Roundtrip Value Governance for Agentic Process Automation

Marco van Hurne

EIGENVECTOR RESEARCH

October 21, 2025

Author Note

Correspondence concerning this article should be addressed to Marco van Hurne, EIGENVECTOR RESEARCH, October 21, 2025. E-mail: marco.vanhurne@eigenvector.eu

Abstract

Enterprises increasingly deploy agentic systems in process environments where the decisive question is not merely whether an agent can act, but whether the selected execution pattern creates defensible value under governance, intervention, and evidence constraints. Earlier work on the Value Office framed value realization for agentic automation as a lifecycle problem connecting measurement, telemetry, operator supervision, and recognition logic. Recent work in the same research family further shows that token-efficiency alone is insufficient to choose the best operating pattern for a process. The economics of execution must be extended from token spend toward measurable orchestration economics, governance burden, intervention burden, and recognized value.[1] [2] [3] [4] [5] [6] This paper updates the roundtrip value framework by integrating these newer insights into a closed-loop operating model for agentic process automation. The revised model connects minimal-information prioritization, case-centric telemetry, governed execution optimization, human intervention design, and value recognition into one measurable architecture. It introduces a formal link between candidate selection and runtime execution choice by distinguishing a **portfolio priority function** from a **runtime pattern optimization function**. The first decides which initiatives deserve attention. The second decides which admissible execution pattern should be used for each case or process state.[1] [7] [8] The paper makes five contributions. First, it redefines the Value Office as a **roundtrip governance capability** rather than a downstream reporting function. Second, it formalizes a minimal-information prioritization mechanism aligned to released FTE, cycle-time reduction, error reduction, compliance improvement, and strategic fit. Third, it introduces a measurement-first optimization layer in which execution patterns are ranked by recognized value per total enterprise cost under hard governance constraints. Fourth, it integrates observability, explainability, and auditing into a single evidence spine centered on cases, trajectories, and interventions. Fifth, it proposes a pilot-ready research agenda and operating model for validating the framework in exception-rich, governance-heavy processes. The result is a

broader and more mathematically explicit academic formulation of roundtrip value governance for agentic process automation.[1] [2] [5] [6] [9] [10]

Keywords: roundtrip value, agentic process automation, Value Office, prioritization, patternomics, observability, recognized value, governed execution

1. Introduction

AI agents are shifting enterprise automation from deterministic workflow execution toward adaptive, tool-using, semi-autonomous process behavior. This shift creates opportunity, but it also creates a new control problem. In classic automation settings, organizations could often justify investment through deterministic throughput improvements or simple labor-substitution narratives. Agentic systems complicate this logic because outcomes depend on planning behavior, retrieval quality, tool selection, topology depth, exception handling, human interventions, and policy enforcement paths that do not fit neatly into conventional business-case accounting.[2] [5] [6]

A second difficulty concerns selection. Most organizations do not lack agentification ideas. They lack a defensible way to decide which process candidates deserve priority and under what governance assumptions they should be pursued. When prioritization is informal, politically driven, or dependent on over-engineered business cases, portfolios become distorted. Attractive but infeasible opportunities crowd out tractable ones, while governance-heavy processes are either accelerated too early or dismissed too quickly.[7] [8] [11] [12]

A third difficulty emerges after deployment. Even when a process candidate is selected sensibly and implemented competently, the apparent benefit of automation may dissolve once hidden intervention work, evidence deficits, governance frictions, or adoption failures are taken into account. For this reason, the enterprise question is not only whether the agent performed efficiently, but whether the process generated **recognized value** under bounded autonomy. A run that is fast and cheap but unusable, unexplainable, or non-recognizable is not economically superior merely because its token bill is low.[3] [4] [13]

This paper argues that the problems of **which candidate to choose**, **which execution pattern to run**, and **which value to recognize** should not be treated separately. They are different views of the same governed optimization problem. The Value Office is therefore reformulated here as a **roundtrip value capability** that begins at candidate intake, continues through governed deployment and runtime supervision, and loops back into portfolio reprioritization through measured evidence. The phrase *roundtrip value* is not rhetorical. It

denotes a closed-loop operating logic in which value assumptions made at intake are later tested against telemetry, intervention burden, evidence maturity, and recognized outcomes, and are then fed back into the next selection cycle.[1] [7]

2. Research Question and Positioning

The updated research question is as follows.

