The Eastasouth Journal of Information System and Computer Science
Vol. 2, No. 02, December 2024, pp. 398-410
ISSN: 3025-566X, DOI: 10.58812/esiscs.v2i02

Inventory Replenishment with Delayed Stock Reports

Danish Mahmud
Addis Ababa University, Addis Ababa, Ethiopia

Article Info

Article history:

Received Dec, 2024
Revised Dec, 2024
Accepted Dec, 2024

Keywords:

Inventory Visibility;
Replenishment;
Safety Stock;
Snapshot Age;
Supply Chain Analytics

ABSTRACT

Inventory dashboards often display the most recent count even when it is several days old. Replenishment decisions based on that count face two distinct errors. Demand since the snapshot reduces the expected stock position, and its unknown size increases uncertainty. We derive an age-conditioned order quantity for a single replenishment cycle with backlogged demand and compare it with a rule that ignores snapshot age and a rule that corrects only expected depletion. An exact Poisson benchmark evaluates five snapshot ages and three reporting environments. At seven days of age, mean correction reduced expected cost from 208.00 to 13.71 per cycle. Conditioning the stock target on the full distribution reduced cost further to 12.72 and increased the probability of no terminal shortage from 69.7% to 80.1%. Under a mixed-age portfolio, the two age-aware policies cost 9.34 and 9.09, respectively. The first gain comes from correcting a biased inventory position. The second comes from allowing uncertainty to grow with the unobserved interval. The benchmark separates these mechanisms and provides a reproducible basis for testing timestamp-aware replenishment using operational data.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Name: Danish Mahmud
Institution: Addis Ababa University, Addis Ababa, Ethiopia
Email: danishmahmud.ai@gmail.com

1. INTRODUCTION

A store can have a reliable daily demand forecast and still order the wrong quantity. The inventory count that appears next to the forecast may have been accurate when it was recorded. If several days of demand have occurred since then, it may no longer describe what is available. The problem becomes harder when systems continue to present the most recent value as if it were current. A planner looking at a plausible stock number sees no obvious signal that an order is too small. The number and the timestamp must be read together to understand what the inventory system actually knows.

A stale observation changes the decision in two ways. The first is a shift in the expected stock position. On average, inventory declines by expected demand during the unobserved interval when there are no other movements. Subtracting this expected consumption addresses the shift. The second is increased uncertainty about the remaining position. Subtracting a mean does not establish how much demand actually arrived. If demand fluctuates, an estimate adjusted by the mean still has an error that generally grows with the length of the interval. A planner who uses the same safety

margin for fresh and old observations has handled only one component of the problem.

This study asks how much of the cost of stale inventory information can be attributed to those two components. We hold the demand model and planning horizon fixed while changing the age of the stock observation. We evaluate three order rules using the same underlying demand distribution. The first rule treats the latest available count as current. The second subtracts expected unobserved consumption, while retaining the safety allowance used for fresh counts. The third combines unobserved and future demand in a single distribution and uses its quantile to determine the order. The comparison makes the role of the timestamp explicit without confounding it with changes in forecasting models.

The research contribution is a transparent decomposition rather than a new forecasting algorithm. We establish a conditional cost benchmark, calculate the corresponding service outcomes and examine how the results aggregate when different stock records are different ages. A small sensitivity grid then checks how the value of using the age-conditioned quantile changes with the demand rate, reported stock and shortage penalty. All numerical results are calculated from a specified probability model. This design supports reproducibility while keeping the scope of the evidence clear. It characterizes a single-cycle inventory decision with one uncertain demand process and an available timestamp.

2. RELATED RESEARCH

2.1 Forecasts and inventory decisions

The forecasting literature supplies tools for estimating what customers may buy, but an ordering policy must also account for what can still be supplied. Comparative work on supermarket forecasting evaluates several learning algorithms, including tree-based models, and illustrates the range of ways to estimate demand from retail records [1].

Predictive analytics research on supply chain resilience situates forecasting within management information systems, where analytics must inform an operational response [2].

Work on smart manufacturing likewise places data-driven decisions in a setting where the measurements and processes supporting them can change over time [3].

The customer experience literature provides a reason to consider service outcomes alongside inventory costs. A stock decision can change whether a customer receives an item even if aggregate demand estimates remain accurate [4].

Energy-market analytics gives a related example of how decisions based on projections must accommodate variability and operational constraints [5].

In this study, demand is described by a deliberately simple model. This keeps the forecasting component constant while the treatment of the inventory observation changes. Replacing a sophisticated forecast with a simpler one might worsen demand estimates in a real deployment. Here the question is different. Even if the distribution of future demand is known, what ordering error remains when the current stock position is uncertain? The design answers that question by making the assumed demand process common to every policy.

2.2 Information systems and observation provenance

A timestamp is useful only if it describes the event that matters. The time a dashboard fetched a record may differ from the time the warehouse counted the units. If an old record is reprocessed, the ingestion time may look fresh even when the physical count remains old. For replenishment, the relevant clock starts at the effective inventory observation. Data governance research provides the broader rationale for assigning ownership and maintaining a shared definition of that field across systems [6].

Explainable decision-support research motivates showing the quantities that entered the order calculation. A planner should be able to distinguish the last reported count, the estimate of unobserved consumption and the allowance for uncertainty [7].

