The Eastasouth Journal of Information System and Computer Science
Vol. 4, No. 02, December 2026, pp. 19-35
ISSN: 3025-566X, DOI: 10.58812/esiscs.v4i02

From Risk Scores to Decision Value: An Evidence-Governed Business Analytics Framework Under Capacity Constraints

Celeste Yvonne Ramirez
BBW University of Applied Sciences, Berlin, Germany

Article Info

Article history:

Received Nov, 2026
Revised Nov, 2026
Accepted Nov, 2026

Keywords:

Business Analytics;
Constrained Optimization;
Data Governance;
Decision Intelligence;
Human Oversight;
Intervention Value

ABSTRACT

Organizations increasingly use machine learning to identify customers, transactions, suppliers, and projects that require attention. Yet a high predicted risk does not establish that an intervention will help, that its benefit exceeds its cost, or that the underlying evidence permits action. This paper develops an analytical framework that connects prediction, intervention value, evidence verification, and limited managerial capacity. The framework combines temporally governed data, calibrated risk estimates, bounded intervention effects, constrained optimization, and selective human review. Its central object is a decision record linking each proposed action to the information available at the decision time, the assumptions behind its expected value, and the authority required for execution. Analytical propositions establish when risk ranking agrees with value ranking, how uncertainty changes the intervention threshold, and when additional review has positive expected value. A fully specified six-case worked example demonstrates the distinction. With capacity for two interventions, risk ranking produces an assumed expected net value of −1 unit, whereas value-based selection produces 54 units. A common uncertainty allowance of 0.10 reduces the optimized lower-bound value to 34 units. These quantities are arithmetic consequences of stated assumptions, not estimates from an organizational experiment. The contribution is an auditable decision formulation and a prospective evaluation protocol that separates predictive performance from intervention effectiveness. The framework offers testable design principles for business analytics while identifying the causal, temporal, and organizational evidence required before deployment.

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

Corresponding Author:

Name: Celeste Yvonne Ramirez
Institution: BBW University of Applied Sciences, Berlin, Germany
Email: celesteramirez.analytics@gmail.com

1. INTRODUCTION

A business prediction becomes consequential when someone acts on it. A retention team contacts a customer, a procurement manager expedites an order, or an analyst sends a transaction for investigation. Each action consumes resources and can create unintended costs. The practical question is therefore broader than whether an event is likely: which feasible action is worth taking now, given the available evidence and the resources that remain?

Predictive systems often answer only the first part of that question. A churn model ranks customers by their probability of

leaving, but the customers most likely to leave may be the least responsive to an offer. A supplier model identifies a severe disruption, but an alternative route may already be unavailable. A project dashboard detects a late milestone after the window for effective intervention has closed. These examples illustrate a gap between estimating an outcome and selecting a useful response.

The supplied business-management literature places artificial intelligence within organizational decision processes rather than treating modeling as an isolated activity. [1] discuss implementation challenges alongside potential business applications. This perspective motivates a framework in which data access, action ownership, and operational constraints are explicit parts of the analytical problem.

Business intelligence also supports competing organizational objectives. The supplied review by [2] addresses decision-making and organizational ambidexterity. Its place in the research agenda is useful here: exploiting a reliable operational rule and exploring a new intervention requires different evidence and different tolerances for uncertainty. A common dashboard should not conceal that distinction.

This paper asks three questions. First, under what conditions can risk scores support economically coherent decisions when intervention capacity is limited? Second, how should incomplete evidence and uncertain intervention effects change the feasible action set? Third, what information must be preserved so that managers and researchers can assess whether a decision process creates value rather than merely predicting events?

The contribution is analytical. We define modular architecture, derive selection and review conditions, and supply a reproducible worked example. We also specify an empirical protocol capable of testing the framework prospectively. No new customers, transactions, suppliers, or project observations are claimed. The paper does not report on a deployed system, a field experiment, or a pooled estimate of effects across the cited literature.

The distinction matters because the available references span substantially different tasks and evidence designs. The result from telecom churn cannot establish procurement savings, and governance architecture cannot establish clinical benefits. The proposed integration therefore concerns the structure of decisions: information availability, expected consequences, feasible actions, and accountable execution. Domain-specific outcomes remain separate.

2. RESEARCH DESIGN AND EVIDENCE BOUNDARIES

2.1 Analytical design

The study combines a structured conceptual mapping of a supplied bibliography with original mathematical reasoning and transparent numerical illustration. The bibliography contains 45 entries across business analytics, finance, artificial intelligence, data engineering, supply chains, and management information systems. It is a purposive source set rather than a systematically sampled literature corpus. No claims about comprehensive coverage, publication trends, or the frequency of successful implementations follow from it.

The unit of analysis is a decision opportunity at a specified time. A decision opportunity contains an entity, the information currently observable about it, a set of permissible actions, and a deadline. An action can be selected only if it satisfies evidence and authorization requirements. Its expected value must then be assessed against alternative uses of capacity.

The analytical method has four stages. We identify design requirements across the source themes; distinguish verified empirical findings from bibliographic context; formulate a constrained decision problem; and examine the resulting rules through a worked example. The example is deliberately small enough for every calculation to be inspected without a

training pipeline. Its purpose is to reveal assumptions, not to approximate a particular industry.

2.2 *Evidence hierarchy*

Three publisher-accessible studies supply specific methodological and empirical anchors. Other supplied references provide thematic context where their detailed methods and findings have not been independently established for this manuscript. A bibliographic reference is not treated as evidence that a reported technique works. Numerical performance claims are restricted to findings inspected on the relevant publisher pages.

The distinction also applies to publication dates. The marketing architecture chapter carries a 2027 bibliographic year while its publisher records first online publication in September 2026. It is retained using the supplied volume year. The Fuel Cells Bulletin article has a 2024 publication date within an issue labeled 2023(9). Both labels are preserved rather than silently made consistent. These differences affect bibliographic interpretation, not the analytical calculations.

2.3 *Empirical anchors and transferable limits*

[3] report a leakage-conscious retail-banking churn pipeline with calibration, cost-sensitive thresholding, and explainability. The publisher reports a test ROC-AUC of 0.8668 for tuned LightGBM. The transferable lesson is to distinguish ranking, probability quality, and the operating threshold. A single-dataset predictive study does not establish the causal benefit of offering retention treatment.

[4] benchmark five tabular learners with nested cross-validation and explanation methods in telecom churn. Their reported multilayer-perceptron accuracy is 92.28%. This motivates separating model tuning from evaluation and making predictions inspectable. Neither that accuracy nor a feature explanation determines which customer will respond to a specific offer.

[5] combine structured fraud prediction with deterministic checking of language-model claims. Their reported D2 AUPRC improves from 0.5137 to 0.5824 relative to TabTransformer, while a separate post-shock comparison exceeds their stated degradation target. This supports examining verification and temporal failure together. It does not establish that verified rationales eliminate fraud or that the same improvement will occur in another business task.