> **How can a roundtrip Value Office jointly optimize candidate prioritization, execution-pattern selection, operational control, and recognized value for agentic process automation using measurable, estimable, derived, and policy-defined variables?**

This formulation extends the earlier roundtrip paper in two important ways. First, it incorporates a measurement-first optimization perspective drawn from newer work on governed execution patterns. Second, it places the Value Office in the same research family as TOKENOMICS while preserving a distinct contribution. TOKENOMICS addresses token expenditure, budget negotiation, and quality-per-token logic. The present paper addresses the broader problem of selecting and governing the best **execution pattern** and the best **portfolio sequence** under enterprise constraints.[13]

In other words, tokenomics remains necessary, but it is not sufficient. Token efficiency is one layer of the economics of agentic execution. Pattern choice, governance burden, intervention burden, and value recognition are additional layers that materially alter which process should be prioritized and which runtime pattern should be preferred.[5] [6] [9]

3. Theoretical Grounding

3.1 Process analytics and baseline reconstruction

Process mining and predictive process monitoring provide the empirical foundation for reconstructing actual business flow, bottlenecks, deviations, and timing patterns from event data. This matters because value claims are unreliable when they are based only on workshops or sponsor estimates. In a roundtrip Value Office, process analytics provides the

basis for baseline formation, candidate validation, attribution analysis, and continuous value monitoring.[2] [3]

3.2 Enterprise AI value beyond simple ROI

Enterprise AI value frameworks increasingly argue that value should not be reduced to time saved or direct labor substitution. Time multiplication, skill multiplication, strategic capability gain, and differentiated stakeholder value all matter. This literature supports a multidimensional Value Office because financial, operational, strategic, and human value often diverge unless measurement logic is made explicit from the outset.[4] [14]

3.3 Observability, explainability, and trajectory transparency

Current work on AI observability argues that consequential agentic systems require end-to-end traceability across prompts, models, tools, memory, and actions. The object of explanation is not just the final answer, but the trajectory by which the answer was reached, including plans, tool invocations, evidence references, retries, and human interventions. This is directly relevant to roundtrip value because recognized value depends on reconstructable execution history, not only on output quality.[5] [6] [10]

3.4 Meaningful human control and governed autonomy

Human control research suggests that oversight must be designed according to purpose, such as legality, accountability, safety, or fairness, rather than added as a generic safeguard. In process environments, this means that intervention points, escalation rights, reversibility rules, and approval logic must be designed as part of the operating model. Value that cannot be governed into recognizability is not enterprise-grade value.[9]

3.5 Multicriteria prioritization and automation portfolios

Prioritization research across business process management, RPA selection, and AI use case management shows that strong portfolios do not optimize only for one metric. Instead, they balance value, pain, feasibility, governance fit, and strategic alignment. This

supports the use of a minimal-information triage layer before detailed validation, particularly when candidate volumes are high and information quality is initially low.[7] [8] [11] [12]

3.6 Governed execution optimization

A more recent body of thinking extends token-centered efficiency into a broader optimization problem that includes execution topology, intervention burden, and recognized value. The key idea is that a token-cheap run may still be expensive once failure, escalation, delay, cleanup, or compliance exposure are considered. This motivates the distinction between **pennywise** optimization of direct inference spend and **poundwise** optimization of total valid completion cost.[13]

4. Roundtrip Value as a Closed-Loop Architecture

The updated roundtrip model defines the Value Office as a closed-loop institutional capability rather than as a reporting office. The loop connects seven stages: intake and framing, triage prioritization, focused validation, controlled deployment, runtime supervision, value recognition, and portfolio feedback.

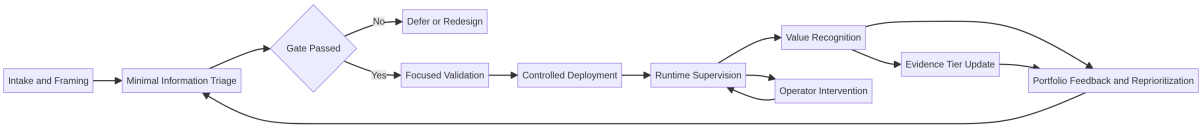

Figure 1

The roundtrip value operating loop connects intake, prioritization, validation, deployment, runtime supervision, value recognition, and portfolio feedback.

The conceptual novelty of the updated model is that it separates two analytical choices that were previously too loosely coupled. The first is a **portfolio choice**, namely whether a process candidate should move forward in the backlog. The second is an **execution choice**, namely which governed pattern should be used for a specific class of work or a live case. These two choices are linked by shared measurement logic, but they are not identical. A

process may deserve prioritization while still requiring different runtime patterns across segments or states.

5. A Four-Layer Economics Model for Roundtrip Governance

The updated framework integrates the earlier Value Office structure with a four-layer economics view.