The cloud project-management literature offers context for separating the systems that transport records from the decision logic that consumes them [8].

Research on agile project management informs how organizations can stage a change to a decision workflow and compare its outputs before broad deployment [9].

Project management systems that integrate analytics with management information systems provide a parallel for putting analytical recommendations into an existing process rather than issuing them in isolation [10].

Research on employee retention analytics gives an example outside inventory in which a predictive output is used to inform an allocation decision. It supports the distinction between estimating an outcome and selecting an action in response to that estimate [11].

Research on microcredit supplies a separate example of why the consequences of a service decision may differ across recipients. That context motivates examining whether stale records cluster in certain locations before drawing conclusions from portfolio averages [12].

2.3 Transferable measurement lessons

Several references from health analytics inform the measurement framing without supplying parameters for the inventory model. Diabetes progression research considers prediction and ranking stability, which illustrates that a model can be evaluated differently depending on the downstream decision [13].

Survey-weighted cardiovascular research illustrates the distinction between population-level characterization and the evidence needed for a decision about an individual observation [14].

Work on chronic disease risk stratification provides another example in which the utility of an estimate depends on the use made of it [15].

Research on cancer analytics and cost optimization motivates treating prediction and the costs of acting on it as separate components of an evaluation [16].

Real-time epidemiological surveillance provides a particularly close analogy for the difference between event time and reporting time [17].

Machine learning for epidemic prediction offers a related context in which the latest reported value may cover a different time window from the one a decision maker intends to assess [18].

Multisource geospatial analytics makes the importance of consistent provenance visible when measurements from different systems are combined [19].

Multi-omics research provides a methodological analogy for integrating sources that describe different aspects of the same underlying process [20].

Bioinformatics for discovery illustrates why an output can be useful for planning a subsequent decision while still requiring independent confirmation [21].

AI-assisted drug discovery provides another example of the distinction between a model output and evidence that a downstream decision worked [22].

2.4 Integrity and operational continuity

A delayed report can arise from mundane batching, a connectivity problem, or a disruption in the systems that carry the report. The inventory model conditions on age without assuming any particular cause. Research on protection in cloud environments supplies context for preserving the

integrity of the records and timestamps on which downstream decisions rely [23].

Work on cloud-integrated management information systems similarly places governance and data protection alongside the analytics that use these records [24].

Cognitive cyber defense research offers an organizational parallel for maintaining visibility when several services are responsible for moving or processing the same records [25].

Defensive threat detection research illustrates the distinction between observing a signal and arranging a response to it [26].

Research on emerging threats provides context for designing continuity measures for information systems that support operational decisions [27].

Work on integrating threat intelligence with project management points to the importance of assigning responsibility for responding when monitoring detects an abnormal condition [28].

Research on training and employee behavior supports treating the interpretation of warnings as part of a workflow, since an alert is useful only when staff know how to act on it [29].

Distributed threat intelligence provides a further analogy for checking whether independently maintained records have compatible definitions and time coverage [30].

Data center research considers the resource costs of maintaining secure infrastructure. It motivates keeping the economics of more frequent reporting separate from the savings from adjusting an order when a report is already stale [31].

3. METHODS

3.1 Decision timing and assumptions

Consider one stockkeeping unit at one location. A snapshot records a net inventory position of x units. The order decision occurs a days after that observation. Demand between the snapshot and the decision is unknown, and there are no unrecorded supply receipts, transfers or inventory adjustments in this interval. The planner chooses a nonnegative integer order q, which is available by the terminal evaluation date. The forward interval from the decision to that date is H days. Demand that cannot be met immediately is backlogged. The objective charges for terminal surplus and terminal shortage only. It does not measure the time a customer waited, and it does not imply that every order arrived before every demand event. Those distinctions are necessary when interpreting the reported service probabilities.

The model assumes independent daily demand with a Poisson rate λ. Total demand from the snapshot to the end of the forward interval, D, then has a Poisson distribution with mean $\lambda(a + H)$. The sum contains demand that already occurred but is not yet visible and demand that has not occurred. Although the two portions differ in time, their effect on terminal inventory is additive under the specified stock-flow assumptions. The ending net position is $x + q - D$. A positive position is surplus and a negative position is backlog. The age is known when the order is placed, and λ is known in the experiment.

3.2 Ordering rules and objective

Let $s = x + q$ denote the total stock available to cover demand measured from the snapshot. Let h be cost per unit of terminal surplus and b be cost per unit of terminal backlog. For a given age, expected loss is $h\,E[(s - D)^+] + b\,E[(D - s)^+]$. Increasing s by one raises holding cost by h when $D \le s$ and avoids shortage cost b when $D > s$. The discrete incremental cost is therefore $(h + b)\,P(D \le s) - b$. The smallest integer s whose cumulative probability is at least $b/(h + b)$ minimizes the unconstrained objective. If the snapshot stock already exceeds this

target, the nonnegative order constraint sets q to zero.

The age-ignorant rule computes a target from the fresh-count distribution with mean λH and orders against the reported stock. Its decision does not vary with age. The mean-correction rule adds λa units to the fresh target. The target increases enough to cover expected unobserved demand, while the safety allowance above the expected total stays constant. The age-specific rule computes the critical-fractile quantile from the full demand distribution. This target reflects both expected depletion and the larger variance associated with additional days. All demand rates and ages used in the reported design are integers, so the mean-correction target is also integer-valued.