Table 1. Verified empirical anchors and limits of transfer

Study	Inspected method or finding	Boundary for this paper
[3]	Bank churn: calibration, cost-sensitive thresholds; tuned LightGBM test ROC-AUC 0.8668.	Predictive performance does not identify retention treatment effects.
[4]	Telecom churn: nested cross-validation; reported MLP accuracy 92.28%.	Single-task accuracy is not a cross-domain decision-value estimate.
[5]	Fraud: structured claims and deterministic verification; D2 AUPRC 0.5137 to 0.5824.	Temporal stress remains consequential; verification is not universal safety.

3. DESIGN REQUIREMENTS ACROSS THE BUSINESS ANALYTICS LITERATURE

3.1 *Customer, market, and public value decisions*

The supplied customer-experience study identifies a relevant application domain for connecting analytical outputs with service decisions [6]. Within the present framework, experience becomes an outcome to measure, rather than a presumed consequence of deploying artificial intelligence. Complaints, offer acceptance, and customer burden should be recorded alongside retention.

Retail channel planning introduces structural change as a decision context [7]. A useful framework must permit demand relationships, action costs, and capacity to change over time. A channel allocation that was optimal under one cost structure may become unattractive even if its demand forecast remains accurate.

Pricing is another setting where predictive associations and intervention effects can diverge. The supplied refurbished-device pricing paper provides a task-specific reference point [8]. Historical price and sales data reflect earlier managerial choices. Estimating the effect of changing a price therefore requires assumptions beyond fitting observed sales.

The U.S. economic-growth framing in [9] extends the bibliography beyond individual operational tasks. The present model nevertheless measures value at a defined decision boundary. Firm-level expected gains cannot be interpreted as national growth without evidence about displacement, spillovers, and distribution.

Sustainable management introduces outcomes that may not be represented adequately by short-term financial return [10]. A business can improve a narrow margin measure while increasing waste or service exclusion. The proposed formulation allows such outcomes to enter explicit constraints, provided their measurement and policy rationale are specified.

The homelessness-cost reference provides a public-value application in which the consequences of an analytical decision are especially uneven [11]. Cost forecasts alone do not determine an appropriate allocation of support. Eligibility, unmet need, and service access must remain visible even when the decision model includes a monetary objective.

The broad healthcare and economic-competitiveness reference crosses several sectors [12]. Its bibliographic scope motivates a modular design, but cross-sector breadth is not evidence of transferability. Each application requires its own outcome definition, intervention mechanism, and validation population.

3.2 Finance and access to services

Credit-risk assessment is a natural example of distinguishing event prediction from action consequences [13]. Default risk may inform an assessment, but the effect of granting, declining, or restructuring credit is a separate question. An outcome model estimated among approved applicants also requires careful treatment before it is applied to applicants who would previously have been declined.

The supplied secure-financial-platform reference places model outputs within a platform setting [14]. In the proposed architecture, access control and separation of duties constrain who may view a score and who may execute an action. A technically correct prediction does not override the authority of the decision owner.

The microcredit reference highlights access and socioeconomic outcomes as distinct dimensions of financial analytics [15]. A selection rule that maximizes a narrow private return can distribute opportunities unevenly. Evaluation must therefore identify whose outcomes are observed and whose interests appear in the objective.

3.3 Data architecture, scale, and governance

The publisher abstract for the DataSphere chapter describes federated management, customer graphs, and marketing decision processes [16]. Its architectural focus helps motivate a shared evidence layer. The present framework does not adopt the chapter's reported performance gains as expected benefits of the proposed system.

Data governance and business analytics are connected explicitly in the supplied corporate-case reference [17]. Here, governance becomes operational through field-level availability times,

ownership, admissible uses, and versioned transformations. These requirements can be checked for a particular decision, rather than assessed only through an organization-wide maturity label.

High-velocity processing is represented in the stream-analytics reference [18]. The corresponding requirement is temporal consistency: an event may have occurred before a decision but arrived afterward. Reconstructing historical decisions requires both event time and the time at which information became available to the system.

Large repositories create a related computational requirement [19]. Expensive feature generation should be separated from the final decision step, and operational latency should be measured across the complete path. Fast model inference is insufficient if entity resolution or policy checks dominate response time.

The privacy-aware blockchain reference provides a governance theme [20]. A record of access can support accountability, but a ledger does not by itself prevent inappropriate collection or disclose whether consent permits a proposed use. The framework therefore specifies access and purpose checks independently of storage technology.

Federated learning and lakehouse design appear together in another supplied reference [21]. These are candidate infrastructure choices rather than mandatory ingredients. An organization should select them only when data distribution, confidentiality, and coordination requirements justify their additional complexity.

The DataOps reference motivates treating transformations and releases as managed operational objects [22]. The proposed decision record links an action to the feature version, model version, calibration version, and policy version in force at that time. Reproducibility requires preserving these relationships after systems change.

3.4 Adaptive models and bounded language-model roles

Hierarchical representation learning is a relevant model family for complex information [23]. Nevertheless, model complexity is not an objective of the framework. A simpler learner remains preferable when it produces comparable decision value with more stable estimation and lower operating cost.

Adaptive decision systems introduce the possibility of changing a model after deployment [24]. The associated requirement is controlled adaptation. Updated parameters and thresholds should first be evaluated against a frozen benchmark and then introduced with a traceable release decision.

Information noise is explicitly present in the managerial-robustness reference [25]. This motivates perturbation tests for missing, stale, duplicated, and inconsistent inputs. Such tests assess whether the action changes in a reasonable way, rather than merely whether the score changes by a small amount.

Human review is relevant to language-model reliability [26]. In the present design, review is a scarce intervention with its own cost and expected informational benefit. Sending every difficult case to a person can produce an unmanageable queue and shift errors into delayed decisions.

Multimodal agents expand the range of evidence that a decision system may process [27]. Each modality should retain its own provenance and availability time. A document summary must not silently replace the structured record from which a financial limit or eligibility condition is checked.

Human–machine collaboration raises a separate execution question [28]. A conversational interface can explain an available action, but permission to

recommend is distinct from permission to execute. The framework assigns these permissions explicitly.

Transformer-based summarization provides a useful communication function for large data flows [29]. A summary should reference the decision record and disclose missing evidence. Its fluency is not an acceptance criterion for the underlying decision.

3.5 Operations and management information systems

Supermarket forecasting provides a setting in which forecast errors have asymmetric consequences [30]. Overstock and stockout costs differ, while replenishment decisions are coupled across products. The empirical protocol therefore compares decisions under matched inventory or service capacity rather than relying solely on forecast accuracy.