Table 1

Summary of data

Layer	Core question	Representative constructs	Primary source systems
Token economics	What is the direct computational cost-efficiency of execution	token count, model cost, retries, budget negotiation	LLM gateway, billing logs
Pattern economics	Which execution topology yields the best valid-completion economics	agent count, verifier depth, critic intensity, escalation logic, memory breadth	orchestration traces, runtime events
Governance economics	What control burden and admissibility conditions attach to a pattern	approvals, policy checks, evidence requirements, reversibility, auditability	workflow logs, policy engines, audit systems
Recognized value economics	What value survives adoption, evidence, intervention, and governance correction	accepted FTE release, validated cycle-time reduction, recognized quality gains	process mining, finance review, Value Office ledger

This layered model preserves the earlier roundtrip architecture but gives it a stronger formal basis. Token economics remains a foundational sublayer and should reuse the cost logic already established in TOKENOMICS. Pattern economics decides which topology should be selected. Governance economics determines which topologies are admissible and

what burdens they impose. Recognized value economics determines which observed benefits count as defensible value.[13]

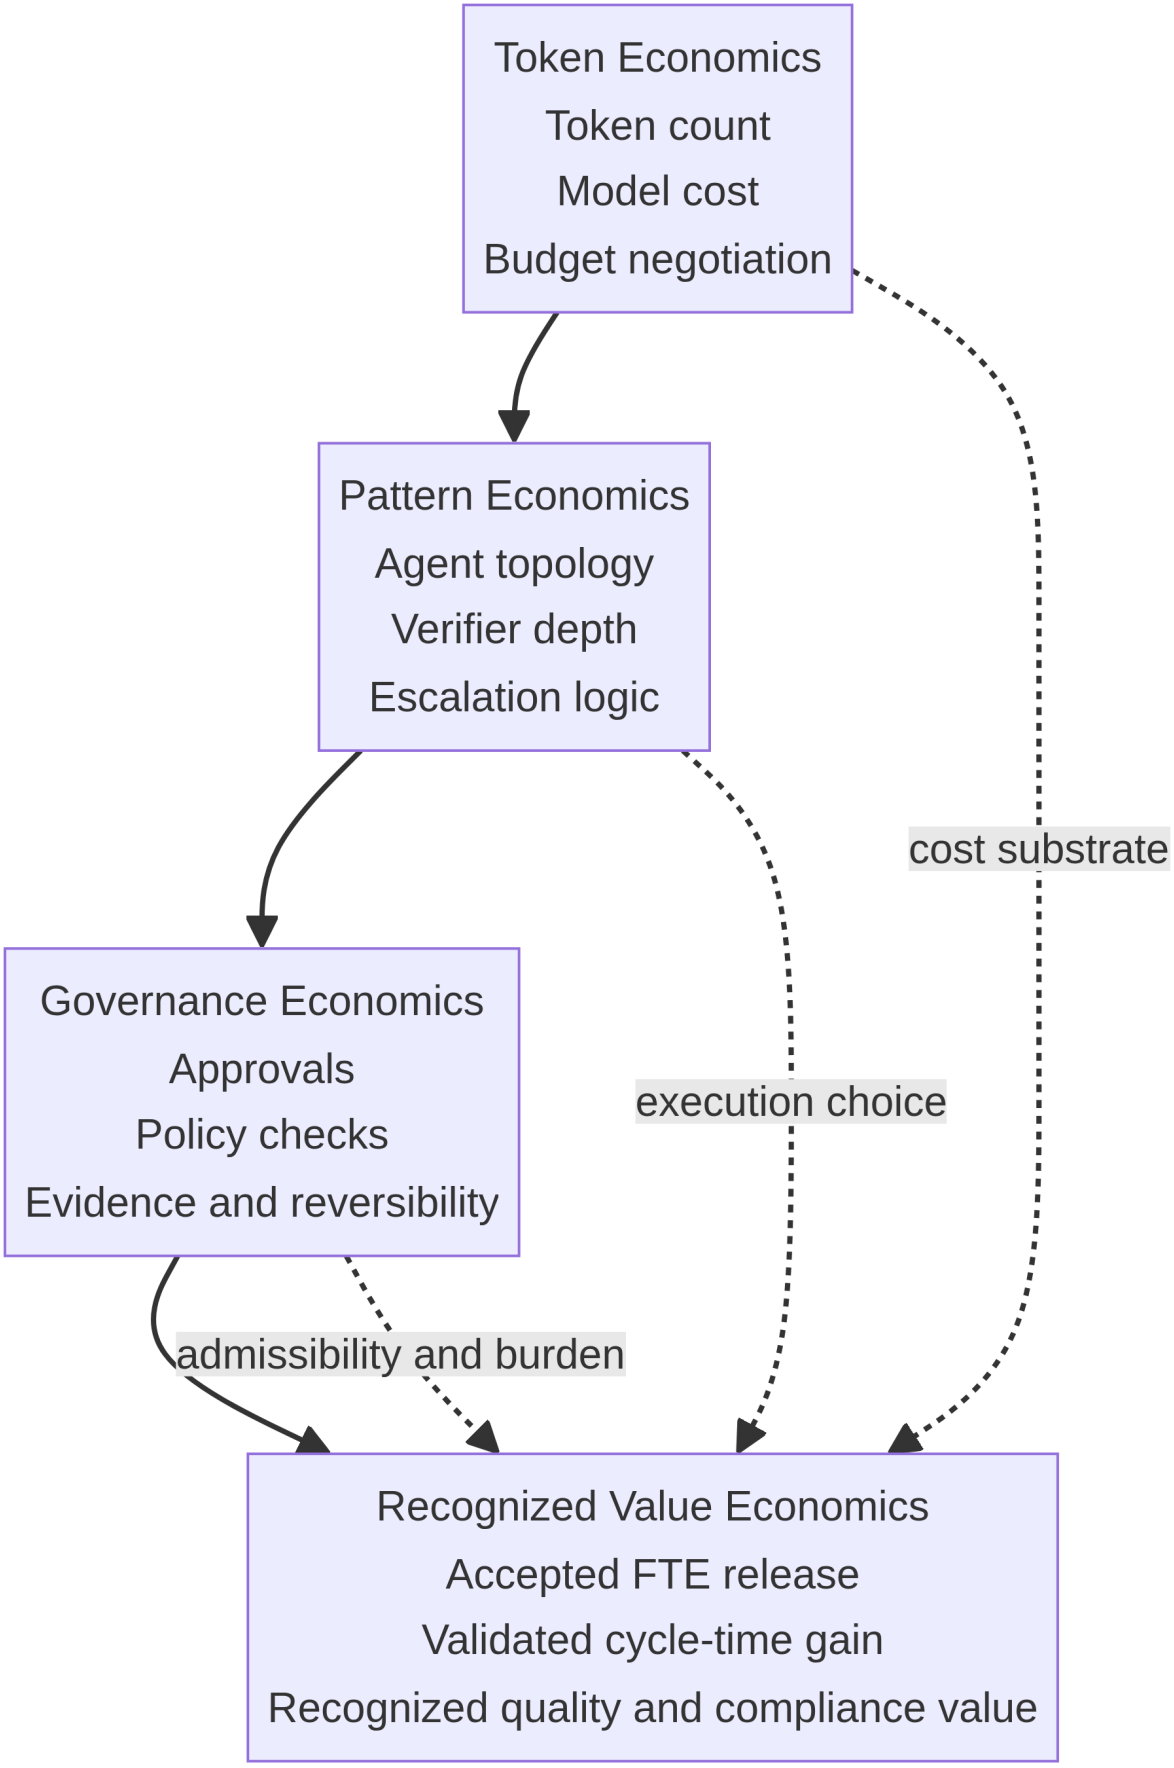
Token Economics
Token count
Model cost
Budget negotiation
Pattern Economics
Agent topology
Verifier depth
Escalation logic
Governance Economics
Approvals
Policy checks
Evidence and reversibility
cost substrate
execution choice
admissibility and burden
Recognized Value Economics
Accepted FTE release
Validated cycle-time gain
Recognized quality and compliance value