The age-specific rule is a model-consistent optimum under the declared objective. Its advantage over the two simpler rules is consequently guaranteed in expected cost within the model. The experiment measures the size of that advantage and its operational components. It does not test whether the assumed demand law forecasts real demand well. A useful implementation would compare actual conditional demand distributions with the chosen model, particularly when reporting delays cluster around events that also change demand.

3.3 Benchmark and sensitivity design

The base case uses $\lambda = 8$ units per day, a three-day forward interval, a reported stock position of 18 units and holding and shortage costs of 1 and 4 per unit, respectively. The critical fractile is 0.80. Five snapshot ages are evaluated: zero, one, two, four and seven days. The ages provide examples from a current observation to one a week old. They are experimental inputs rather than estimates of how often a particular retailer experiences reporting delays.

The portfolio analysis weights the five age-specific results to represent three environments. In ascending age order, the mostly fresh environment uses weights 0.65, 0.20, 0.10, 0.04 and 0.01. The mixed environment uses weights 0.20, 0.25, 0.25, 0.20 and 0.10. The mostly delayed environment uses weights 0.05, 0.10, 0.15, 0.30 and 0.40. These are mixtures of independent decision cases with otherwise identical parameters. They do not describe repeated inventory cycles for one item. Weighting the age-specific values lets the portfolio contain a realistic mixture of record ages without assuming a particular time series for outages.

The sensitivity grid changes the daily demand rate to 3, 8 or 16 units, reported stock to 6, 18 or 36 units, and the shortage penalty to 2, 4 or 9. It retains the three-day forward interval, the holding cost of 1 and the mixed reporting environment. The grid produces 27 combinations and 81 policy evaluations. It examines whether the approximate rule performs similarly across scales and costs, and whether a high starting stock position limits the practical effect of changing the target.

3.4 Outcomes and computation

We calculate order size, expected surplus, expected backlog, expected terminal cost and the probability of no terminal shortage. For a Poisson variable with mean μ, expected backlog at s is $\mu\, P(D \geq s) - s\, P(D > s)$. Expected surplus is $s - \mu$ plus expected backlog. The no-shortage probability is $P(D \leq s)$. This calculation uses cumulative and survival probabilities directly. It avoids simulation noise and does not require a finite truncation of the demand support. The reported values are exact up to the numerical precision of the probability functions.

Terminal no-shortage probability is a cycle-level outcome. A cycle with one backlogged unit and a cycle with ten are both counted as having a shortage. Expected backlog preserves that difference. Similarly, expected holding cost reflects the size of surplus when it occurs. These measures should

be read together. A policy can achieve a high probability of no terminal shortage by buying substantially more stock, but that strategy may be undesirable when carrying costs, storage constraints or obsolescence are important.

4. RESULTS

At age zero, all policies order ten units and produce the same expected cost of 7.03. The service probability is 82.3%, slightly higher than the 80% critical fractile because a discrete demand distribution cannot generally attain a specified cumulative probability exactly. This agreement serves as an internal check. Once the observation becomes stale, only the rules that use age change their order quantity. At one day of delay, the age-ignorant rule still orders ten units. Mean correction orders eighteen and the age-specific rule orders nineteen.

Table 1. Expected terminal cost and no-shortage probability by observation age

Age	Ignore age cost	Mean correction cost	Age-specific cost	Ignore age service	Mean correction service	Age-specific service
0	7.03	7.03	7.03	82.3%	82.3%	82.3%
1	19.78	8.15	8.10	27.4%	79.0%	83.5%
2	48.25	9.21	9.03	2.9%	76.6%	81.0%
4	112.00	11.14	10.66	0.0%	73.1%	80.9%
7	208.00	13.71	12.72	0.0%	69.7%	80.1%

The difference between the two corrections becomes more visible as age grows. At seven days, mean correction adds 56 units to the original fresh target and orders 66. The age-specific rule orders 69. Those three additional units reduce expected backlog from 1.94 to 1.14 while increasing expected surplus from 5.94 to 8.14. Under a four-to-one shortage penalty, the reduction in backlog outweighs the added surplus. Expected cost falls by about one unit per cycle. The no-shortage probability rises from 69.7% to 80.1%, which is a larger relative change than the cost improvement.