Domestic timber supply chains introduce sector-specific constraints and economic-security concerns [31]. Such objectives require concrete operational definitions. A resilience claim should identify the disruption scenario, service requirements, and resource sacrifice that support it.

Procurement automation introduces contracts, approval limits, and supplier eligibility into the feasible action set [32]. These constraints should be checked before optimization. Selecting an economically attractive action that procurement policy prohibits does not produce an executable recommendation.

Predictive analytics and disruption management are linked in the supplied supply-chain MIS reference [33]. The present formulation adds a clear distinction between forecasting a disruption and estimating the benefit of a response. That distinction makes it possible to compare alternative mitigation actions.

Language models in operations motivate a structured interface between textual evidence and numerical planning [34]. Contract text can supply candidate facts, but a typed verifier should check quantities, dates, and constraints before those facts alter a decision.

The integrated supply-chain and trade reference broadens the operational setting [35]. Interdependence is consequential: rerouting one order can consume the transport capacity needed by another. The independent-case model developed below is therefore a starting formulation, with coupled constraints required for such applications.

End-to-end optimization introduces coordination across several decision stages [36]. Local score improvements can shift congestion downstream. Evaluation should measure service outcomes across the full process and retain a record of where capacity became binding.

Explainable decision-support architecture provides an MIS reference for connecting model output to organizational use [37]. In this paper, explanations identify the evidence and assumptions relevant to an action. They do not substitute for an estimate of the action's consequences.

Agile IT project management highlights changing responsibilities and short decision cycles [38]. The proposed system preserves ownership at the decision level so that a model alert can be accepted, deferred, or rejected by an identifiable role within the available response window.

The smart-manufacturing reference provides another MIS application with process dependencies [39]. Machine availability, workforce constraints, and quality requirements can make nominally valuable actions infeasible. These belong in the optimization model rather than in an informal implementation note.

The U.S. digital-transformation reference connects analytics with IT project outcomes [40]. Organizational improvement remains a hypothesis to test. The adoption of an analytical tool is

an exposure or process measure, not proof of project success.

Cloud-based MIS emphasizes collaboration and project governance [41]. A shared platform should support a common decision history while preserving role-specific access. Stakeholders need to distinguish a model recommendation, a manager's override, and an executed action.

The big-data and cloud-project framework similarly concerns analytical infrastructure [42]. Infrastructure can make evidence accessible, but the value of an intervention still depends on its timing and effectiveness. These components are separated in the proposed architecture.

Project-risk intelligence introduces a task where intervention outcomes can take months to mature [43]. Evaluation must allow sufficient follow-up and record changes in scope. Otherwise, an apparent reduction in risk can reflect a revised project definition rather than a better outcome.

MIS-integrated project management adds an execution context for predictive recommendations [44]. The framework accordingly distinguishes a recommendation from its acceptance and execution.

The supplied AI-powered business-intelligence reference reinforces the broader prediction-to-decision theme [45]. Its role here is contextual rather than quantitative validation of the integrated model.

Across these themes, the common requirement is a traceable connection between observable evidence, feasible action, expected consequence, and measured outcome.

4. PROPOSED FRAMEWORK

4.1 Decision records and temporal information

Let i identify a decision opportunity and t its decision time. Let I(i,t) contain only information that was available and permissible to use at t. An entity's record contains event timestamps, ingestion timestamps, source identifiers, and transformation versions. A feature is eligible only when its availability time is no later than t. Historical reconstruction therefore uses an as-of join, not the most recent version of a table.

This distinction prevents a familiar leakage mechanism. A payment reversal may refer to a transaction that occurred yesterday but become known only tomorrow. Using the reversal to score yesterday's transaction gives a model access to future information. Similarly, a project status field overwritten at closure cannot be used to reconstruct what a manager knew during an earlier escalation.

The decision record contains the opportunity identifier, allowed actions, predicted outcome risk, intervention-effect estimate, uncertainty specification, action cost, evidence-check status, capacity use, and policy version. After selection, it records the chosen action, execution status, any override, and the eventual outcome. The same identifier links offline evaluation with operational monitoring.

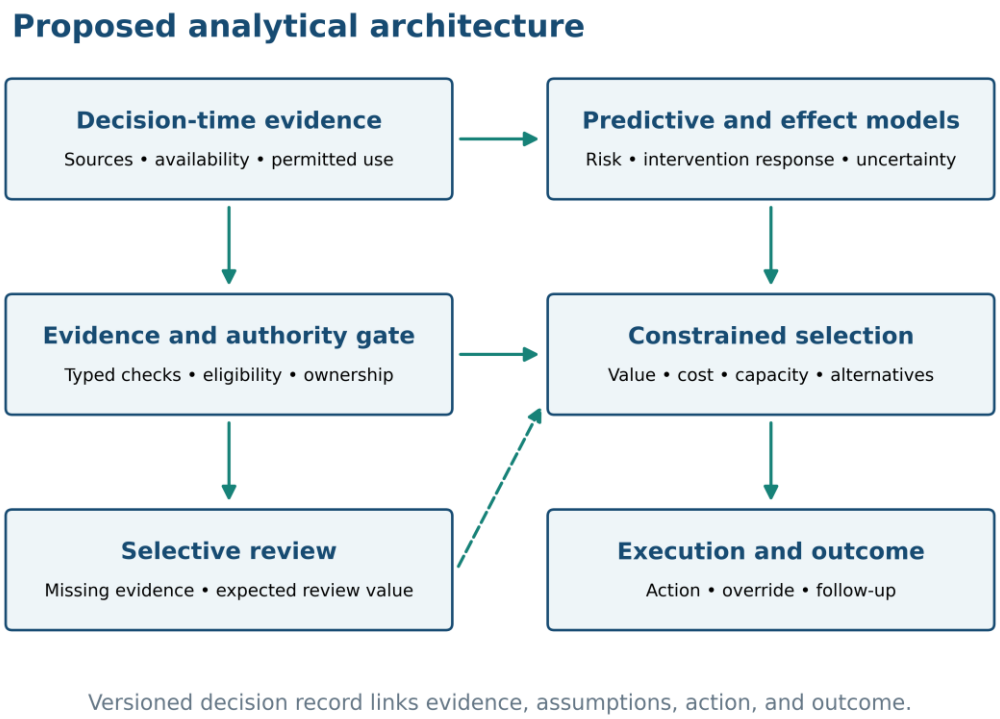


Figure 1. Proposed evidence-governed decision architecture. Solid arrows indicate routine flow; the dashed arrow indicates review-dependent resolution before selection. The diagram is conceptual, not a deployed system.

4.2 Risk and intervention effect

For a binary undesirable outcome, define $Y_i(0)$ as the outcome without an intervention and $Y_i(1)$ as the outcome under the specified intervention. Define $p_i = E[Y_i(0) \mid I(i,t)]$ and $\delta_i = E[Y_i(0) - Y_i(1) \mid I(i,t)]$. Risk p_i and effect δ_i answer different questions. The former describes expected harm under the baseline; the latter describes the expected reduction attributable to an action.