Figure 2

The four-layer economics model links token efficiency, pattern choice, governance burden, and recognized value in the updated roundtrip framework.

6. Measurement Model Before the Main Equations

The measurement-first rule is simple: if a construct materially affects prioritization, ranking, admissibility, replanning, or recognition, it must be observable, estimable, derived, or policy-defined.

Table 2

Summary of data

Variable	Definition	Class	Data source	Unit
`Freq`	Case frequency	Observable	workflow logs	cases/month
`Eff_manual`	Current manual effort	Estimated or observable	time study, task logs	minutes/case
`Pain`	Current operational pain	Estimated	intake rubric	1–5
`T_case`	Tokens per case	Observable	LLM gateway	tokens/case
`C_model`	Model cost per case	Derived	token logs and price tables	currency/case
`C_tool`	Tool and API cost per case	Observable	API logs, vendor invoices	currency/case
`C_orch`	Orchestration overhead cost	Derived	runtime traces, infra allocation	currency/case
`L_e2e`	End-to-end latency	Observable	case timestamps	minutes/case
`R_retry`	Retry or replan count	Observable	runtime events	count/case
`I_count`	Intervention events per case	Observable	operator workbench	count/case
`I_min`	Intervention minutes	Observable	case handling logs	minutes/case
`A_count`	Approval count	Observable	workflow engine	count/case
`P_check`	Policy checks	Observable	policy engine	count/case
`E_score`	Evidence completeness score	Estimated or derived	evidence rubric	0–1
`S_valid`	Valid completion probability	Derived		0–1

			accepted completions / total cases	
`CT_red`	Realized cycle-time reduction	Derived	baseline vs post-state	percentage
`ER_red`	Realized error reduction	Derived	baseline vs post-state	percentage
`FTE_rel`	Released FTE minutes or hours	Derived	time studies, queue load, staffing analysis	minutes/case or hours/month
`Comp_inc`	Compliance incident rate	Derived	incident systems	0–1
`Rev_score`	Reversibility score	Policy-defined or estimated	control architecture review	0–1
`Adopt`	Adoption rate	Derived	usage logs, override rate	0–1
`RV_state`	Recognized value approval state	Policy-defined	Value Office decision register	binary or tiered

This table matters because it prevents conceptual elegance from outrunning operational measurability. Evidence completeness, reversibility, and recognized value approval are not raw natural quantities, but they can still be modeled credibly if their rubric and governance ownership are explicit.

