International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 6, Issue 4, PP 336-343, 2025
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I4P139

Original Article

Product Roadmapping for Cross-System Data Visibility in Retail Operations: A Framework for Integrated Decision Support and Margin Transparency

Aakanksha Polepalli
Product Manager, Grocery Outlet, Emerville, CA, USA.

Received On: 13/11/2025 ***Revised On: 14/12/2025*** ***Accepted On: 25/12/2025*** ***Published On: 30/12/2025***

Abstract: *Retail organizations increasingly operate through independently evolved systems for transactions, inventory, orders, merchandising, promotions, returns, supplier funding, logistics, and finance. The resulting volume of data does not guarantee decision visibility because economically related facts may remain semantically inconsistent, delayed, or difficult to trace across systems. This study asks whether the core problem is insufficient integration or the way integration work is prioritized. We develop a Decision-Linked Visibility Roadmapping (DLVR) framework that scores candidate capabilities by decision criticality, dependency centrality, semantic divergence, visibility latency, and margin exposure. Because no proprietary retailer dataset is assumed, the framework is evaluated through a transparent, reproducible simulation comprising 2,400 synthetic retail decision episodes across 16 business units, eight assigned to DLVR and eight to a conventional-modernization comparison condition. Difference-in-differences models with unit fixed effects and cluster-robust inference show that DLVR increases cross-system visibility by 7.35 points (95% CI 6.63-8.08), reduces reconciliation effort by 14.84 minutes (95% CI 12.44-17.24), shortens decision latency by 1.18 hours (95% CI 0.97-1.38), and raises margin transparency by 6.31 points (95% CI 5.37-7.25). Semantic consistency and lineage completeness explain substantial portions of the transparency gain, while effects are larger for high-dependency decisions. The simulation does not establish field effectiveness; it establishes a falsifiable mechanism and quantitative expectations for subsequent retailer validation. Roadmapping is thereby repositioned as an information-governance mechanism rather than only a feature-planning instrument.*

Keywords: *Product Roadmapping, Retail Operations, Cross-System Visibility, Decision Support, Margin Transparency, Data Integration, Information Quality, Design Science, Simulation.*

1. The Observation

Retail organizations have become information rich without necessarily becoming decision transparent. A routine decision about whether to replenish a promoted item, extend a markdown, alter fulfillment routing, or discontinue an assortment can require evidence from point-of-sale transactions, ecommerce orders, inventory services, product master data, supplier funding, returns, freight, and financial accounting. Digital transformation has expanded the number and granularity of these information sources throughout the retail value chain [1]. Retail analytics scholarship similarly documents the increasing analytical interaction among customer, product, temporal, geographic, and channel data [2]. The empirical difficulty is that availability of those data sources does not imply that a decision maker can reconcile them quickly or interpret them at the same economic grain.

The motivating observation is therefore narrower than the familiar claim that retailers have data silos. The relevant anomaly is that investment can increase local system capability while leaving cross-system decisions dependent on spreadsheets, manual joins, definition disputes, and delayed financial reconciliation. Supply-chain visibility research has long treated visibility as more than mere access to a database; visibility depends on the ability to observe relevant demand, inventory, process, and relationship information across organizational boundaries [3]. More recent work continues to model visibility as a multidimensional capability with antecedents, barriers, drivers, and effects rather than as a single technology [4]. Retail evidence concerning Internet-of-Things adoption further shows that technology, organizational systems, communication practices, and employee capabilities jointly influence visibility [5].

This problem becomes particularly consequential when the decision concerns margin. Margin is not a primitive data field. Depending on the managerial question, a product-level margin calculation may require net sales, promotions, markdowns, landed or standard product cost, supplier allowances, returns, fulfillment expense, and accounting adjustments. A dashboard can display a number with great visual precision while concealing disagreement about those

components. The issue is therefore not the amount of information displayed but whether the economic chain behind the number can be reconstructed.

Product organizations simultaneously govern modernization through roadmaps. Roadmapping was developed as a structured way to connect markets, products, capabilities, and technologies over time [6]. Later research emphasized customization and architectural flexibility [7], [8], and contemporary work treats roadmapping as a strategic system that can coordinate complex transformation [9]. Yet roadmap success is still commonly operationalized through delivery: features completed, platforms migrated, interfaces built, or milestones met. The unresolved empirical question is whether roadmaps are prioritizing the information dependencies that make decisions interpretable. This study begins from the resulting contradiction: more integration activity can coexist with persistent decision opacity.