For a beneficial binary intervention effect, $0 \le \delta_i \le p_i$ is a useful feasibility restriction. Negative effects should be permitted in a more general model when an intervention can cause harm. If outcomes are continuous, the effect and value units must be redefined accordingly. An observed risk reduction between treated and untreated groups is not automatically an estimate of δ_i because treatment may have been assigned selectively.

Let v_i be the value of avoiding one undesirable outcome and c_i the full incremental intervention cost. The expected incremental net value is $u_i = v_i \delta_i - c_i$. Costs include delivery, incentives, processing, expected customer burden when monetized, and incremental operating expenses. Fixed platform costs are assessed separately when comparing adoption with continued use of an existing process.

When causal effect estimates are unavailable, the model can still support scenario analysis, but its result is conditional on assumed effects. A manager may explore break-even values without claiming that a campaign will achieve them. This paper's numerical illustration uses exactly that interpretation.

4.3 Evidence gates and constrained selection

Let z_i be a binary selection variable, g_i a binary evidence-and-authorization gate, b_i the resource consumption of action i, B the available budget, and K the maximum number of actions. The nominal selection problem is:

Maximize $\Sigma_i z_i(v_i \delta_i - c_i)$, subject to $\Sigma_i z_i \leq K$, $\Sigma_i b_i z_i \leq B$, and $z_i \leq g_i$ for every i. (1)

The gate equals one only when the required evidence is present, the proposed use is permitted, and the execution authority is valid. A failed gate does not imply that the entity has low risk. It means that the proposed action is currently outside the feasible set. The system can instead request missing evidence or offer a separately authorized fallback.

With equal resource consumption and a count constraint alone, the solution selects up to K positive-value opportunities in descending u_i. With unequal resource use, simple ranking by value or value-to-cost ratio is not generally exact for binary decisions. An integer optimizer is required unless the problem has additional structure. Shared supplier or household constraints can also couple cases that otherwise appear independent.

4.4 Uncertainty and conservative action

Suppose δ_i belongs to a prespecified interval $[l_i, h_i]$. When v_i is nonnegative, the worst-case incremental value over that interval is $u_i^L = v_i l_i - c_i$. Replacing u_i with u_i^L in Equation 1 gives a robust selection problem under a rectangular uncertainty set. Rectangularity means that every combination of within-interval effects is admitted; it may be conservative when errors are strongly dependent.

The intervals must have an explicit interpretation. They may be sensitivity ranges or statistical intervals obtained using a documented procedure. A sensitivity range has no automatic confidence level. Even a statistically valid individual interval does not necessarily provide simultaneous protection for a selected portfolio. The present example uses assumed sensitivity ranges only.

A useful boundary follows immediately: an individually attractive conservative action requires $l_i > c_i/v_i$ when $v_i > 0$. Capacity introduces an additional opportunity cost. A case can pass its own break-even threshold and still be excluded because another eligible case uses the same resource more productively.

4.5 Review and explanation

Human review can alter the information state or resolve an evidence gate. Let V_0 be the optimal feasible portfolio value before review. Let V_1 be the random optimal value after the review result becomes available. If review costs r and has no additional delay cost, its expected net value is $E[V_1 - V_0] - r$. Review is attractive when this quantity is positive and review capacity is available.

This expression compares complete portfolios. It avoids counting an action's gross value while ignoring the action it displaces. If review consumes intervention capacity or delays the response, those effects must be included. Review is not assumed to be correct; its result can confirm, revise, or reject an action.

A language model may translate an accepted decision record into a readable explanation. Each factual claim should point to a typed field or an approved calculation. Verification checks units, entity identity, dates, arithmetic, and the relationship between a claim and its cited evidence. Statements about causation or future benefit must retain the assumptions under which they were calculated.

5. ANALYTICAL RESULTS

5.1 Proposition 1: Conditions for risk ranking

Suppose all cases have the same avoided-outcome value $v > 0$, the same intervention cost c, identical resource use, and an intervention effect $\delta_i = s\, p_i$ for a common $s > 0$. Then $u_i = vs\, p_i - c$ is strictly increasing in p_i. Subject to common gates and a count constraint, ranking by risk and ranking by expected

value agree among cases with positive value.

The proof follows from comparing any two cases: $u_i - u_j = vs(p_i - p_j)$. Its sign equals the sign of the risk difference. This is a sufficient condition rather than a necessary one. If v, c, or responsiveness varies, agreement can still occur for a particular portfolio, but it is no longer guaranteed by risk ordering.

The condition is demanding in practice. High-risk customers may be unresponsive; lower-risk customers may be inexpensive to help; some cases may require substantially more staff time. Consequently, the empirical question is whether the risk ordering preserves the value ordering under the relevant action and cost structure.

5.2 Proposition 2: Uncertainty changes break-even decisions

For nonnegative v_i and an interval effect assumption, worst-case utility equals $v_i\, l_i - c_i$. This follows because utility is a nondecreasing affine function of δ_i, so its minimum occurs at the lower endpoint. Under equal resource use and a count constraint, sorting positive lower-bound utilities is therefore optimal for the stated robust problem.

An uncertainty allowance can change which cases are selected and how many actions should be taken. An optimizer with a capacity ceiling is not required to fill every slot. Once every remaining conservative value is nonpositive, further selection does not improve its objective. This result depends on the model's cost accounting; a separately imposed service obligation may require an action even when private financial utility is negative.

5.3 Proposition 3: Review value depends on displaced alternatives

Suppose one failed-gate case can become eligible with probability q after review. If successful review changes the portfolio value from V_0 to $V_0 + D$, unsuccessful review leaves it unchanged, and review costs r, then expected net review value equals $qD - r$. For $D > 0$, review is worthwhile when $q > r/D$.

This result is obtained by conditioning on the two possible review outcomes. It assumes that q and D are appropriate ex ante quantities and that no delay, capacity, or execution changes occur during review. If other cases arrive while the review is pending, a dynamic optimization is required. The expression is therefore a transparent benchmark rather than a universal queueing rule.

5.4 Worked example: six intervention opportunities

Consider six hypothetical opportunities labeled A through F. Each uses one intervention slot and costs 8 value units. Avoiding an undesirable outcome is worth 100 units for every case. Baseline risks are 0.90, 0.80, 0.70, 0.60, 0.50, and 0.40. Assumed absolute effect reductions are 0.05, 0.10, 0.30, 0.40, 0.25, and 0.10. All effects satisfy the binary feasibility restriction relative to baseline risk.