7. Work Package 1 as Value Taxonomy and Recognition Logic

WP1 defines what counts as value, how it is evidenced, and when it may be recognized. The updated version retains the earlier value families while aligning them to the new mathematics.

Table 3

Summary of data

Value family	Definition	Typical evidence
Released work	Human effort removed from the process	touch-time comparison, task logs, manager validation
Time value	Reduced elapsed time, queue time, or SLA exposure	process events, queue metrics
Quality value	Reduced errors, rework, and leakage	defect records, correction volumes
Compliance value	Reduced expected or observed governance loss	policy events, audit records, control assessments
Throughput value	More completed work without proportionate added labor	completion counts, backlog changes
Capability value	Ability to handle more complex or scarcer work	expert dependency change, delayed-work avoidance

A crucial distinction is retained among **hours removed**, **capacity released**, and **financially captured FTE**. Effort reduction does not automatically imply captured economic value. The updated framework therefore treats FTE release as the most important value family, but only when supported by evidence and recognition rules.

The revised value states are defined as follows.

Table 4

Summary of data

Value state	Definition	Required evidence	Validator
Claimed value	projected benefit before deployment	business case assumptions, benchmark claims	sponsor
Realized value	measured operational gain after deployment	baseline comparison, stabilized process evidence	process owner
Recognized value	realized value accepted after evidence, intervention, and governance correction	validated metrics, evidence packet, control review	Value Office, finance, risk

8. Work Package 2 as Telemetry and Evidence Spine

WP2 creates the case-centric evidence backbone that allows the roundtrip loop to function empirically. Every consequential case must be joinable across business events, agent traces, policy events, and human intervention records.

Table 5

Summary of data

Event class	Typical contents	Function
Process events	start, complete, queue, SLA state	reconstruct business flow
Agent events	plan, tool use, retry, retrieval, verifier result	reconstruct execution trajectory
Policy events	checks, blocks, approvals, rule outcomes	prove bounded autonomy
Human intervention events	approve, correct, enrich, reroute, suspend, escalate	measure supervisory burden
Value events	evidence-tier change, value posting, recognition decision	link execution to value acceptance

The **Minimal Explanation Packet** remains central in the updated design. For high-consequence cases, the packet should contain case summary, execution trajectory, evidence sources, policy context, intervention history, confidence indicators, final outcome, and version context. This packet unifies observability, explainability, auditing, and value defensibility.[5] [6] [9] [10]

9. Work Package 3 as Dashboard and Control Environment

WP3 turns telemetry and calculations into role-specific decision environments. The updated version adds stronger economic and operational differentiation across users.

Table 6

Summary of data

Dashboard	Core question	Representative metrics
Executive Value Dashboard	Are we creating defensible value	recognized value, released FTE, evidence maturity, portfolio hit rate
Process Performance Dashboard	Is the process materially improving	cycle time, backlog, quality, manual touches
Agent Operations Dashboard	Are digital workers stable and economical	valid completion, retries, token cost, orchestration cost, intervention burden
Control and Compliance Dashboard	Are we within governance bounds	policy breaches, overrides, evidence gaps, reversibility issues
Operator Workbench	Which case requires action and why	queue rank, explanation packet, trace timeline, recommended action

The operator is not merely an approver but a **digital process controller**. Cases should be ranked by business criticality, value at risk, SLA exposure, and control severity. Standardized actions such as approve, correct, enrich, reroute, suspend, and escalate form both an intervention mechanism and a learning signal for the optimizer.