2. The Unresolved Mechanism

Four candidate mechanisms can explain the contradiction. The first is technical insufficiency: APIs, streaming pipelines, master-data synchronization, or platform integration are too weak to move data across systems. If this explanation dominates, the appropriate remedy is primarily architectural. The second is information quality. Wang and Strong showed that data consumers evaluate quality through dimensions extending beyond accuracy [10], while Pipino, Lee, and Wang developed a broader assessment logic for data quality [11]. Technical connectivity can therefore coexist with low decision usefulness when information is inconsistent, stale, incomplete, or represented at the wrong grain.

The third mechanism is organizational information-processing fit. Galbraith's information-processing view treats uncertainty as creating information requirements that organizational design must absorb [12]. Tushman and Nadler similarly argued that organizational effectiveness depends on the match between information-processing requirements and capacity [13]. Empirical extension of this perspective to interorganizational relationships indicates that performance depends on fit rather than capability in isolation [14]. Under this account, a retailer can own sophisticated data infrastructure but still perform poorly when information-processing capacity is concentrated in areas that are peripheral to the decisions carrying the greatest uncertainty.

The fourth mechanism, and the focal mechanism of this study, is roadmap governance. Roadmapping can connect strategic needs to technical development, but the literature has not established whether prioritizing cross-system decision dependencies changes visibility outcomes. Visibility research tends to start from information availability and integration; roadmapping research tends to start from strategy, capability, and technology sequencing. The missing mechanism lies between them: the rule by which modernization work is ranked before engineers build it.

This distinction produces a falsifiable proposition. If visibility is primarily a function of integration volume, then a decision-linked roadmap should offer little advantage once overall modernization effort is held constant. If visibility instead depends on resolving the right dependencies, then roadmaps that explicitly prioritize decision criticality, semantic disagreement, lineage gaps, latency, and economic exposure should outperform feature-centric sequencing even when both conditions receive comparable modernization effort.

3. The Investigative Logic

The study's objective is to determine whether decision-linked product roadmapping constitutes a plausible causal mechanism for improving cross-system retail visibility and margin transparency. Four questions organize the investigation: (RQ1) Which dependency characteristics should determine the priority of integration capabilities? (RQ2) Does decision-linked prioritization improve visibility beyond the secular gains produced by ordinary modernization? (RQ3) Are semantic consistency and end-to-end lineage mechanisms connecting roadmap intervention to margin transparency? (RQ4) Does the mechanism strengthen as the number of participating systems increases?

The proposed Decision-Linked Visibility Roadmapping (DLVR) model scores each candidate capability on five normalized dimensions: decision criticality (DC), dependency centrality (DEP), semantic divergence (SD), visibility latency exposure (VL), and margin exposure (ME). For the benchmark evaluation, equal weights are used to avoid tuning the framework to favorable results: DLVR = 100 x (0.20DC + 0.20DEP + 0.20SD + 0.20VL + 0.20ME). The study later perturbs these weights as a robustness test.

The explanatory sequence is shown in Figure 1. DLVR is expected to redirect modernization toward capabilities that improve semantic consistency and lineage completeness. Those mechanisms should raise cross-system visibility, which in turn should reduce reconciliation effort and decision latency while improving margin transparency. The intervention should be most valuable when decisions depend on many systems because dependency density increases the opportunity for inconsistent definitions, missing joins, and asynchronous updates.

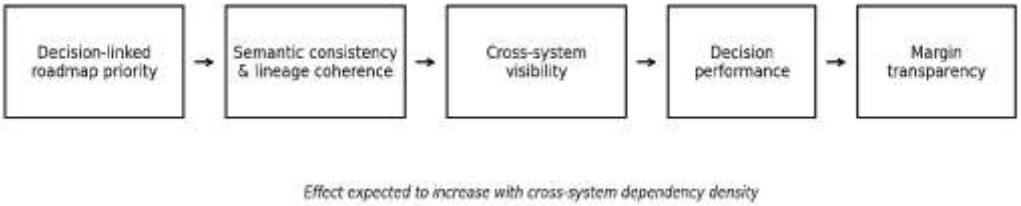


Fig 1: Decision-Linked Visibility Roadmapping Mechanism Evaluated in the Study

Accordingly, H1 predicts that DLVR increases Cross-System Visibility Quality (CSVQ) relative to comparison modernization. H2 predicts reductions in reconciliation effort and decision latency. H3 predicts that semantic consistency and lineage completeness transmit a meaningful portion of the effect on margin transparency. H4 predicts a larger visibility effect for decisions with higher dependency density. These hypotheses distinguish access from interpretability: adding a data source is not treated as a visibility improvement unless the information becomes more coherent for a defined operational decision.