With K = 2, a highest-risk rule selects A and B. Their assumed net values are −3 and 2, yielding −1. Value-based selection chooses C and D, whose values are 22 and 32, yielding 54. The difference is 55 units. These are expected incremental values under the specified scenario; they are not realized profits, estimated causal effects, or measured gains from a trained model.

Table 2. Fully specified hypothetical intervention opportunities

Case	Baseline risk	Assumed effect	Net value	Lower-bound value
A	0.90	0.05	-3	-8
B	0.80	0.10	2	-8
C	0.70	0.30	22	12
D	0.60	0.40	32	22
E	0.50	0.25	17	7
F	0.40	0.10	2	-8

Note. Avoided-outcome value = 100 units; intervention cost = 8 units; one slot per action; maximum two actions. Lower-bound values use an assumed absolute effect allowance of 0.10. All quantities are illustrative inputs or derived arithmetic.

Now impose a common uncertainty allowance $\varepsilon = 0.10$ and define $l_i = \max(0, \delta_i - \varepsilon)$. The conservative utilities are −8, −8, 12, 22, 7, and −8. C and D remain selected, giving a portfolio lower-bound value of 34. Under the unchanged central assumptions, the same selected portfolio still has nominal value 54. A conservative objective and a nominal evaluation of the selected actions are different quantities.

If D fails its gate, the nominal feasible optimum becomes C and E, with value 39. The conservative value under $\varepsilon = 0.10$ becomes 19. The missing eligibility reduces both optima by 15 units relative to the corresponding fully eligible portfolio. This arithmetic does not imply that bypassing the gate is justified; it quantifies the opportunity associated with resolving the missing evidence.

Suppose a review of D costs 3 units and restores eligibility with probability q. The expected nominal improvement is $15q - 3$, so the break-even probability is 0.20. At $q = 0.60$, expected net review benefit is 6 units. If review fails, C and E remain the selected fallback. If review succeeds, C and D become feasible. The calculation assumes the same action capacity and no deadline penalty in both branches.

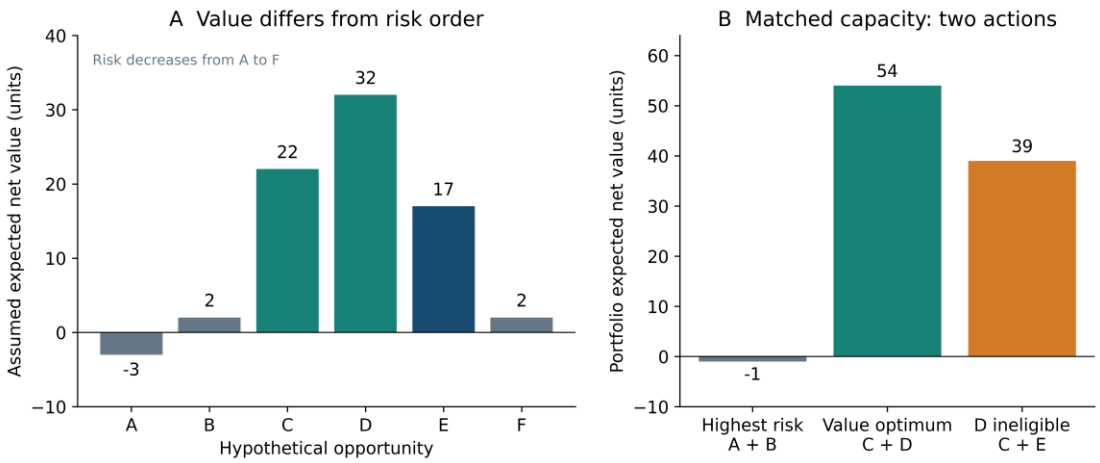


Figure 2. Hypothetical opportunity and portfolio values. A: individual assumed utility in descending baseline-risk order. B: nominal portfolio values under matched capacity. Values derive from Table 2 and do not represent empirical performance.

5.5 Sensitivity to uncertainty and capacity

As ε increases from 0 to 0.20, the conservative values of C and D remain positive, and the optimized total falls from 54 to 14. At $\varepsilon = 0.25$, only D has positive conservative value, yielding 7.

At $\varepsilon = 0.30$, the optimum is 2. By $\varepsilon = 0.35$, selecting no action yields the maximum conservative value of zero. The optimizer is allowed to leave capacity unused throughout.

With zero uncertainty and all gates open, increasing K from one through six produces optimized nominal values of 32, 54, 71, 73, 75, and 75. Capacity has diminishing marginal value in this example, and the sixth slot adds nothing because A has negative utility. This curve offers a concrete way to discuss staffing: a larger queue-clearing target is not automatically a better economic objective.

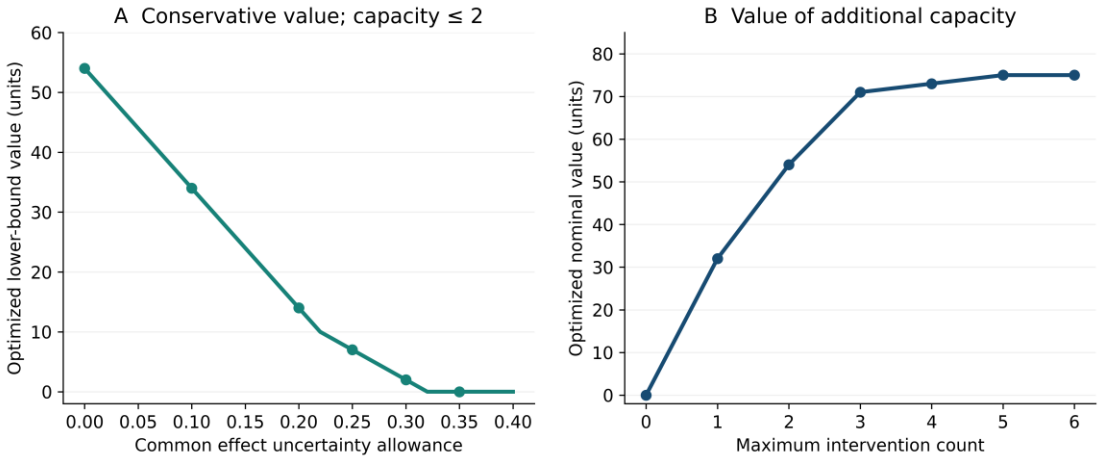


Figure 3. Analytical sensitivity of the worked example. A: optimized worst-case value under a common absolute effect allowance, with at most two actions. B: optimized nominal value as the intervention-count limit changes. Unused capacity is permitted.

The example establishes possibility, not prevalence. It shows that plausible heterogeneity can reverse a risk-based allocation and that uncertainty can make additional action unattractive. It does not show how frequently these conditions occur in banking, retail, or procurement. Answering that question requires observations and a credible intervention-effect design.

