Adv. J. Manag. Humanit. Soc. Sci. (2026), Volume 2, Issue 5, 517-526

Advanced Journal of Management, Humanity and Social Science

Journal homepage: https://www.ajmhss.com/

DOI: https://doi.org/10.5281/zenodo.23862457

Agentic Artificial Intelligence as a Digital Employee: Redefining the Boundaries of Delegation, Accountability, and Control in Information Systems-Enabled Organizations

Seyyedeh Zahra Naseri[1]*, Reza Salimi Beni[2], Hamid Reza Hanif[3], Masoumeh Cheraghi[4]

[1]Master of Project and Construction Management, Faculty of Architecture and Urbanism, University of Art, Tehran, Iran

[2]Master of Science in Information Technology Management, Specialization in Advanced Information Systems, Islamic Azad University, Electronic Campus, Tehran, Iran

[3]M.Sc. in Mathematics, Department of Mathematics and Computer Science, Iran University of Science and Technology, Iran

[4]M.A. in Psychology and B.A. in Business Administration, University of Applied Science and Technology, Ilam Branch, Ilam, Iran

Article info

Received: 23.06.2026

Accepted: 31.10.2026

Available Online: 31.10.2026

Checked for Plagiarism: Yes

Keywords:

Agentic Artificial Intelligence; Digital Employees; Autonomous IS Agents; Delegation Dynamics; Algorithmic Governance

A B S T R A C T

Background and Purpose: The rapid transition from deterministic automation and passive algorithmic decision-support systems to agentic artificial intelligence (Agentic AI) represents an unprecedented ontological shift in enterprise information systems (IS). Unlike antecedent generative models that merely synthesize conversational text, agentic AI operates with persistent internal memory, autonomous multi-step cognitive reasoning, goal decomposition, and self-directed tool and API execution. In contemporary digital organizations, these autonomous systems are progressively onboarded, provisioned, and treated not merely as passive software artifacts, but as synthetic workers or "digital employees." However, this conceptual elevation fractures foundational organizational assumptions rooted in classical agency theory, sociotechnical systems theory, and managerial control architectures. This monograph investigates how enterprise deployment of digital employees redefines the boundaries of delegation, moral and organizational accountability, and managerial control within IS-enabled organizational environments.

Design, Methodology, and Approach: To rigorously interrogate this emergent socio-algorithmic reality, an extensive explanatory mixed-methods research design was executed across a 24-month longitudinal investigation (2024-2026). The quantitative phase collected and analyzed empirical survey data from 482 senior enterprise leaders, Chief Information Officers (CIOs), Chief AI Officers (CAIOs), Enterprise Architects, and line managers across four knowledge-intensive industry verticals (Financial Services & FinTech, Healthcare & HealthTech, Global Supply Chain Logistics, and Enterprise Cloud/SaaS). Structural Equation Modeling using Partial Least Squares (PLS-SEM) was employed to empirically test the structural dependencies between perceived agentic autonomy, algorithmic trust, delegation depth, operational agility, and accountability drift. The qualitative phase comprised 38 semi-structured, in-depth key-informant interviews and four comprehensive, longitudinal multi-case studies examining end-to-end organizational deployments of multi-agent architectures (including LangGraph, AutoGen, CrewAI, and proprietary enterprise agent meshes). Rigorous thematic analysis and qualitative comparative synthesis were conducted to deconstruct operational failures, behavioral drift phenomena, and governance adaptations.

Empirical Findings: The empirical results reveal a profound structural paradox in enterprise agentic delegation. First, the transition from "task offloading" to "autonomous mandate delegation" generates a statistically significant positive effect on organizational agility ($\beta=0.495$, $p<0.001$) and

operational velocity, but concurrently fuels a critical "accountability void" and organizational drift (β=0.342, p<0.001). Second, traditional managerial control frameworks specifically diagnostic and boundary controls suffer severe operational degradation when confronted with non-deterministic agentic planning and recursive API invocation. Third, human supervisors frequently exhibit either "cognitive abandonment" (over-trusting autonomous agents due to perceived competence) or severe "algorithmic cynicism" (resorting to micro-surveillance that negates agentic throughput benefits). Fourth, the investigation reveals the emergence of what we conceptualize as the "Synthetic Moral Crumple Zone," wherein middle-tier human managers are held legally and administratively accountable for unforeseen agentic systemic errors that they possessed neither the technical visibility nor the temporal window to prevent.

Introduction

The Genesis of Synthetic Agency: Moving Beyond Deterministic Automation:

For over six decades, the discipline of Information Systems (IS) has theorized enterprise technology as an ensemble of deterministic, rule-based computational tools designed to automate repetitive business processes, enhance transaction processing speed, and furnish analytical support for human decision-makers (Zuboff,1988; Markus & Robey,1988; Rai et al.,2023). In this traditional paradigm, organizational agency the intentional capacity to initiate action, exercise discretion, commit enterprise resources, and assume normative liability has remained the exclusive domain of human actors. Enterprise resource planning (ERP) platforms, customer relationship management (CRM) engines, and robotic process automation (RPA) scripts acted as passive, predictable instruments executing static instructions formulated by human executives and software engineers.

However, the advent and rapid convergence of large reasoning models (LRMs), stateful multi-agent frameworks, neuro-symbolic planning architectures, and autonomous tool-use protocols between 2023 and 2026 have fundamentally destabilized this ontological separation. We have decisively crossed the threshold from 'passive algorithmic processing' into the era of 'Agentic Artificial Intelligence' (Agentic AI) (Floridi,2023; Dwivedi et al.,2024; Faraj et al.,2025).