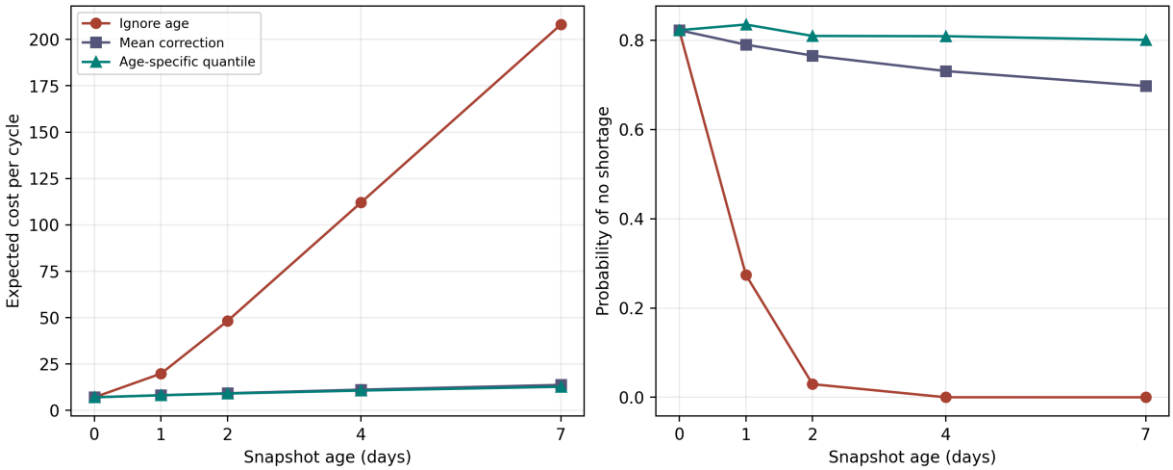


Figure 1. Expected terminal cost and service by snapshot age

The dominant effect is correction of the expected inventory position. At age seven, applying the fresh rule to a week-old observation leaves a total supply position of 28 units against a demand mean of 80. Expected terminal shortage is about 52 units. The resulting expected cost of 208.00 reflects a large mismatch between the quantity ordered and the time span it must cover. Mean correction addresses almost all of that mismatch. The quantile rule then adjusts the remaining difference caused by a safety allowance that stayed fixed even though the uncertain demand interval grew.

Table 2. Portfolio results by reporting environment

Environment	Policy	Expected cost	No terminal shortage
Mostly fresh	Ignore age	19.91	59.2%
Mostly fresh	Mean correction	7.70	80.5%
Mostly fresh	Age-specific quantile	7.64	82.3%
Mixed	Ignore age	61.61	24.0%
Mixed	Mean correction	9.34	76.9%
Mixed	Age-specific quantile	9.09	81.8%
Mostly delayed	Ignore age	126.37	7.3%
Mostly delayed	Mean correction	11.38	73.3%
Mostly delayed	Age-specific quantile	10.80	80.9%

Averaging by age changes the size of the benefit. In the mostly fresh environment, moving from mean correction to the age-specific target reduces expected cost from 7.70 to 7.64. In the mixed environment, the reduction is from 9.34 to 9.09. In the mostly delayed environment, it is from 11.37 to 10.80. The age-specific rule also raises the probability of no terminal shortage in each environment. It orders more units on average, so its service gain comes with a measurable increase in terminal surplus. The objective values show when that trade is favorable under the assumed shortage penalty.

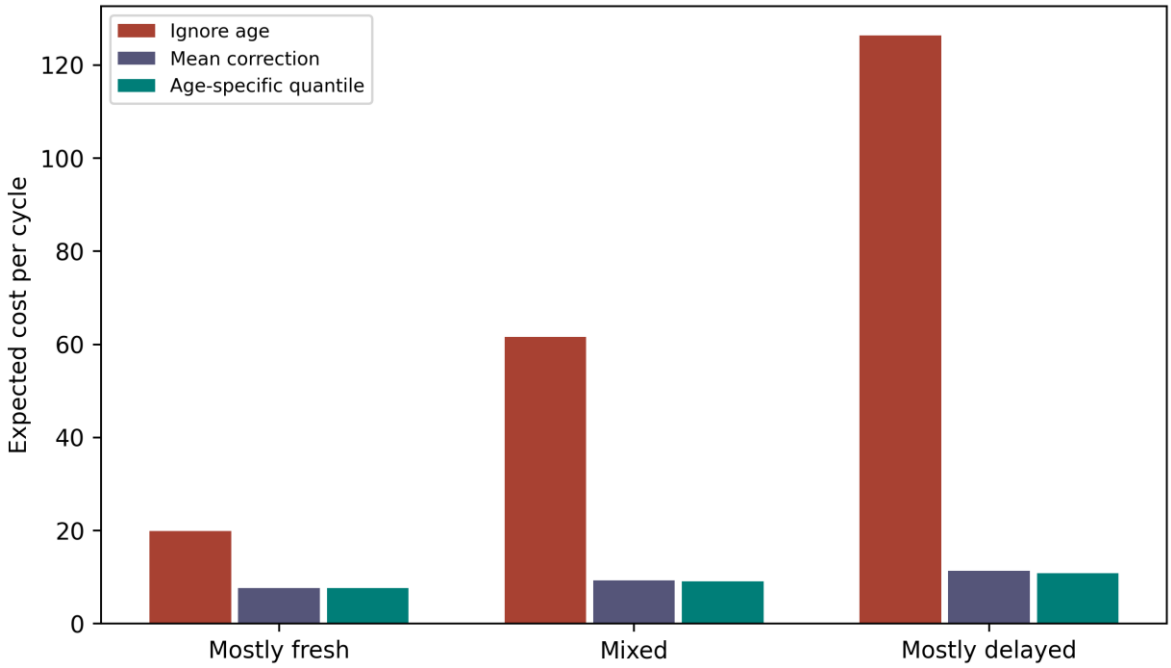


Figure 2. Expected terminal cost across reporting environments

The sensitivity grid confirms that the age-specific quantile weakly minimizes modeled cost throughout the evaluated combinations. The value of that adjustment is not constant. Starting with more stock can make order rules agree because none can order a negative quantity. Where the order constraint does not bind, the difference depends on how far the mean-correction target lies from the conditional quantile and how much cost changes across nearby integers. The sensitivity file reports every combination, including cases with little or no incremental value. It can be used to distinguish situations where a more detailed implementation changes the decision from situations where both rules recommend nearly the same order.