6. PROSPECTIVE EMPIRICAL EVALUATION

6.1 Data and identification

A first empirical study should select one operational domain and a well-defined intervention. Suitable records include decision-time features, eligibility, assigned action, execution, treatment cost, outcome horizon, and competing interventions. Repeated observations of the same customer, supplier, or project require entity-aware partitions. Data availability, rather than the timestamp of the eventual outcome alone, determines which features are admissible.

Randomized intervention assignments within an ethically and operationally eligible population provides a clear starting point for estimating average response differences. Heterogeneous effects require additional estimation care and honest evaluation. If only observational data are available, the protocol must state the assumptions used to address confounding, overlapping, and consistency. Cases without adequate treatment overlap should be identified rather than assigned confident effect estimates through extrapolation.

The outcome window should match the decision's purpose. An offer may delay churn without improving longer-term retention; an expedited order may protect one delivery while worsening another. Multiple horizons can expose these differences. Outcome ascertainment should remain comparable across action groups so that additional review does not merely increase the probability that an adverse event is documented.

6.2 Temporal validation and comparators

Training, calibration, policy selection, and final evaluation should use separate temporal windows. Hyperparameters and action rules are chosen before the test period is examined. The final period is locked. A second organization or operating region provides a stronger transfer test than additional random splits of the original source. Performance under a disruption period should be reported separately from routine operation.

Comparators should include the current organizational rule, highest-risk selection, estimated-value selection, and conservative selection. All policies receive the same admissible information, action budget, and outcome horizon. An additional ablation can remove the evidence gate in an offline analysis, but it must not execute prohibited actions. Another can replace verified explanations with unverified summaries to measure factual support without altering treatment assignment.

Prospective randomization of eligible policies can estimate differences in outcomes directly when feasible. Offline policy evaluation requires documented action propensities and sufficient overlap. Neither a high classification score nor a favorable retrospective simulation provides a substitute for identifying the value of the intervention policy.

6.3 Outcomes and uncertainty reporting

Predictive outcomes include discrimination, probability error, and calibration by time and relevant population group. Decision outcomes include incremental net value, missed benefit, intervention burden, service delay, and resource consumption. Evidence outcomes include missing fields, verification failures, unsupported claims, and execution mismatches. These dimensions should be reported together, because an improvement in one may conceal deterioration in another.

Statistical intervals must respect the unit of assignment and dependence structure. Repeated customer observations cannot be treated as independent when the intervention is assigned at customer level. Site-level deployment may require site-level analysis. Small numbers of temporal blocks create uncertainty that should be disclosed rather than hidden through repeated resampling of the same records.

Table 3. Prospective evaluation plan

Dimension	Primary assessment	Design safeguard
Prediction	ROC-AUC, PR-AUC, Brier score, calibration	Locked future period; entity-aware partitions
Intervention	Incremental outcomes and net value	Randomized eligible allocation or justified causal identification
Capacity	Value at fixed budgets; queue delay	Matched information and resource constraints
Evidence	Unsupported claims; gate failures	Typed verification; separate evidence audit
Human use	Comprehension; decision time; override quality	Task-based evaluation; record review costs
Robustness	Temporal and subgroup performance	Separate shock periods; dependence-aware intervals

6.4 Human and operational evaluation

A useful explanation should help a decision owner understand the recommended action, its assumptions, and its alternatives. A study can measure decision time, comprehension, factual error detection, override quality, and perceived workload. User confidence alone is inadequate: persuasive explanations can increase confidence without improving correctness.

The operational trial should record the full latency distribution, queue length, review abandonment, and time to

execution. A low average response time can coexist with severe tail delays. Resource use includes feature computation, evidence retrieval, optimization, and explanation verification. The economic evaluation should also include maintenance and governance costs when comparing deployment strategies.

Overrides deserve specific analysis. An override can reveal missing local information, an incorrect model assumption, or an organizational incentive conflict. Logging the reason allows investigators to distinguish these mechanisms. Override frequency by itself is ambiguous; fewer overrides can indicate better alignment or excessive deference to the system.

7. DISCUSSION

7.1 What integration contributes

The framework makes four relationships explicit. Risk describes an expected event under a baseline. Intervention effect describes the change attributable to a specified action. Governance defines which actions can be considered. Capacity determines which eligible, attractive actions can be executed together. Collapsing these relationships into a single score makes assumptions difficult to inspect and failures difficult to diagnose.

The proposed decision record is the practical connection among these components. It permits a reviewer to ask what was known, what action was available, why the action appeared useful, and what actually happened. This supports investigation after a poor outcome without assuming that every poor outcome implies a poor decision. Decisions are made under uncertainty; their quality depends on the information and alternatives available at the time.

The analytical propositions also clarify the place of prediction improvements. Better discrimination can help when it improves the ordering of valuable opportunities. Better calibration can improve cost calculations when probabilities enter the effect or loss model. Yet neither guarantees a better policy if action costs, responsiveness, or constraints are misspecified. Predictive progress has to be translated into the decision objective being tested.

7.2 Implications for managers

Managers should begin with the action rather than the algorithm. A useful specification names who can act, what the action changes, what it costs, when it must occur, and how success will be measured. This creates a tractable analytical question and reveals where data are missing. A vague ambition to automate decision-making provides little basis for evaluating value.

Capacity should be treated as an economic choice as well as a technical limit. The marginal value of an additional intervention slot can inform staffing, but only under the assumed effect and cost model. Review capacity has a different marginal value because its purpose is to improve information or resolve eligibility. Conflating these resources can create a system that identifies many opportunities but cannot execute them in time.

An evidence gate also changes how failure is communicated. A blocked recommendation should identify the missing requirement and the responsible role. The user interface should make the next permissible step clear, such as obtaining a current contract limit or correcting an entity match. An unexplained rejection encourages workarounds and prevents useful learning about recurring data problems.

7.3 Limits and external validity

The paper offers an analytical construction and a worked example, not empirical validation of an integrated platform. Its illustrative effects and costs are assumed. There is no evidence here that the proposed policy outperforms an existing organization, generalizes across sectors, or satisfies any particular

regulatory requirement. The proposed evaluation protocol is intended to produce that missing evidence.

The source set is purposive and heterogeneous. Detailed empirical use is restricted to inspected publisher material, while the broader bibliography is used to map application themes. No quality equivalence is assumed among publication venues, and no aggregate effect is calculated. A targeted review with fuller access to all source texts would be needed before making systematic claims about the state of the literature.

The optimization model also simplifies interactions. Customer contacts can influence one another, supplier actions can compete for shared logistics, and interventions can alter future risk. Static formulation does not capture these feedback effects. Multi-period models require state transitions, a discounting convention, and an explicit treatment of uncertainty in future actions.

Finally, a monetary objective cannot settle every organizational or social question. Some service obligations and protections should be constraints rather than tradable terms. Defining them requires domain expertise and accountable governance. The framework makes these choices visible but does not determine their legitimacy from data alone.