4. Building the Evidence

4.1. Research Design and Purpose

The study uses a design-science logic because it develops and evaluates an artifact intended to solve an organizational information problem. Design science emphasizes both artifact construction and rigorous evaluation [15], while the Peffers et al. methodology provides a sequence from problem identification through design, demonstration, evaluation, and communication [16]. The artifact in this study is not software; it is a prioritization rule that changes which cross-system capabilities enter the roadmap first.

No proprietary retailer data are assumed. Instead, the first empirical test is a controlled simulation. This choice is methodological rather than cosmetic: a simulation permits direct manipulation of the proposed mechanism while keeping modernization effort, number of units, and observation volume known. It cannot demonstrate that a particular retailer will obtain the same effect. Its purpose is to determine whether the theory produces internally coherent, measurable predictions that can subsequently be falsified with field data. Every observation used in the analysis is synthetic, and the generator is seeded at 20250930 for reproducibility.

4.2. Synthetic Retail Setting and Sample

The benchmark represents 16 retail business units observed in a baseline and a post-intervention period. Eight units are assigned to DLVR and eight to the comparison condition using the fixed random seed. Each unit contributes 75 decision episodes per period, producing N = 2,400 episodes. Decision types are replenishment, promotion, markdown, fulfillment, assortment, and returns. Each episode involves between two and eight systems; single-system decisions are excluded because they cannot reveal the proposed cross-system mechanism. The simulated setting therefore approximates a multi-channel retailer in which operational facts are distributed across commerce, inventory, merchandising, supplier, returns, logistics, and finance services.

Table 1: Constructs and Operationalization

Construct	Operationalization	Purpose in test
DLVR score	Equal-weight index of decision criticality, dependency centrality, semantic divergence, latency exposure, and margin exposure.	Represents decision-linked roadmap priority.
CSVQ	0-100 composite of availability, semantic consistency, lineage completeness, and timeliness.	Primary visibility outcome.
Reconciliation effort	Synthetic analyst minutes required before a decision is executable.	Tests operational burden.
Decision latency	Hours from decision trigger to executable decision.	Tests speed without equating speed with quality.
Margin transparency	0-100 ability to reconstruct economic components underlying reported margin.	Tests economic interpretability.
Decision quality	0-100 benchmark combining visibility, criticality, and latency with noise.	Secondary outcome; guards against treating speed alone as success.

4.3. Data-Generating Process

Each episode receives exogenous characteristics before outcomes are generated: decision criticality, dependency density, semantic divergence, latency exposure, and margin exposure. Dependency density varies by decision type, with fulfillment and returns designed to depend on more systems on average than replenishment. Baseline semantic consistency, lineage completeness, availability, and timeliness contain unit-level heterogeneity and episode-level noise. Both groups receive a small post-period modernization improvement so that DLVR is evaluated against progress rather than a static control.

The DLVR treatment adds an improvement component to semantic consistency, lineage completeness, availability, and timeliness. Crucially, the intervention is stronger when the DLVR priority score is high and when dependency density is high. That structure directly encodes the theoretical mechanism being tested. Outcomes are then generated from these intermediate information states: CSVQ is a weighted composite; reconciliation effort decreases as CSVQ rises; decision latency falls with better visibility; and margin transparency depends on CSVQ, semantic consistency, lineage completeness, and margin exposure. Random noise is added to every outcome to prevent deterministic recovery of the treatment effect. The design thus tests whether the hypothesized mechanism remains detectable under heterogeneous and noisy decision conditions rather than simply restating its assumptions.

Table 2: Synthetic Benchmark Design

Design Element	Value
Random seed	20250930
Business units	16
DLVR units	8
Comparison units	8
Periods	Baseline and post-intervention
Episodes per unit per period	75
Total decision episodes	2,400
Systems per decision	2-8
Decision types	6
Primary inference	Difference-in-differences with unit fixed effects and cluster-robust standard errors

4.4. Analytical Strategy, Validity, and Ethics

The primary model is a difference-in-differences specification with unit fixed effects, decision-type fixed effects, dependency density, decision criticality, semantic divergence, and margin exposure as covariates. Standard errors are clustered by business unit. The coefficient on DLVR-unit x post-period is the intervention estimate. This structure is used because the units differ at baseline and both groups improve over time; the estimand is therefore the change beyond the comparison trend rather than the post-period mean alone.