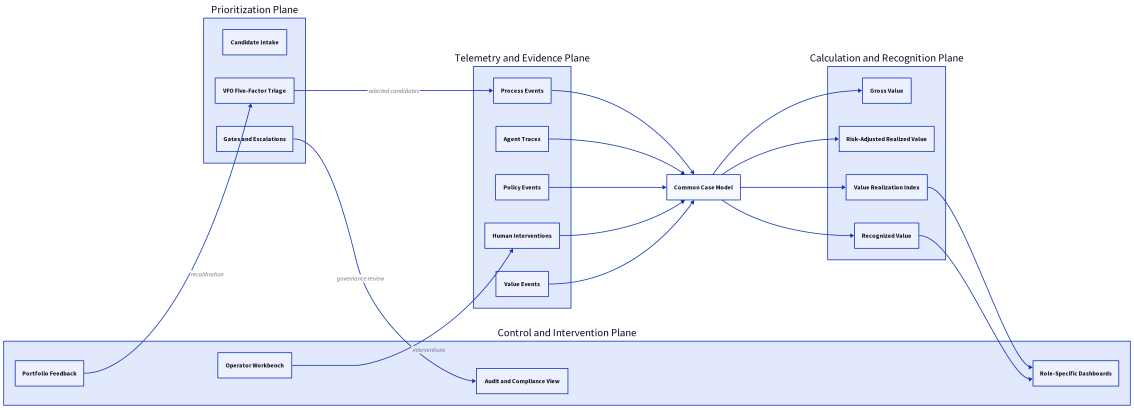

Figure 3

The roundtrip Value Office architecture links prioritization, telemetry, calculation, dashboards, and control through a common case model.

10. Minimal-Information Prioritization Model

Prioritization is the front end of roundtrip value. The updated model retains the low-information design while linking it more explicitly to recognized value and governed execution.

Table 7

Summary of data

Dimension	Core question	Why it matters
`V` Value potential	If this works, how much value could it create	aligns with FTE release, cycle time, quality, compliance, throughput
`P` Process pain and manuality	How painful and manually intensive is the process now	captures urgency and burden
`F` Feasibility	Can this be agentified with current data, tooling, and process stability	prevents value-only bias
`G` Governance fit	Can this operate safely under required controls	prevents unsafe selection
`S` Strategic alignment and ownership	Does it matter strategically and is it owned	improves sponsorship and realization
Gate	Should the candidate be blocked, deferred, or escalated before scoring	handles hard constraints

The practical triage formula remains:

$$PriorityScore = 100 \times \frac{0.35V + 0.20P + 0.20F + 0.15G + 0.10S}{5}$$

The value-potential factor is itself decomposed to align with Value Office priorities:

$$V = 0.40FTE + 0.20CT + 0.15ER + 0.15CR + 0.10TH$$

where `FTE` is released FTE potential, `CT` is expected cycle-time reduction, `ER` is expected error reduction, `CR` is expected compliance-risk reduction, and `TH` is expected

throughput effect. This makes the portfolio front end explicitly consistent with downstream value recognition.

11. Runtime Pattern Optimization

Prioritization alone does not determine how a selected process should run. The updated paper therefore introduces a second optimization layer for governed execution patterns.

A candidate execution pattern is represented as:

$$p = (n_a, v, c, h, o, m, e)$$

where `n_a` is the number of agents, `v` verifier depth, `c` critic intensity, `h` human gating level, `o` ontology or policy-enforcement level, `m` memory breadth, and `e` escalation threshold.

The master ranking function for admissible patterns is:

$$Score(p) = G(p) \times \frac{RV(p) \times S(p) \times CX(p)^{\alpha}}{C_{total}(p)^{\beta}}$$

where `G(p)` is the admissibility gate, `RV(p)` recognized value, `S(p)` expected valid-completion probability, `CX(p)` complexity successfully handled or a measurable complexity proxy, and `C_total(p)` total enterprise cost.

Total enterprise cost is defined as:

$$C_{total}(p) = C_{token}(p) + C_{tool}(p) + C_{orch}(p) + C_{human}(p) + C_{delay}(p) + C_{failure}(p)$$

Recognized value is defined as:

$$RV(p) = GrossValue(p) \times AdoptionFactor(p) \times EvidenceFactor(p) \times GovernanceFactor(p) \times InterventionPenalty(p)$$

with an intervention penalty such as:

$$InterventionPenalty(p) = e^{-\lambda I(p)}$$

This is the main mathematical integration introduced into the roundtrip paper. It makes explicit that a runtime pattern is preferable not because it is architecturally elegant, but because it maximizes recognized value per total enterprise cost under admissibility constraints.

11.1 Admissibility constraints

Certain patterns must be excluded regardless of local efficiency. The optimizer therefore applies hard constraints such as:

$$Risk(p) \leq R_{max}$$

$$Latency(p) \leq L_{max}$$

$$Intervention(p) \leq I_{max}$$

$$Budget(p) \leq B_{remaining}$$

$$Governance(p) = approved$$

$$Reversibility(p) \geq T_{min}$$

A pattern with low token spend but poor reversibility or insufficient control fit is therefore inadmissible, not merely suboptimal.

11.2 Runtime replanning

The active pattern at time `t` can be selected dynamically as:

$$p^*(t) = \arg\max_{p \in \mathcal{A}(t)} \mathbb{E}[Score(p) \mid \mathcal{I}_t]$$

where `\mathcal{A}(t)` is the currently admissible pattern set and `\mathcal{I}_t` the information available at runtime. Replanning triggers include retry spikes, latency drift, evidence-score deterioration, intervention spikes, contradiction rises, and governance-state tightening.