What an additional day changes

The marginal effect of another day is easier to see by separating the center of the demand distribution from its upper tail. In the base case, each day adds eight units to

expected demand. The fresh target is 28, which is four units above the mean of 24. Mean correction keeps this four-unit margin regardless of age. At age seven, the mean is 80, and the mean-correction target is 84. The age-specific target is 87, which puts seven units above the same mean. Both rules incorporate the same expected depletion. The difference lies entirely in the buffer and the discrete placement of the quantile.

The baseline numbers also show why an improvement in the probability of avoiding any shortage can coexist with a smaller improvement in expected cost. The objective prices the units that are left over or backlogged. The probability counts any positive shortage equally. Moving the terminal target across a region where many demand outcomes occur can improve that probability appreciably even if the marginal change in unit costs is modest. For an organization with an explicit service target, both metrics are relevant. One describes the financial trade under the selected weights. The other describes how often the stock position meets all demand by the terminal date.

There are two distinct limits to the value of adjusting the order. The first is the lower bound of zero. When a reported position exceeds both targets, the stock already present determines the terminal supply. Adjusting a target then has no effect on the actual order. The second limit is the local shape of expected cost near the optimum. Two distinct order quantities can produce similar expected costs when each increment trades off nearly equal changes in surplus and backlog. It would therefore be misleading to judge the importance of a policy solely by whether the two targets differ. The decision difference should be considered alongside the realized cost surface and service requirement.

The computed sensitivity results illustrate these limits. The age-specific rule and mean correction tied in two of 27 settings. Both ties occurred at a daily demand rate of three and reported inventory of 36, with shortage penalties of two and four. The largest improvement was approximately 1.08 cost units per cycle. It occurred when demand averaged 16 units per day and the shortage penalty was nine, with the same improvement at all three tested stock positions. This repetition across stock positions is consistent with a setting where the relevant targets remain above the reported position, so changing that position changes the order quantity without changing terminal supply under either rule. The full grid is included so these patterns can be checked directly.

5. DISCUSSION

The most consequential error in this benchmark is treating an old observation as a current count. That practice leads the ordering system to overlook units likely consumed since the snapshot. The mistake is separate from forecasting accuracy. Every policy in the benchmark has access to the same demand rate. One policy fails because it applies that rate to the wrong interval. This distinction matters when diagnosing stockouts. Improving a forecast can reduce error about future demand, but even a correct forward forecast leaves the inventory position biased if the decision ignores elapsed demand.

The two age-aware policies differ in how they treat the remaining uncertainty. A fixed safety allowance can be adequate when the report is recent. It becomes less protective as the inventory estimate must absorb more unobserved demand. Under the benchmark assumptions, the mean and variance both rise linearly with the length of the interval. Mean correction addresses the former. The quantile rule reflects both. In practice, the decision to implement full conditioning depends on whether the added service gain is worth the engineering and validation effort. A mostly fresh operation might obtain most of the available modeled benefit from a simple correction. A location with persistent delays may have more to gain from estimating the conditional distribution carefully.

The difference in terminal service can matter even where the cost gap appears small. The model converts a stock shortage into a linear cost and a discrete event indicator. Real organizations may also face service

agreements, expedited shipments or repeat orders that make small shortfalls disproportionately consequential. Conversely, inventory with a short shelf life may have a higher surplus cost than assumed here. Changing either side of the trade changes the quantile that the planner should target. The base case illustrates the effect of age under one stated trade-off, while the sensitivity grid makes the role of the penalty explicit.

An operational test would start with a timestamp audit. A system can record the time an order decision was produced, the effective time of its stock snapshot and the time the snapshot became visible. These fields distinguish elapsed physical demand from processing delay. An analyst can then reconstruct what the system knew before each historical decision. Using a later reconciled balance to generate the historical order would leak information unavailable at the time. That balance can still be valuable as a reference for evaluating the estimated inventory position. Historical demand may also need correction where stockouts censor observed sales.

Freshness is only one input to a reliable inventory position. A stock snapshot may omit purchase orders in transit, customer reservations or outstanding backorders. Even a fresh count can then disagree with the decision-relevant position. An implementation should document which movements are reflected in x and which are modeled separately. If unobserved receipts are possible, the simple subtraction of expected demand is insufficient. The decision should account for net movement and any dependence between receipts, delays and demand. That change preserves the idea of conditioning on age while altering the probability distribution used for ordering.

The study also suggests a way to interpret dashboard information. A planner benefits from seeing the observation time and an estimate of what happened since, along with a clear distinction between expected depletion and the quantity held for uncertainty. If a recommended order jumps when the count becomes older, the planner can determine whether the change reflects more expected consumption or a wider uncertainty range. That explanation is operationally useful when decisions must be approved or audited. It also helps identify whether an apparent issue originates in the demand model, the data feed or the replenishment rule.

