



MANAGEMENT VARIANCE BRIEF / CC-410

ACTUAL

\$4.82M

PLAN

\$4.30M

LBE

\$4.58M

PLAN VAR.

+\$520K

MODELED DRIVER BRIDGE

RATE



\$210K

VOLUME



\$160K

MIX



\$90K

TIMING



\$60K

ASSESSMENT

MATERIAL

82% COVERAGE

HUMAN REVIEW

FINANCE RESEARCH MEMORANDUM / SYSTEM 007

Financial Variance Command Center

An inspectable control room for financial variance investigation.

Fatima Aguilar · Decision Systems Lab · July 2026

PUBLIC RESEARCH RELEASE · SYNTHETIC FINANCIAL DATA

Live: fa366193.github.io/financial-variance-command-center/

Source: github.com/fa366193/financial-variance-command-center

Abstract

A financial result is not yet an explanation.

The Site Financial Health Engine establishes a standardized analytical model across SAP actuals, general-ledger detail, Plan, Latest Business Estimate, cost-center hierarchies, and operational measures. Financial Variance Command Center is an experimental extension that turns this trusted model into a visible investigation environment.

Six bounded roles validate the record, decompose volume/rate/mix/timing, connect operational signals, challenge the forecast narrative, inspect risk and controls, and synthesize a management brief. Visitors can change assumptions, observe an 18-second investigation, inspect the trace, record a disposition, and export the synthetic brief.

The current release is deterministic and browser-based. It does not connect to SAP or Power BI, use external language models, post journal entries, or claim measured decision improvement.

Research question

Can visible, adversarial agent roles improve the completeness and accountability of financial variance explanations without weakening deterministic controls?

Layer	Function	Status
Unified financial model	Actual, Plan, LBE and mappings	In development
Agentic review	Visible variance investigation	Public prototype
ERP write-back	Journal or forecast action	Not implemented
Controlled evaluation	Hybrid vs rules-only baseline	Proposed
External LLM agents	Tool-using reasoning roles	Future research

Development notebook

Six versions of one financial problem

V1 - Fragmented reporting

KSB1 Actuals, FBL3N GL detail, Dodeca Plan/LBE, master data, and operating measures were prepared independently.

V2 - Automated pipeline

PowerShell validation, normalization, enrichment, mapping, consolidation, and Power BI refresh replaced recurring manual preparation.

V3 - Unified semantic model

Cost center, lead, function, GL account, category, scenario, period, month, and amount became a consistent analytical schema.

V4 - Financial health surface

Executive, finance-analyst, and operations views referenced one validated model.

V5 - Driver intelligence

Variance was decomposed into volume, rate, mix, and timing so numbers could carry an explanatory structure.

V6 - Variance Command Center

Six bounded roles investigate one material variance, challenge the leading theory, simulate corrective assumptions, and prepare a human-reviewable brief.

The agentic layer was added after the deterministic data and logic foundation. It is a contestable review surface, not a replacement for the semantic model or financial controls.

Hybrid architecture

A trusted backbone with a challenge layer

MANAGEMENT VARIANCE BRIDGE



Source and model layers remain deterministic: SAP FI/CO and planning extracts are validated, normalized, mapped, and consolidated before any narrative is generated. The prototype then assigns bounded analytical responsibilities around the same record.

Agent team

Role	Bounded responsibility
Data Steward	Completeness, schema, GL and cost-center mapping
Variance Analyst	Volume, rate, mix, timing, residual bridge
Operations Liaison	Labor, throughput, utilization, operational context
Forecast Challenger	Alternate baselines, counterarguments, outlook fragility
Risk Auditor	Cutoff, accrual, anomaly, control and evidence risks
Finance Synthesizer	Evidence, disagreement, management action, trace

The Forecast Challenger is structurally required before synthesis. Agreement is therefore an output to be earned, not the default behavior of cooperating agents.

Financial logic

Variance is the language. Drivers are hypotheses.

The prototype computes Actual minus Plan and Actual minus LBE, ranks cases by materiality, and displays a four-driver bridge. The supplied volume, rate, mix, and timing components are treated as modeled explanations rather than ground truth.

Measure	Definition
Plan variance	Actual - Plan
LBE variance	Actual - Latest Business Estimate
Variance percent	(Actual - Plan) / Plan
Driver residual	Plan variance - volume - rate - mix - timing
Materiality	variance percent > 8% = Material; > 4% = Watch
Operational risk	Variance magnitude plus utilization and throughput deviation

Residual transparency

The unexplained residual remains visible. A polished narrative must not erase what the driver bridge fails to account for. In a production system, source-level provenance would connect every driver to a reproducible calculation and evidence record.