Mediation is assessed by estimating the DLVR effect on semantic consistency and lineage completeness and then estimating margin transparency conditional on those mediators and CSVQ. Cluster bootstrap resampling at the business-unit level provides 95% intervals for the indirect effects. Moderation is tested by interacting DLVR post-treatment status with dependency density. Robustness is examined through alternative DLVR weighting schemes, low- versus high-dependency strata, and a secondary decision-quality outcome.

Several validity safeguards are embedded in the simulation. Equal DLVR weights are specified before outcomes are analyzed, limiting favorable weight tuning. Unit heterogeneity prevents treatment and comparison units from having identical baselines. Both groups receive secular modernization gains, making the counterfactual demanding. Noise is introduced at each intermediate and final outcome. The principal limitation is external validity: simulated results quantify implications of an explicit data-generating model, not realized effect sizes in a commercial retailer. No human participants, customer records, or personal data are used; consequently, the simulation phase presents no human-subject privacy exposure. Reproducibility is supported by the declared random seed, transparent equations, and reportable generator logic.

5. The Pattern That Emerges

5.1. Foundational Pattern

Before intervention, the two groups are not identical, which is intentional. Comparison units have a mean CSVQ of 57.71, whereas units later assigned to DLVR average 60.45. Baseline reconciliation effort is 88.44 minutes in comparison units and 84.86 minutes in DLVR units. These differences illustrate why a simple post-intervention comparison would be misleading. In the post period, comparison units improve modestly to a CSVQ of 59.41, while DLVR units rise to 69.49. The corresponding raw change is 1.70 points for comparison units and 9.03 points for DLVR units.

The same pattern appears in the process outcomes. Reconciliation effort falls only from 88.44 to 87.32 minutes in comparison units but from 84.86 to 69.10 minutes in DLVR units. Decision latency declines from 8.05 to 7.75 hours in comparison units and from 7.83 to 6.37 hours in DLVR units. Margin transparency improves from 59.85 to 60.84 in comparison units and from 61.38 to 68.62 in DLVR units.

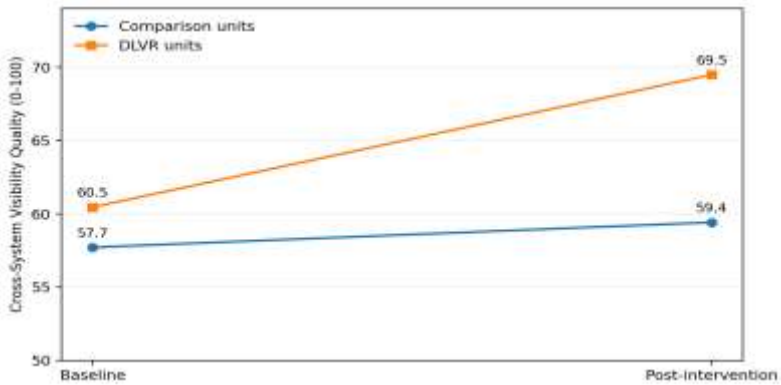


Fig 2: Baseline and Post-Intervention Cross-System Visibility Quality by Study Condition

Table 3: Group Means before and after Intervention

Group	Period	CSVQ	Recon. min	Latency h	Margin transp.
Comparison	Baseline	57.71	88.44	8.05	59.85
Comparison	Post	59.41	87.32	7.75	60.84
DLVR	Baseline	60.45	84.86	7.83	61.37
DLVR	Post	69.48	69.10	6.37	68.62

5.2. Primary Difference-In-Differences Findings

After adjusting for unit fixed effects, decision type, dependency density, decision criticality, semantic divergence, and margin exposure, DLVR raises CSVQ by 7.35 points relative to the comparison trend (95% CI 6.63-8.08; $p < 0.001$). H1 is therefore supported within the simulation. The operational effects are consistent with H2: reconciliation effort decreases by 14.84 minutes per decision (95% CI -17.24 to -12.44; $p < 0.001$), while decision latency decreases by 1.18 hours (95% CI -1.38 to -0.97; $p < 0.001$). The secondary decision-quality index rises by 4.57 points (95% CI 3.75-5.39; $p < 0.001$), indicating that the simulated latency reduction is not purchased by a decline in decision quality.