5.1 How to evaluate a reporting intervention

A replenishment policy and an improvement in reporting speed intervene at different points. More frequent reporting changes the distribution of ages at which decisions are made. An age-aware order policy changes what happens conditional on the age that the planner actually sees. To compare them, an organization would need separate estimates of the cost of making a report fresher and the change in order outcomes that a fresher report would permit. The portfolio scenarios provide a way to organize that comparison. They reweight conditional outcomes without changing the demand rate or the stock observation in each hypothetical case. In live data, those characteristics may move together with age, so a simple reweighting based on age alone could misstate the benefit.

An empirical design could evaluate a reporting change first through a retrospective reconstruction. For each order, the analyst would identify the stock value visible at the decision time, the effective time of that observation, any known supply movements, and the subsequent demand or backlog over the defined horizon. The reconstructed baseline should reproduce the incumbent order or document the role of manual overrides. The age-aware policy would then be evaluated from the same information set. Separating forecast training, calibration and evaluation periods would prevent later sales or reconciliation records from determining the order that the policy would have recommended earlier.

Where a historical order affects what can be observed, reconstruction

alone is insufficient. Recorded sales can fall below demand during a stockout. If the replenishment policy changes availability, it may also change the measured series used to evaluate the policy. Backorder records, availability logs and data on substitutions can help distinguish those processes. A prospective comparison may be needed when the historical records cannot resolve them. A trial could assign policies by site or product group, account for dependencies between items, and predefine both service and cost outcomes. Those design choices would depend on how products share warehouse resources and how customers substitute between items.

Several diagnostics would help assess whether the distribution behind the ordering rule is credible. Conditional quantile coverage can be measured within bands of snapshot age, so that adequate performance on fresh observations does not hide undercoverage when reports are delayed. Observed demand variance can be compared with the Poisson assumption after accounting for changes in the mean. Event records can show whether reporting delays are more common during promotions or disruptions. Each of these diagnostics addresses an assumption used in the benchmark. None can be settled by the expected cost comparison itself, because that comparison takes the demand law and available observations as given.

The value of freshness can also differ by product even when an information system has uniform reporting delays. A product that sells slowly may experience little depletion over several days. A high-volume product can consume the same stock in a much shorter period. The delay should therefore be interpreted in units of exposure to demand, not simply in clock time. This suggests a testable distinction for operational planning. A system may need to prioritize improvements for records whose age is large relative to their replenishment horizon and demand rate. Ranking those improvements by estimated cost would require empirical distributions and infrastructure costs beyond the scope of the benchmark.

When timestamps themselves are uncertain, an analyst must distinguish a known reporting lag from a missing or questionable effective date. The model treats age as observed. A production system that silently assigns age zero to undated records would reproduce the central error in the baseline even if its ordering formula otherwise uses timestamps. Such records could be identified separately in the evaluation to understand their frequency and operational significance. A policy for handling them would need to reflect what other evidence is available, including audited counts or independently recorded inventory movements.

Finally, a product portfolio can have common causes of delay. An outage affecting several sites may make reporting ages increase together while the same disruption changes demand. The average expected cost can still be calculated from conditional cases, but the distribution of aggregate losses depends on those correlations. A planning team interested in the risk of unusually expensive days would need a joint model or a suitable resampling design. The current portfolio comparison addresses the mean effect of changing the age mix. It leaves the risk of simultaneous shortages across products and locations for a multi-item extension.

5.2 Interpreting cost in a pilot

The units of cost in this benchmark are relative. A holding cost of one fixes the scale, and a shortage cost of four sets the ratio that controls the ordering target. Multiplying both by the same positive constant rescales expected losses without changing the selected quantity. A pilot should estimate the ratio for its own setting. For a backlogged

item, shortage cost might include processing an additional shipment or compensating a customer. For a perishable item, the relevant surplus cost might include disposal or markdowns. The cost categories should be defined for the same period and applied consistently to the alternative policies. Otherwise the numerical comparison may look more precise than the underlying accounting permits.

A pilot should also separate changes in recommendation from changes in execution. A planner may override a larger order because of a supplier constraint or a known promotion missing from the demand inputs. That override is informative if the reason can be recorded and evaluated. A useful comparison would report how often the timestamp-aware policy changes the recommended quantity, how often the recommendation is adopted, and how adoption affects the two terminal measures. These quantities answer different questions about the decision rule and the workflow that carries it out. They also prevent a technically correct calculation from being credited with improvements that arose from a separate operational change.

6. LIMITATIONS

The experiment assumes Poisson demand with a constant known rate. Retail demand often varies by day of week, promotion, weather and season. Demand can also be overdispersed or intermittent, and reporting delays can coincide with the same events that change demand. An age-specific Poisson quantile may then be poorly calibrated. The principle can be extended by replacing the Poisson sum with a predictive distribution for net demand that conditions on observed context. Validation would need to test whether coverage remains adequate for old snapshots and whether predictive uncertainty grows appropriately with age.

The study uses a single cycle and terminal backlogging. It cannot establish how the policies affect multi-period inventory, stockout duration, immediate fill rate or customer abandonment. Repeated orders may overlap when reports remain stale for several decisions. If the ordering system does not track its own pending receipts, using an old snapshot repeatedly can lead to duplicate orders. A multi-period extension should track the order pipeline, expiry and replenishment arrival times. It should measure waiting and stock availability during the cycle as well as at the terminal date.