8. CONCLUSION

Business analytics create value through choices made under constraints, not through prediction scores alone. This paper formalizes the connection between decision-time evidence, intervention effects, action costs, eligibility, and scarce capacity. Its analytical results identify conditions under which risk ranking is coherent, establish a conservative break-even rule, and derive a portfolio-based criterion for human review.

The worked example demonstrates why these distinctions matter: the highest-risk cases need not offer the greatest expected benefit, and additional capacity need not justify additional action. The next research step is a prospective, domain-specific study that estimates intervention effects credibly and compares policies under matched operational conditions. Such evidence would establish whether the framework improves outcomes in practice and where its assumptions require revision.

REFERENCES

[1] P. Chakraborty *et al.*, "Leveraging Artificial Intelligence and Machine Learning for Decision-Making in Business Management: A Comprehensive Analysis," *J. Manage.*, vol. 2, pp. 46–56, 2025, doi: https://doi.org/10.53935/jomw.v2024i4.867.

[2] N. Das *et al.*, "The strategic impact of business intelligence tools: A review of decision-making and ambidexterity," *Membr. Technol.*, vol. 2025, no. 1, pp. 542–553, 2025, doi: https://doi.org/10.52710/mt.307.

[3] F. Karim *et al.*, "Calibrated, cost-sensitive, explainable ensemble learning for customer churn prediction in retail banking: A leakage-safe pipeline and single-dataset case study," *Front. Artif. Intell.*, vol. 9, 2026, doi: https://doi.org/10.3389/frai.2026.1889639.

[4] I. Zerine *et al.*, "Explainable churn prediction in telecom with tabular ML five model benchmark and SHAP analysis," *Discov. Artif. Intell.*, vol. 6, no. 1, 2026, doi: https://doi.org/10.1007/s44163-026-00983-0.

[5] M. S. A. Sourav *et al.*, "Grounding language models with deterministic verifiers for fraud detection," *Discov. Artif. Intell.*, vol. 6, no. 1, 2026, doi: https://doi.org/10.1007/s44163-026-02355-0.

[6] U. Haldar *et al.*, "The role of artificial intelligence in enhancing customer experience in business," *J. Comput. Anal. Appl.*, vol. 33, no. 8, pp. 7138–7155, 2024, doi: https://eudoxuspress.com/index.php/pub/article/view/4068.

[7] M. J. Hossain, S. S. Utsab, S. A. Jarir, E. Rozario, Z. A. Tanni, and T. Roy, "Business analytics for United States retail channel planning under structural change," *Eastasouth J. Inf. Syst. Comput. Sci.*, vol. 3, no. 1, pp. 211–223, 2025, doi: https://doi.org/10.58812/esiscs.v3i01.1263.

[8] N. Prajapati, "Leveraging artificial intelligence to optimize pricing strategies in the refurbished handheld device industry: An engineering perspective," *2025 International Conference on Information, Implementation, and Innovation in Technology (I2ITCON)*. IEEE, pp. 1–7, 2025. doi: https://doi.org/10.1109/I2ITCON65200.2025.11210682.

[9] U. Haldar *et al.*, "AI-Driven Business Analytics for Economic Growth Leveraging Machine Learning and MIS for Data-Driven Decision-Making in the U.S. Economy," *J. Posthumanism*, vol. 5, no. 4 SE-, pp. 932–957, Apr. 2025, doi:

https://doi.org/10.63332/joph.v5i4.1178.

[10] S. Hossain *et al.*, "From Data to Value: Leveraging Business Analytics for Sustainable Management Practices," *J. Posthumanism*, vol. 5, no. 5 SE-, pp. 82–105, Apr. 2025, doi: https://doi.org/10.63332/joph.v5i5.1309.

[11] M. M. R. Bhuiyan *et al.*, "Economic implications of homelessness in the U.S. applying advanced business analytics to forecast costs and develop sustainable solutions," *2025 5th International Conference on Electrical, Computer and Energy Technologies (ICECET)*. IEEE, pp. 1–5, 2025. doi: https://doi.org/10.1109/ICECET63943.2025.11472202.

[12] M. H. Rahman, M. M. Anwar, and F. Hossain, "AI-driven big data and business analytics: Advancing healthcare, precision medicine, supply chain resilience, energy innovation and economic competitiveness," *J. Med. Heal. Stud.*, vol. 6, no. 3, pp. 205–215, 2025, doi: https://doi.org/10.32996/jmhs.2025.6.3.30.

[13] M. M. T. G. Manik, A. S. M. Saimon, M. S. Islam, M. Moniruzzaman, E. Rozario, and E. Hossin, *Big Data Analytics for Credit Risk Assessment*. 2025. doi: https://doi.org/10.1109/ICMLAS64557.2025.10967667.

[14] H. Imam, M. Mustafizur, T. Roy, A. Lalafaryan, A. Mkrtchyan, and V. Roy, "Intelligent and secure financial analytics platforms using advanced machine learning," *2026 IEEE 18th International Conference on Computational Intelligence and Communication Networks (CICN)*. IEEE, pp. 769–775, 2026. doi: https://doi.org/10.1109/CICN70047.2026.11594305.

[15] 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, 2024, doi: https://doi.org/10.22399/ijcesen.4074.

[16] D. Alam, S. Saha, M. I. Alam, N. Sayeed, U. S. Saha, and M. Mustafizur, "Data-centric architectures for digital business and marketing technologies," in *Innovative Computing and Communications*, vol. 2046, Springer, 2027, pp. 546–561. doi: https://doi.org/10.1007/978-3-032-30438-4_44.

[17] 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, 2024, doi: https://doi.org/10.52783/jisem.v9i4s.14807.

[18] M. Adnan, P. A. Momotaj, T. Roy, R. Anzum, Z. A. Tanni, and M. M. T. G. Manik, "Stream data analytics for high-velocity data processing applications," *2026 6th International Conference on Intelligent Technologies (CONIT)*. IEEE, pp. 1–9, 2026. doi: https://doi.org/10.1109/CONIT69683.2026.11621735.

[19] M. M. T. G. Manik, "Scalable data mining algorithms for large-scale data repositories," *2026 6th International Conference on Intelligent Technologies (CONIT)*. IEEE, pp. 1–9, 2026. doi: https://doi.org/10.1109/CONIT69683.2026.11621696.

[20] S. Bauskar, R. K. Sahoo, S. S. Boda, H. Singhai, M. M. Bakhsh, and M. Adnan, "Privacy-aware big data governance framework using blockchain," *2025 IEEE International Conference on Emerging Trends in Computing and Communication (ETCOM)*. IEEE, pp. 1–9, 2025. doi: https://doi.org/10.1109/ETCOM66606.2025.11436976.

