather than the arithmetic mean of weekly percentages. The case's 95% target is an evaluation criterion, not an attained performance level. Weekly observations belong to one continuing factory ledger and are not treated as 26 independent intervention replications.

B.4. Additional workbook registers are not additional outcome banks

The source describes nine production-structure definitions: dedicated, shared and mixed labels at sizes 2 products/3 materials, 4 products/5 materials and 8 products/9 materials. Each definition lists one resource, with nominal weekly hours of 110, 220 and 440 by size; 84 associated BOM rows are reported. These labels do not prove independently observed factories, distinct material topologies or successful comparative evaluations.

The episode register contains 30 definitions of 26 weeks, with ten each labeled Training, Validation and Test, and numeric seeds 20261003–20261032. Demand multipliers span 0.85–1.15, normal-capacity fractions 0.3–0.7 and extra transit times 7, 14 or 21 days. Eighteen settings combine freezes of 1, 8 and 15 days with ETA refreshes of 1, 3 and 7 days and service targets of 95% or 98%.

The structure and setting registers contain configuration fields rather than additional per-policy cost, service or environmental series. They are not included as completed external-benchmark, commitment/ETA, broader-structure or comparable-service outcome banks. The reported evidence remains the controlled computational comparisons and the continuing factory baseline.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- DiscoverSustainabilitySupplementaryProtocolS1.docx