Finally, the benchmark assumes unrestricted ordering at a linear terminal loss. Suppliers may impose minimum quantities, pack sizes, budget ceilings or shared capacity constraints across products. In those cases, independent quantile targets may be infeasible or too costly when applied across the assortment. A constrained version would allocate shared capacity based on the expected marginal cost of another unit. The present study isolates the effect of age before such complications are introduced. Its exact results support a mechanism and a benchmark for comparison. Field effectiveness remains a question for operational data.

7. CONCLUSION

A stale inventory record changes both the expected stock position and the uncertainty surrounding it. In the exact benchmark, correcting expected depletion accounted for most of the cost reduction. Using an age-specific quantile delivered an additional improvement in expected cost and terminal service, especially in reporting environments with many old snapshots. This distinction helps explain why a simple adjustment can work well in one environment and leave meaningful service losses in another. The contribution is a decision model, a reproducible comparison and a clear specification of the conditions under which the results apply.

REFERENCES

[1] S. Sultana *et al.*, "A Comparative Review of Machine Learning Algorithms in Supermarket Sales Forecasting with Big Data," *J. Ecohumanism*, vol. 3, no. 8, pp. 14457 – 14467, Nov. 2024, https://doi.org/10.62754/joe.v3i8.6762.

[2] M. A. Goffer *et al.*, "Leveraging predictive analytics in management information systems to enhance supply chain resilience and mitigate economic disruptions," *Educ. Adm. Theory Pr.*, vol. 30, no. 4, pp. 11134–11144, 2024, doi: https://doi.org/10.67145/kuey.v30i4.9641.

[3] M. E. Hossin *et al.*, "Harnessing business analytics in management information systems to foster sustainable economic growth through smart manufacturing and Industry 4.0," *Educ. Adm. Theory Pr.*, vol. 30, no. 10, pp. 730–739, 2024, doi: https://doi.org/10.67145/kuey.v30i10.9643.

[4] U. Haldar *et al.*, "The Role of Artificial Intelligence in Enhancing Customer Experience in Business," *J. Comput. Anal. Appl.*, vol. 33, no. 08, pp. 7138–7155, Dec. 2024, [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4068

[5] F. B. Khair *et al.*, "Sustainable economic growth through data analytics: The impact of business analytics on U.S," *energy Mark. green Initiat. 2024 Int. Conf. Progress. Innov. Intell. Syst. Data Sci. (pp. 108–113). IEEE*, 2024, doi: https://doi.org/10.1109/ICPIDS65698.2024.00026.

[6] M. A. R. Chy *et al.*, "Understanding the relationship between data governance and business analytics success: A case study of global corporations," *J. Inf. Syst. Eng. Manag.*, vol. 9, no. 4s, pp. 4120–4136, 2024, doi: https://doi.org/10.52783/jisem.v9i4s.14807.

[7] P. Chakraborty *et al.*, "Toward autonomous decision intelligence: Integrating explainable AI and scalable DSS architectures in modern management information systems," *J. Inf. Syst. Eng. Manag.*, vol. 9, no. 4s, pp. 3893–3917, 2024, doi: https://doi.org/10.52783/jisem.v9i4s.14591.

[8] F. Mahmud *et al.*, "Big data and cloud computing in IT project management: A framework for enhancing performance and decision-making," *Fuel Cells Bull.*, vol. 2023, no. 9, pp. 1–18, 2024, doi: 10.52710/fcb.166.

[9] N. Das *et al.*, "Leveraging management information systems for agile project management in information technology: A comparative analysis of organizational success factors," *J. Bus. Manag. Stud.*, vol. 5, no. 3, pp. 161–168, 2023, doi: https://doi.org/10.32996/jbms.2023.5.3.17.

[10] K. B. Siddiqa *et al.*, "AI-Driven Project Management Systems: Enhancing IT Project Efficiency Through MIS Integration," in *2024 International Conference on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS)*, 2024, pp. 114–119. https://doi.org/10.1109/ICPIDS65698.2024.00027.

[11] M. Hossain, M. M. T. G. Manik, A. Tiwari, J. Ferdousmou, N. Vanu, and A. Debnath, "Data Analytics for Improving Employee Retention in the U.S. Technology Sector," in *2024 International Conference on Intelligent Cybernetics Technology & Applications (ICICyTA)*, IEEE, 2024, pp. 344–349. https://doi.org/10.1109/ICICYTA64807.2024.10913216.

[12] F. Karim *et al.*, "Empowering Women through Microcredit: An Analysis of the Grameen Bank Model's Impact on Socio-Economic Outcomes," *Int. J. Comput. Exp. Sci. Eng*, vol. 10, no. 4, pp. 2578–2589, 2024, doi: https://doi.org/10.22399/ijcesen.4074.

[13] M. Adnan *et al.*, "Interpretable Data Analytics for One-Year Diabetes Progression: Comparative Prediction, Feature Reliance, and Ranking Stability," *Eur. J. Med. Heal. Res.*, vol. 2, no. 6, pp. 257–269, 2024, doi: https://doi.org/10.59324/ejmhr.2024.2(6).34.

[14] E. Hossain *et al.*, "Cardiovascular Disease Burden and Clinical Associations Among US Adults Using Survey-Weighted Data Analytics," *Eur. J. Med. Heal. Res.*, vol. 2, no. 3, pp. 269–281, Jun. 2024, https://doi.org/10.59324/ejmhr.2024.2(3).31.