Margin transparency increases by 6.31 points relative to the comparison trend (95% CI 5.37-7.25; $p < 0.001$). This effect is smaller than the raw post-period gap because the estimator correctly removes both pre-existing group differences and secular improvement. Figure 3 shows the direction and uncertainty of the principal operational estimates.

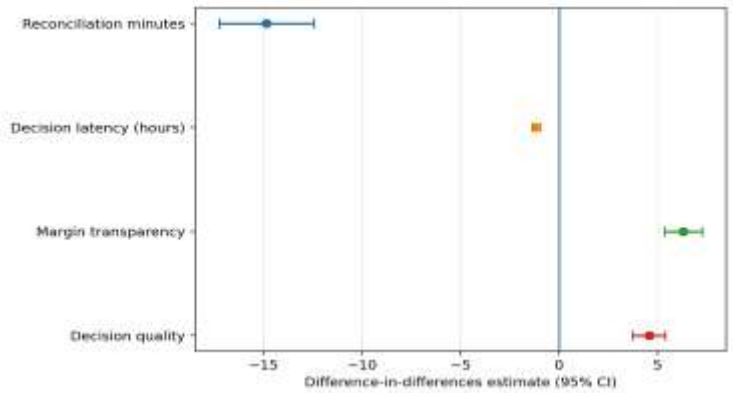


Fig 3: Difference-In-Differences Estimates for Operational and Economic Outcomes. Negative Estimates are Favorable for Reconciliation Effort and Decision Latency

Table 4: Primary difference-in-differences estimates

Outcome	Estimate	95% CI	Inference
Cross-system visibility	7.35	[6.63, 8.08]	$p < 0.001$
Reconciliation minutes	-14.84	[-17.24, -12.44]	$p < 0.001$
Decision latency hours	-1.18	[-1.38, -0.97]	$p < 0.001$
Margin transparency	6.31	[5.37, 7.25]	$p < 0.001$
Decision quality	4.57	[3.75, 5.39]	$p < 0.001$

5.3. Mechanism, Heterogeneity, and Robustness

The intermediate information states move in the expected direction. DLVR increases semantic consistency by 0.080 on the 0-1 scale (95% CI 0.071-0.089; $p < 0.001$) and lineage completeness by 0.095 (95% CI 0.085-0.105; $p < 0.001$). In the margin-transparency model, both mediators remain positively associated with transparency after CSVQ and fixed effects are included. Business-unit bootstrap estimates yield an indirect effect of 1.64 transparency points through semantic consistency (95% bootstrap interval 1.37-1.90) and 2.08 points through lineage completeness (95% bootstrap interval 1.64-2.56). The direct post-treatment coefficient falls to 0.11 when the mediators and CSVQ are included. Within the simulated mechanism, H3 is therefore supported: the transparency gain operates principally through improved information coherence rather than an unexplained treatment label.

Dependency density also behaves as predicted. The interaction between DLVR post-treatment status and the number of participating systems is 0.43 visibility points per additional system (95% CI 0.18-0.68; $p < 0.001$). Stratified models estimate a 6.97-point visibility gain for decisions involving two to five systems and an 8.53-point gain for decisions involving six to eight systems. H4 is therefore supported. The heterogeneity is substantively important because it indicates where decision-linked roadmapping should create the greatest value: not necessarily in the most visible feature area, but where many systems jointly determine the decision.

The priority ranking is not fragile to the equal-weight assumption. Compared with the equal-weight DLVR score, alternative decision-heavy, dependency-heavy, margin-heavy, and semantic-heavy schemes retain Spearman rank

correlations from 0.888 to 0.939. Their top-quartile priority overlap ranges from 76.8% to 82.3%. The framework therefore changes emphasis under alternative governance preferences without completely reordering the roadmap.