Interaction design

A decision review, not a financial chatbot

Visitors begin with five synthetic cost-center cases ranked by absolute variance. Selecting a case exposes Actual, Plan, LBE, driver amounts, and operating indices. Editing any input deterministically changes the assessment.

Stage	Visible work
01 - Validate	Check schema, mappings, and completeness
02 - Decompose	Build the volume/rate/mix/timing bridge
03 - Connect	Compare labor, throughput, and utilization
04 - Challenge	Contest the leading explanation and expose residuals
05 - Simulate	Test alternate driver and operating assumptions
06 - Brief	Prepare assessment, evidence, action, and owner

Management retains three terminal choices: accept the explanation, request additional evidence, or escalate the forecast. The exported JSON brief records the synthetic case, analysis, timestamp, disposition, and research boundary. It does not create a journal entry or forecast submission.

Counterfactual laboratory

What would change the outlook?

Four sliders vary volume, rate, labor, and throughput. The interface recomputes plan variance, largest driver, materiality, and operational risk beside the baseline.

Counterfactuals make the decision boundary inspectable. A management action that reverses under a small assumption change is fragile; an explanation that survives a wide range of plausible inputs is more robust.

Example: Aseptic Operations

The synthetic case reports \$4.82M Actual against \$4.30M Plan and \$4.58M LBE. Rate is the largest modeled driver, while throughput and utilization provide operating context. The Challenger warns that the attribution may absorb cutoff or mapping effects and calls attention to any residual outside the bridge.

Current limitation

The prototype accepts entered driver values rather than deriving them from transaction-level volume and rate records. Production research should calculate each component from source data, preserve lineage, and test whether the four categories remain mutually exclusive and collectively useful.

Evaluation protocol

The system must be able to fail the experiment.

A controlled study should compare the hybrid interface with the existing financial dashboard and a rules-only explanatory baseline. Historical cases should be held out from interface tuning and adjudicated by finance and operations experts.

Metric	Operational definition
Time to explanation	Minutes from opening a material case to disposition
Driver accuracy	Agreement with adjudicated causal decomposition
Evidence coverage	Case-critical data and caveats acknowledged
Residual visibility	Unexplained bridge value recognized before action
Forecast calibration	Confidence versus later realized financial result
Decision consistency	Agreement across reviewers with the same record
Comprehension	Ability to explain why the action was proposed
Business outcome	Forecast error, corrective action latency, avoided leakage

Ablations

Rules-only: isolates the value of specialized review. **No Challenger:** tests mandatory dissent. **No Operations Liaison:** tests cross-functional context. **Hidden trace:** tests whether visibility improves comprehension or only adds time.

Governance and limitations

The prototype stops before financial authority begins.

Synthetic data

No production financial or employee data appears in the public app.

No autonomous accounting

No journal, accrual, forecast, or budget transaction is created.

Human authority

Finance owns interpretation; operations owns context; authorized leaders own disposition.

Controls

Production requires RBAC, segregation of duties, approval thresholds, period locks, and reversible write-back.

Security

Source documents and free text are untrusted data; prompt and data injection must not become instruction.

Lineage

Every mapping, rule, model, prompt, evidence source, and disposition requires durable versioning.

Correlated errors

Specialized agents may still repeat the same mistaken assumptions or persuasive narrative.

Override learning

Human overrides are review signals, not automatically correct labels for retraining.

Research boundary

The objective is not an autonomous CFO. It is a financial investigation system that makes uncertainty, evidence, disagreement, and authority easier to see.

Conclusion

From financial visibility to accountable explanation

The Site Financial Health Engine establishes the prerequisite for agentic finance: a standardized, validated semantic model. Financial Variance Command Center explores what becomes possible once that foundation exists.

The contribution is a hybrid architecture in which deterministic financial logic remains inspectable, specialized roles make competing interpretations visible, counterfactuals expose decision boundaries, and the final disposition stays human.

Public artifacts**Live application**

<https://fa366193.github.io/financial-variance-command-center/>

Public source

<https://github.com/fa366193/financial-variance-command-center>

Research notebook

<https://decisionsystemslab.org/systems/site-financial-health-engine>



Decision Systems Lab · System 007 · Fatima Aguilar · 2026