12. Formalizing Value Realization in Roundtrip Form

The earlier Value Realization Index remains useful, but it now sits inside a richer analytical environment:

$$VRI = 100 \times (0.35S_{fte} + 0.20S_{cycle} + 0.15S_{quality} + 0.15S_{compliance} + 0.10S_{throughput} + 0.05S_{capability})$$

This index reflects the strategic preference for released FTE while still preserving substantial weight for time, quality, and compliance outcomes. However, the updated framework makes one further move. It explicitly links **priority formation**, **runtime pattern choice**, and **recognized value**.

Conceptually, a strong candidate should not only score well in triage but also demonstrate three downstream characteristics: high evidence maturity, controllable intervention burden, and stable recognized value. If high-priority candidates repeatedly show weak evidence or excessive intervention, the front-end weighting is miscalibrated. The Value Office must then learn from runtime truth, not merely report outcomes.

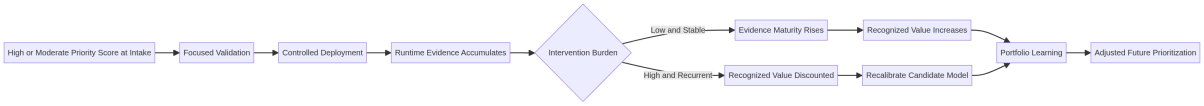

Figure 4

The lifecycle relation among Priority Score, evidence maturity, intervention burden, and recognized value in the roundtrip Value Office model.

13. Pennywise and Poundwise Optimization in the Value Office

The roundtrip framework now distinguishes between two forms of economic reasoning. **Pennywise optimization** minimizes direct run cost, especially token spend and immediate inference expenditure. **Poundwise optimization** minimizes total enterprise cost of valid completion.

$$\text{TotalEnterpriseCost} = \text{DirectRunCost} + \text{HumanCleanupCost} + \text{DelayCost} + \text{FailureCost} + \text{ComplianceExposureCost}$$

This distinction matters for the Value Office because a candidate that looks attractive in a simplistic business case may become unattractive once intervention time, remediation burden, evidence generation, and delayed approval cycles are counted. The Value Office should therefore not mistake cheap inference for cheap operations.

14. Operational Roundtrip Way of Working

The updated roundtrip way of working can be summarized as a seven-step operating rhythm.

Table 8

Summary of data

Step	Main actor	Core output
Intake	process owner	candidate description and minimal-information intake
Triage	Value Office analyst	priority score and gate status
Validation	architect and analyst	refined feasibility, telemetry design, baseline view
Deployment	delivery team	governed implementation and admissible pattern library
Supervision	operator and control functions	intervention decisions, stability and evidence signals
Recognition	Value Office, finance, risk	recognized value decision
Feedback	portfolio lead	reprioritized backlog and updated weighting guidance

This rhythm makes the Value Office a standing portfolio-and-control function rather than a downstream measurement committee. The organization no longer asks only whether the agent succeeded technically. It asks whether the candidate was selected wisely, instrumented adequately, controlled appropriately, and recognized credibly.

15. Platform Implications and Architecture Choices

The current platform landscape suggests at least three implementation patterns. A suite-led approach may use workflow-centric platforms such as UiPath or ServiceNow for orchestration and governance-led operations. A process-intelligence-led approach may use Celonis to anchor baselines and continuous process visibility. A composable approach may combine process intelligence, workflow execution, policy enforcement, agent observability, and business intelligence into a custom multi-plane architecture.[15] [16] [17] [18]

The implication of the updated framework is that no platform should be selected solely on the basis of orchestration features or agent demos. The architecture must support case identity, telemetry joinability, evidence completeness, intervention logging, cost allocation, and value recognition.

16. Research Hypotheses and Empirical Agenda

Table 9

Summary of data

Hypothesis	Statement
H1	A roundtrip Value Office with multidimensional recognition rules produces more accurate value reporting than conventional ROI tracking
H2	Minimal-information prioritization with governance gates selects candidates with higher recognized value than ad hoc selection
H3	Case-centric observability and explanation packets reduce intervention time and improve audit readiness
H4	Runtime pattern optimization improves recognized value per total enterprise cost relative to fixed-pattern deployment
H5	Feedback from recognized value into reprioritization improves portfolio hit rate over time
H6	Purpose-specific intervention design improves control quality while reducing unnecessary human touches

A mixed-method program is appropriate. Quantitative work should compare baseline and post-automation outcomes across several processes and patterns, while qualitative work should examine trust calibration, operator usability, and governance acceptance. The first empirical setting should consist of Zone III processes with clear case identity, meaningful manual effort, high exception burden, and compliance relevance.[1] [9]

17. Discussion

The updated paper strengthens the original roundtrip claim in three ways. First, it shows that the Value Office is not only a measurement office but also a governed optimization office. Second, it integrates prioritization and runtime selection into one evidence logic rather than leaving them institutionally fragmented. Third, it introduces a more mathematically explicit basis for understanding why some automation initiatives appear attractive at intake but fail to create recognized value at scale.