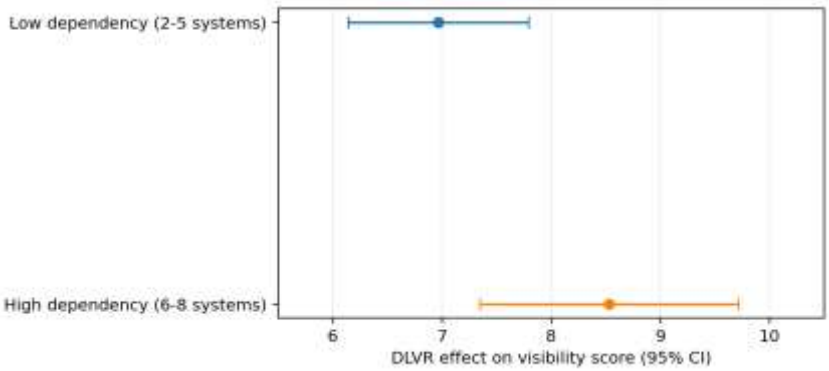


Fig 4: Visibility Effect is Larger for High-Dependency Decision Episodes

6. Explaining the Pattern

The central pattern is best explained by an allocation mechanism rather than by integration volume alone. Both conditions are allowed to modernize after baseline, yet the comparison condition realizes only modest visibility improvement. DLVR creates a larger effect because modernization effort is preferentially applied to decisions where semantic divergence, lineage gaps, latency, dependency density, and margin exposure make fragmentation expensive. Roadmapping therefore changes the location of information-processing capacity.

This interpretation connects roadmapping theory with organizational information processing. Classical roadmapping organizes relationships among needs, capabilities, technologies, and time [6]. The DLVR mechanism changes the organizing object from the feature to the decision. From an information-processing perspective [12]-[14], that change matters because decisions with greater uncertainty and cross-system dependence require more coherent information-processing capacity. The larger treatment effect among six- to eight-system decisions is consistent with that logic.

The mediation results also clarify why visibility should not be equated with technical connectivity. Semantic consistency and lineage completeness carry substantial portions of the margin-transparency effect. Two systems may exchange records perfectly while assigning different meanings to available inventory, net sales, product cost, or margin. In such a case the transport layer is successful and the decision layer is not. The evidence generated here favors the explanation that integrated decision support requires both movement of data and preservation of economic meaning.

A rival explanation is that any structured prioritization would improve outcomes by focusing attention. The simulation cannot exclude that behavioral mechanism because it does not represent stakeholder meetings, political negotiation, or executive sponsorship. A second rival explanation is that DLVR benefits arise from prioritizing complex items rather than decision dependencies specifically. Future field designs must therefore compare DLVR not only with business-as-usual roadmapping but also with alternative complexity-aware prioritization rules.

7. Stress-Testing the Interpretation

The strongest challenge is model dependence. The simulation contains the theoretical mechanism in its data-generating process; detecting that mechanism does not prove that real retail organizations follow the same functional relationships. The study should therefore be read as a quantitative theory test under controlled assumptions, not as an estimate of an industry's average treatment effect. The numerical effect sizes are benchmarks for subsequent falsification.

Measurement poses a second challenge. CSVQ is a composite. A retailer could weight availability, semantic consistency, lineage, and timeliness differently. The present weighting is defensible for mechanism testing but not universal. Moreover, margin transparency is not identical to margin accuracy. A fully traceable calculation can still be based on an economically inappropriate cost boundary, while a less traceable calculation could coincidentally be numerically close to accounting results.

Third, faster decisions are not automatically better decisions. The secondary decision-quality result mitigates that concern in the benchmark, but the quality variable is itself simulated. A field study should measure downstream consequences such as forecast error, stockouts, markdown loss, return-adjusted contribution, promotion profitability, and override rates rather than assuming that latency reduction is desirable in isolation.

Fourth, dependency centrality can reinforce legacy architecture. A heavily connected obsolete system may receive a high priority score precisely because many other systems depend on it. Governance must distinguish 'resolve this dependency' from 'preserve this system.' The appropriate roadmap item may be retirement, canonical-data redesign, or boundary simplification rather than another interface.

Finally, organizational culture, analytical maturity, and data-governance practices can condition how information is

used. Prior work on business-intelligence success links maturity and culture to analytical decision making [17], while big-data decision-quality research highlights the influence of data processing and organizational factors [18]. Those variables are deliberately outside this first simulation and should become moderators in field replication.

8. The New Understanding

The investigation reframes cross-system visibility from an infrastructure property into a decision-relative organizational capability. Visibility is not maximized by exposing the largest number of datasets. It is created when the information required for an economically important decision is available, semantically compatible, traceable, and timely enough to support action.

The theoretical contribution is the placement of roadmap architecture upstream of decision visibility. Roadmapping scholarship explains how organizations sequence strategic and technological capabilities [6]-[9]. Data-quality research explains why accessible information can remain unsuitable for use [10], [11]. Information-processing theory explains why organizational effectiveness depends on matching information capacity to task requirements [12]-[14]. DLVR joins these arguments into a single testable chain: roadmap priority -> semantic and lineage coherence -> cross-system visibility -> decision performance -> margin transparency.