[15] M. J. Hossain *et al.*, "Data-Driven Predictive Modeling for Early Diagnosis and Risk Stratification of Chronic Diseases," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 2, no. 01, pp. 162–172, Aug. 2024, https://doi.org/10.58812/esiscs.v2i01.1220.

[16] M. H. Rahman *et al.*, "Harnessing big data and predictive analytics for early detection and cost optimization in cancer care," *J. Comput. Sci. Technol. Stud.*, vol. 6, no. 5, pp. 278–293, 2024, doi: https://doi.org/10.32996/jcsts.2024.6.5.22.

[17] M. I. Faruk *et al.*, "Global health intelligence platforms data science infrastructure for real-time epidemiological surveillance and pandemic forecasting," *Spectr. Eng. Sci.*, vol. 2, no. 5, pp. 749–761, Dec. 2024, [Online]. Available: https://thesesjournal.com/index.php/1/article/view/3425

[18] S. Hossain *et al.*, "Big Data Analysis and prediction of COVID-2019 Epidemic Using Machine Learning Models in Healthcare Sector," *J. Ecohumanism*, vol. 3, no. 8, pp. 14468 – 14477, Nov. 2024, https://doi.org/10.62754/joe.v3i8.6775.

[19] M. H. Rahman *et al.*, "Scalable AI Models for Climate Change Mitigation Using Multisource Geospatial Big Data," *J. Comput. Anal. Appl.*, vol. 33, no. 08, pp. 5836–5856, Dec. 2024, [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/3431

[20] M. M. T. G. Manik, "Multi-Omics Integration with Machine Learning for Early Detection of Ischemic Stroke Through Biomarkers Discovery," *J. Ecohumanism*, vol. 2, no. 2, pp. 175 – 187, Jun. 2023, https://doi.org/10.62754/joe.v2i2.6800.

[21] M. K. Ahmed *et al.*, "Big Data in Plant Biotechnology: Leveraging Bioinformatics to Discover Novel Anticancer Agents from Flora," *J. Med. Heal. Stud.*, vol. 4, no. 6, pp. 126–133, 2023, doi: https://doi.org/10.32996/jmhs.2023.4.6.15.

[22] S. I. Nilima *et al.*, "Advancement of Drug Discovery Using Artificial Intelligence and Machine Learning," in *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)*, 2024, pp. 1–7. https://doi.org/10.1109/COMPAS60761.2024.10796748.

[23] P. Chakraborty *et al.*, "An Intelligent Cybersecurity Framework for Data Protection in Cloud Computing Environments," *Spectr. Eng. Sci.*, vol. 2, no. 3, pp. 657–677, Oct. 2024, doi: https://doi.org/10.5281/zenodo.21538159.

[24] K. S. Mudaliyar and S. S. Kumar, "Intelligent Cybersecurity Frameworks for Data Protection in Cloud-Integrated Management Information Systems," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 2, no. 02, pp. 305–323, Dec. 2024, https://doi.org/10.58812/esiscs.v2i02.1165.

[25] M. D. Hossain, M. S. Sikder, M. S. Uddin, R. M. Ahsan, B. Uddin, and T. Hossen, "Cognitive Cyber Defense: AI–MIS Integration through Big Data and Cloud Frameworks for Next-Generation Digital Resilience," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 1, no. 02, pp. 140–152, Dec. 2023, https://doi.org/10.58812/esiscs.v1i02.764.

[26] S. N. Hasan *et al.*, "Enhancing cybersecurity threat detection and response through big data analytics in management information systems," *Fuel Cells Bull*, vol. 2023, no. 12, 2023, doi: https://doi.org/10.52710/fcb.137.

[27] J. Kaur *et al.*, "Advanced cyber threats and cybersecurity innovation: Strategic approaches and emerging solutions," *J. Comput. Sci. Technol. Stud.*, vol. 5, no. 3, pp. 112–121, 2023, https://doi.org/10.32996/jcsts.2023.5.3.9.

[28] S. M. Orthi *et al.*, "AI-Driven Cyber Threat Intelligence as a Management Information System: Integrating Cybersecurity Governance and IT Project Management for Organizational Resilience," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 1, no. 02, pp. 194–213, Dec. 2023, https://doi.org/10.58812/esiscs.v1i02.873.

[29] A. Shan-A-Alahi, M. Mustafizur, K. M. R. Hossan, A. Al Zaiem, and M. M. Rahman, "Cybersecurity training and its influence on employee behavior in business environments," *Comput. Fraud Secur.*, vol. 2024, no. 12, pp. 506–515, 2024, doi: https://doi.org/10.52710/cfs.689.

[30] M. K. Tuhin *et al.*, "Federated Machine Learning for Privacy-Preserving Cyber Threat Intelligence in Global Cloud-Integrated MIS Platforms," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 2, no. 02, pp. 284–293, Dec. 2024, https://doi.org/10.58812/esiscs.v2i02.1123.

[31] M. D. Hossain, M. S. Uddin, M. S. Sikder, T. Hossen, B. Uddin, and R. M. Ahsan, "Green and secure data centers: Balancing energy efficiency with advanced cybersecurity measures," *J. Comput. Sci. Technol. Stud.*, vol. 6, no. 5, pp. 300–315, 2024, doi: https://doi.org/10.32996/jcsts.2024.6.5.24.