A second implication is conceptual. Observability alone is insufficient. It can reveal what the system did, but not whether the process should have been prioritized, whether the chosen pattern was admissible, or whether the resulting benefit is economically recognizable. The roundtrip framework therefore adds a portfolio-learning layer that connects runtime truth back to upstream selection logic.

18. Limitations

This paper remains a design proposition grounded in literature, earlier framework work, and cross-paper synthesis rather than in broad multi-enterprise validation. Several limitations are material. Parameter estimation for delay cost, failure cost, and recognized value may be unstable in early pilots. Proxy variables such as evidence completeness can become performative if poorly governed. Intervention burden is measurable but may not be comparable across teams with different work structures. Governance thresholds vary by sector and over time. Finally, the constructs *roundtrip value* and *patternomics* are introduced here as design constructs rather than mature canonical fields. Their usefulness must therefore be tested empirically rather than assumed.

19. Conclusion

This paper has updated the roundtrip Value Office architecture by integrating newer research on measurable execution optimization, governed autonomy, and recognized value economics. The central insight is that organizations need a capability that connects candidate

selection, execution telemetry, explanation, intervention, runtime pattern choice, and value recognition into one closed loop. Without such a capability, enterprises risk mistaking cheap inference for cheap operations, attractive ideas for viable portfolio choices, and technical activity for defensible value.

The updated framework provides a structured answer. WP1 defines the vocabulary and rules of value recognition. WP2 builds the case-centric evidence spine. WP3 creates role-specific decision environments. The prioritization plane governs which initiatives should move forward. The runtime pattern optimizer governs how selected initiatives should execute under real constraints. Together, these mechanisms turn the Value Office into a learning system for governed agentic automation.

References

[1] van Hurne, M. H. K. ResearchGate profile and publications. ResearchGate. https://www.researchgate.net/profile/Marco-Hk-Van-Hurne

[2] Advances in Process Optimization A Comprehensive Survey of Process Mining Predictive Process Monitoring and Process Aware Recommender Systems. arXiv. https://arxiv.org/html/2301.10398v2

[3] Why Process Mining Is Key To Continuous Automation Value Realization. UiPath. https://www.uipath.com/blog/automation/process-mining-key-continuous-automation-value-realization

[4] The Capability Leap A framework for measuring the enterprise value of AI tools. Thomson Reuters. https://www.thomsonreuters.com/en-us/posts/innovation/the-capability-leap-a-framework-for-measuring-the-enterprise-value-of-ai-tools/

[5] Why observability is essential for AI agents. IBM. https://www.ibm.com/think/insights/ai-agent-observability

[6] AI agent observability The measured leap. Deloitte. https://www.deloitte.com/us/en/services/consulting/articles/ai-agent-observability-human-in-the-loop.html

[7] Guide to identifying and investing in AI projects. CSIRO. https://www.csiro.au/-/media/D61/Investing-in-AI-projects-guide/Guide-to-identifying-and-investing-in-AI-projects-CSIRO-2025.pdf

[8] AI use case management. EY. https://www.ey.com/en_us/cro-risk/ai-use-case-management

[9] On the purpose of meaningful human control of AI. PubMed Central. https://pmc.ncbi.nlm.nih.gov/articles/PMC9868906/

[10] Transparency in Agentic AI A Survey of Interpretability Explainability and Governance. EngrXiv. https://engrxiv.org/preprint/view/6451

[11] A Method for Prioritizing Robotic Process Automation Projects with the Application of Multicriteria Decision Making and Process Mining. Springer. https://link.springer.com/chapter/10.1007/978-3-031-77723-3_1

[12] Business process prioritization criteria a case study in the financial market. SciELO. https://www.scielo.br/j/rmj/a/qGGJmxQqpPDhbT3zbmyy8QK/?lang=en

[13] van Hurne, M. Measured Patternomics for Governed Agentic Execution. Unpublished working manuscript in the same research family as TOKENOMICS, available from the author.

[14] The AI value paradox measuring what matters not just whats easy. CGI. https://www.cgi.com/en/blog/artificial-intelligence/ai-value-paradox-measuring-what-matters-not-just-whats-easy

[15] UiPath Agentic Automation. UiPath. https://www.uipath.com/platform/agentic-automation

[16] ServiceNow AI Agents. ServiceNow. https://www.servicenow.com/products/ai-agents.html

[17] Celonis Process Intelligence. Celonis. https://www.celonis.com/

[18] AI Agent Observability with Langfuse. Langfuse. https://langfuse.com/blog/2024-07-ai-agent-observability-with-langfuse