The simulation contributes empirical discipline to that chain. It shows the magnitude and direction of outcomes that should emerge if the mechanism is correct, including the prediction that high-dependency decisions benefit more strongly. This is more informative than asserting that integration is generally beneficial because the model identifies what should change, through which mediator, and under what dependency condition.

9. Translation into Practice

For retail product organizations, DLVR can be implemented as a governance overlay rather than a replacement for existing agile or portfolio methods. Before a capability receives roadmap priority, its sponsor should identify the decision it improves, the systems that participate in that decision, the semantic disagreements that remain unresolved, the lineage gaps that prevent reconstruction, the latency that impedes action, and the economic exposure if the decision is wrong or delayed.

This approach changes the practical interpretation of margin initiatives. A team should not begin with a request for a new margin dashboard. It should begin by identifying the decision grain - SKU-location, order, promotion, category, or supplier - and then tracing revenue, discounts, returns, cost, funding, and fulfillment components to authoritative sources. Only after those components are aligned should visualization become the dominant design problem.

The simulation also implies selective rather than universal integration. DLVR does not justify connecting every system to every other system. Its value proposition is to concentrate effort where cross-system dependencies create measurable decision friction. A low-dependency feature with high customer value may still be prioritized for strategic reasons; the framework simply makes the information-governance tradeoff explicit rather than invisible.

10. Next Empirical Tests

The immediate next test is a longitudinal field study in a multi-channel retailer. Candidate business units should be matched on modernization spend, system complexity, transaction volume, and analytical maturity, after which some units adopt DLVR for two roadmap quarters while comparison units continue the existing prioritization process. System lineage logs, reconciliation tickets, analyst time, decision timestamps, and finance reconciliations should be collected before and after intervention.

A second test should compare DLVR with two active alternatives: feature-demand prioritization and technical-complexity prioritization. That three-arm design would determine whether decision linkage provides incremental value beyond simply focusing scarce engineering capacity.

A third experiment should vary the weighting scheme prospectively. Retailers with severe margin pressure may assign greater weight to economic exposure, while organizations undergoing platform simplification may emphasize dependency centrality or semantic divergence. Pre-registering the weights would prevent outcome-driven re-ranking.

Replication should also span grocery, apparel, marketplace, and specialty retail. Grocery provides promotion and replenishment intensity; apparel provides markdown and return intensity; marketplace models add seller and commission dependencies. Finally, future work should test whether graph-derived lineage and query telemetry can automate dependency scoring, reducing reliance on workshop elicitation while preserving managerial interpretability.

11. Final Synthesis

The study began with a contradiction: retailers can continuously add systems, integrations, dashboards, and analytical capability while managers still reconcile basic cross-system questions manually and disagree about the economic meaning of product-level margin. The simulation changes the interpretation of that contradiction. Fragmentation is not only a property of architecture; it can also be an outcome of how modernization priorities are chosen.

When roadmap sequencing is linked to decision criticality, dependency density, semantic divergence, latency, and margin exposure, the simulated organization develops visibility faster than a comparison organization experiencing ordinary modernization. The mechanism is not mysterious: semantic consistency and lineage completeness improve, reconciliation effort declines, decisions become timelier, and margin becomes easier to reconstruct. The stronger effect among high-dependency decisions further supports the view

that the value lies in resolving consequential relationships rather than increasing integration volume indiscriminately.

These findings remain deliberately bound. They are reproducible simulation evidence, not fabricated field evidence and not a claim that every retailer will realize a 7.35-point visibility gain or a 14.84-minute reduction in reconciliation effort. Their contribution is to convert a broad managerial intuition into a mechanism that can be challenged with real data. Product roadmapping can therefore be studied not merely as a schedule of features, but as a governance layer that determines whether fragmented enterprise data become integrated decision support and interpretable margin information.