[21] P. Chakraborty *et al.*, "Trustworthy data lakehouse design using federated learning and blockchain," in *In 2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT) (pp. 1–8). IEEE*, 2025, pp. 1–8. doi: https://doi.org/10.1109/ICAFT66710.2025.11453041.

[22] S. M. Orthi *et al.*, "DataOps-oriented big data governance for automated decision pipelines," in *2025 1st International Conference on Advancement in Futuristic Technologies (ICAFT)*, 2025, pp. 1–8. doi: https://doi.org/10.1109/ICAFT66710.2025.11452860.

[23] R. Anzum, R. Gentyala, N. Mukhi, N. Karri, P. Krishnan, and E. Rozario, "Hierarchical learning architectures for complex data representation," *2026 International Conference on Emerging in Computing Technologies (INCECT)*. IEEE, pp. 1–8, 2026. doi: https://doi.org/10.1109/INCECT72370.2026.11742577.

[24] M. J. Hossain, M. M. Bakhsh, M. A. Sami, M. I. Alam, M. M. H. Melon, and M. M. T. G. Manik, "Self-adaptive artificial intelligence systems for large-scale data-driven decision making," *2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON)*. IEEE, pp. 1–8, 2026. doi: https://doi.org/10.1109/I3CTCON68242.2026.11507252.

[25] T. Roy, M. H. Rijvi, M. M. Bakhsh, M. Adnan, M. Nazmussakib, and E. Rozario, "Decision robustness under information noise: A machine learning perspective on managerial decision systems," *2026 IEEE 18th International Conference on Computational Intelligence and Communication Networks (CICN)*. IEEE, pp. 327–333, 2026. doi: https://doi.org/10.1109/CICN70047.2026.11594307.

[26] R. Dandigam, M. Patel, C. S. Pareek, S. M. Faysal, G. T. Alam, and M. A. Siam, "Human-in-the-loop learning for improving language model reliability," *2026 6th International Conference on Intelligent Technologies (CONIT)*. IEEE, pp. 1–8, 2026. doi: https://doi.org/10.1109/CONIT69683.2026.11620666.

[27] M. E. Hossin *et al.*, "Adaptive multi-modal AI agents: Leveraging large language models for dynamic real-world environments," in *Proceedings of International Conference on Computational Technologies for Research in Data Analytics*, vol. 1968, Springer, 2026, pp. 106–119. doi: https://doi.org/10.1007/978-3-032-25694-2_10.

[28] F. B. Khair *et al.*, "Bridging natural language and autonomous action: LLM-driven AI agents for human–machine collaboration," in *Proceedings of International Conference on Computational Technologies for Research in Data Analytics*, vol. 1968, Springer, 2026, pp. 92–105. doi: https://doi.org/10.1007/978-3-032-25694-2_9.

[29] E. Rozario, M. Nazmussakib, G. T. Alam, T. Roy, M. Adnan, and M. I. Faruk, "Dynamic big data summarisation using transformer-based generative models," *2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON)*. IEEE, pp. 1–8, 2026. doi: https://doi.org/10.1109/I3CTCON68242.2026.11507387.

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

[31] G. T. Alam *et al.*, "AI-Driven Optimization of Domestic Timber Supply Chains to Enhance U.S. Economic Security," *J. Posthumanism*, vol. 5, no. 1 SE-, pp. 1581–1605, Jan. 2025, doi: https://doi.org/10.63332/joph.v5i1.2083.

[32] M. Bashir, R. A. M. Rashed, M. Mustafizur, U. S. Saha, S. M. M. Uddin, and H. Imam, "Digital procurement transformation through intelligent automation," *2026 International Conference on Computing Theory and Wireless Communications (ICCTWC)*. IEEE, pp. 1–7, 2026. doi: https://doi.org/10.1109/ICCTWC68241.2026.11583410.

[33] M. A. Goffer, "Leveraging Predictive Analytics In Management Information Systems To Enhance Supply Chain Resilience And Mitigate Economic Disruptions," *Educ. Adm. Theory Pract.*, vol. 30, no. 4, pp. 11134–11144, 2024, doi: https://doi.org/10.53555/kuey.v30i4.9641.

[34] F. Karim, M. Nazmussakib, R. M. Haque, U. S. Saha, and M. M. Bakhsh, "Data-driven operations and supply chain management using large language models," *2026 IEEE 18th International Conference on Computational Intelligence and Communication Networks (CICN)*. IEEE, pp. 321–326, 2026. doi: https://doi.org/10.1109/CICN70047.2026.11594225.

[35] M. A. Miah *et al.*, "An integrated big data and predictive analytics framework for supply chain resilience, international trade and economic growth," *J. Intell. Decis. Mak. Inf. Sci.*, vol. 3, no. 2, pp. 1676–1694, 2026, doi: https://doi.org/10.59543/jidmis.v3.2223.

[36] N. Sayeed, U. S. Saha, F. Karim, M. Mustafizur, S. M. M. Uddin, and M. A. R. Chy, "Artificial intelligence models for end-to-end supply chain optimization," *2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON)*. IEEE, pp. 1–8, 2026. doi: https://doi.org/10.1109/I3CTCON68242.2026.11507407.

[37] 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.

[38] 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, p. 161, 2023, doi: https://doi.org/10.32996/jbms.2023.5.3.17.

[39] 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 Pract.*, vol. 30, no. 10, pp. 730–739, 2024, doi: https://doi.org/10.53555/kuey.v30i10.9643.

[40] M. E. Hossin *et al.*, "Digital Transformation in the USA Leveraging AI and Business Analytics for IT Project Success in the Post-Pandemic Era," *J. Posthumanism*, vol. 5, no. 4 SE-, pp. 958–976, Apr. 2025, doi: https://doi.org/10.63332/joph.v5i4.1180.

[41] F. Mahmud, M. A. Goffer, H. Rahman, and G. T. Alam, "The Role of Cloud-Based Management Information Systems in Enhancing IT Project Governance and Stakeholder Collaboration," 2025, doi: https://doi.org/10.1007/978-3-032-05548-4_1.

[42] 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: https://doi.org/10.52710/fcb.166.

[43] S. M. Orthi *et al.*, "AI-augmented risk intelligence in IT project management: An empirical MIS-driven evaluation using machine learning and neural networks," *Int. J. Appl. Math.*, vol. 38, no. 12s, pp. 2869–2890, 2025, doi: https://doi.org/10.12732/ijam.v38i12s.1594.

[44] 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. doi: https://doi.org/10.1109/ICPIDS65698.2024.00027.

[45] S. M. M. Uddin *et al.*, "AI-powered business intelligence: Enhancing decision-making through predictive analytics and big data," *J. Posthumanism*, vol. 6, no. 3, pp. 218–237, 2026, doi: https://doi.org/10.63332/joph.v6i3.4083.