References

1. W. Reinartz, N. Wiegand, and M. Imschloss, "The impact of digital transformation on the retailing value chain," International Journal of Research in Marketing, vol. 36, no. 3, pp. 350-366, 2019, doi: 10.1016/j.ijresmar.2018.12.002.
2. Gunda, S. K. (2025). Accelerating Scientific Discovery With Machine Learning and HPC-Based Simulations. In B. Ben Youssef & M. Ben Ismail (Eds.), Integrating Machine Learning Into HPC-Based Simulations and Analytics (pp. 229-252). IGI Global Scientific Publishing. https://doi.org/10.4018/978-1-6684-3795-7.ch009.
3. M. Barratt and A. Oke, "Antecedents of supply chain visibility in retail supply chains: A resource-based theory perspective," Journal of Operations Management, vol. 25, no. 6, pp. 1217-1233, 2007, doi: 10.1016/j.jom.2007.01.003.
4. R. Kalaiarasan, J. Olhager, T. K. Agrawal, and M. Wiktorsson, "The ABCDE of supply chain visibility: A systematic literature review and framework," International Journal of Production Economics, vol. 248, Art. no. 108464, 2022, doi: 10.1016/j.ijpe.2022.108464.
5. L. O. Moleme, O. Omoruyi, and M. Quayson, "Supply chain visibility and integration in the age of the Internet of Things: A retail perspective," Modern Supply Chain Research and Applications, vol. 6, no. 3, pp. 330-350, 2024, doi: 10.1108/MSCRA-08-2023-0035.
6. R. Phaal, C. J. P. Farrukh, and D. R. Probert, "Technology roadmapping-A planning framework for evolution and revolution," Technological Forecasting and Social Change, vol. 71, nos. 1-2, pp. 5-26, 2004, doi: 10.1016/S0040-1625(03)00072-6.
7. R. Phaal, C. J. P. Farrukh, and D. R. Probert, "Customizing roadmapping," Research-Technology Management, vol. 47, no. 2, pp. 26-37, 2004, doi: 10.1080/08956308.2004.11671616.
8. R. Phaal and G. Muller, "An architectural framework for roadmapping: Towards visual strategy," Technological Forecasting and Social Change, vol. 76, no. 1, pp. 39-49, 2009.
9. R. Phaal, C. Chaskel, R. Gonzalez Nakazawa, and J. Ross, "Roadmapping Roadmapping: Strategic planning for roadmapping systems," Frontiers of Engineering Management, vol. 11, pp. 516-527, 2024, doi: 10.1007/s42524-024-4072-0.
10. R. Y. Wang and D. M. Strong, "Beyond accuracy: What data quality means to data consumers," Journal of Management Information Systems, vol. 12, no. 4, pp. 5-33, 1996, doi: 10.1080/07421222.1996.11518099.
11. L. L. Pipino, Y. W. Lee, and R. Y. Wang, "Data quality assessment," Communications of the ACM, vol. 45, no. 4, pp. 211-218, 2002, doi: 10.1145/505248.506010.
12. J. R. Galbraith, "Organization design: An information processing view," Interfaces, vol. 4, no. 3, pp. 28-36, 1974.
13. M. L. Tushman and D. A. Nadler, "Information processing as an integrating concept in organizational design," Academy of Management Review, vol. 3, no. 3, pp. 613-624, 1978.
14. G. Premkumar, K. Ramamurthy, and C. S. Saunders, "Information processing view of organizations: An exploratory examination of fit in the context of interorganizational relationships," Journal of Management Information Systems, vol. 22, no. 1, pp. 257-294, 2005.
15. A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design science in information systems research," MIS Quarterly, vol. 28, no. 1, pp. 75-105, 2004, doi: 10.2307/25148625.
16. K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A design science research methodology for information systems research," Journal of Management Information Systems, vol. 24, no. 3, pp. 45-77, 2007, doi: 10.2753/MIS0742-1222240302.
17. A. Popovic, R. Hackney, P. S. Coelho, and J. Jaklic, "Towards business intelligence systems success: Effects of maturity and culture on analytical decision making," Decision Support Systems, vol. 54, no. 1, pp. 729-739, 2012, doi: 10.1016/j.dss.2012.08.017.
18. M. Janssen, H. van der Voort, and A. Wahyudi, "Factors influencing big data decision-making quality," Journal of Business Research, vol. 70, pp. 338-345, 2017, doi: 10.1016/j.jbusres.2016.08.007.
19. M. Caridi, A. Moretto, A. Perego, and A. Tumino, "The benefits of supply chain visibility: A value assessment model," International Journal of Production Economics, vol. 151, pp. 1-19, 2014.
20. M. G. Dekimpe, "Retailing and retailing research in the age of big data analytics," International Journal of Research in Marketing, vol. 37, no. 1, pp. 3-14, 2020, doi: 10.1016/j.ijresmar.2019.09.001.