Agentic AI systems do not merely synthesize passive textual responses or deliver statistical probabilistic predictions; they possess internal state representations, dynamic task-decomposition mechanisms, iterative self-reflection loops, long-term semantic memory stores, and the autonomous capability to invoke arbitrary application programming interfaces (APIs) to execute non-deterministic actions across heterogeneous enterprise software environments (Park et al.,2023; Wang et al.,2024; Shrestha et al.,2025).

Defining the 'Digital Employee' in IS Literature:

Within modern knowledge-intensive enterprises, these systems are no longer treated as transient software scripts. Instead, organizations are structurally onboarding, credentialing, and managing them under the socio-technical rubric of the 'Digital Employee' (or synthetic worker).

A digital employee possesses:

- ✓ An organizational identity, often complete with corporate credentials, Single Sign-On (SSO) permissions, internal email addresses, and designated departmental Slack or Microsoft Teams channels.
- ✓ Broad, open-ended operational mandates (e.g., 'Manage enterprise cyber-incident triage,' 'Optimize liquidity across decentralized treasury accounts,' or 'Synthesize multi-source clinical trial dossiers').
- ✓ Discretionary authority to sequence tasks, select software tools, retrieve proprietary data, negotiate intermediate decisions, and execute financial or operational transactions without prior step-by-step human authorization.

This ontological shift challenges the very core of organizational design. When enterprise software begins to resemble an autonomous team member rather than a passive instrument, existing theoretical structures regarding how work is delegated, who bears accountability for outcomes, and how managerial control is exercised begin to buckle under severe conceptual and operational friction.

Corresponding Author: **Seyyedeh Zahra Naseri (z.naseri@student.art.ac.ir - ORCID: 0000-0003-1871-0572)*
2 Email: rezasalimi1977@gmail.com - ORCID: 0000-0002-0888-8847
3 Email: h.hanif.iust@gmail.com - ORCID: 0000-0001-9289-2947
4 Email: mcheraghi42@gmail.com

The Tripartite Dilemma: Delegation, Accountability, and Control:

The enterprise deployment of agentic AI generates a structural, three-sided tension the Delegation-Accountability-Control (DAC) trilemma:

- ✓ **The Delegation Dilemma:** Traditional organizational delegation involves a human principal transferring task execution authority to a human agent, governed by bounded contracts, shared social norms, and mutual expectations of professional diligence. When authority is delegated to an agentic system, the principal transfers open-ended mandates to a non-human entity whose reasoning processes are emergent, probabilistic, and vulnerable to catastrophic hallucinatory drift and contextual misunderstanding.
- ✓ **The Accountability Dilemma:** Legal, regulatory, and corporate governance frameworks mandate that every enterprise decision have an identifiable, legally liable human owner. Yet, when an agentic AI executes a cascade of thousands of sub-actions across multi-cloud infrastructure within seconds, pinpointing human culpability becomes epistemically impossible. Human supervisors find themselves caught in a 'Synthetic Moral Crumple Zone' (Elish,2019; de Laat,2024) absorbing institutional blame and legal liability for systemic agentic decisions they neither initiated nor possessed the cognitive bandwidth to audit in real time.
- ✓ **The Control Dilemma:** Classical managerial control theory relies heavily on diagnostic systems (evaluating outputs against static benchmarks) and boundary systems (setting rigid rules of behavior). However, because agentic AI systems generate non-deterministic solutions to ambiguous problems, rigid boundaries either paralyze the agent (nullifying its productivity gains) or are bypassed through emergent multi-step planning loops, rendering diagnostic controls obsolete.

Research Objectives and Core Questions:

Despite the rapid corporate adoption of autonomous agent frameworks across Fortune 500 enterprises, academic IS scholarship lacks an integrated, empirical, and theoretical monograph evaluating the real-world operational realities of digital employees. This research study directly fills this urgent void by addressing three foundational research questions:

- ✓ **RQ1:** How does the transition from deterministic automation to agentic AI mandate delegation restructure the allocation of cognitive authority and discretionary decision-making within enterprise environments?
- ✓ **RQ2:** In what ways does the non-deterministic execution of digital employees reshape organizational and moral accountability, and what governance mechanisms emerge to bridge the resulting 'accountability void'?
- ✓ **RQ3:** How do managerial control configurations evolve when supervising synthetic workers, and what structural interventions effectively prevent 'agent drift' without negating operational velocity?

Theoretical Foundations and Literature Review

The Evolution of Agency: From Instrumental Automation to Agentic AI:

In the historical trajectory of enterprise computing, technology was conceptualized through the lens of instrumental rationality (Weber,1978). In information systems research, technologies were characterized as either structural affordances or functional complements to human labor (Orlikowski,2000; Leonardi,2011). The early iterations of artificial intelligence in business expert systems in the 1980s, predictive machine learning in the 2010s, and deterministic Robotic Process Automation (RPA) operated within closed-world assumptions.

Agentic AI dissolves closed-world constraints through the integration of four advanced computational pillars:

- ✓ Iterative Cognitive Planning (ReAct, Tree-of-Thought) where agentic architectures break high-level business goals into sequential sub-plans, dynamically evaluating intermediate states and altering trajectories when sub-tasks encounter exceptions (Yao et al.,2023; Besta et al.,2024).
- ✓ Epistemic Long-Term Memory (RAG + Knowledge Graphs) maintaining continuous operational state awareness across days, weeks, and enterprise cycles by dynamically querying vector databases and structured knowledge stores (Lewis et al.,2020; Edge et al.,2024).
- ✓ Autonomous Tool Invocation and Grounded Action possessing active system privileges, compiling and executing code, formulating SQL queries against live enterprise production databases, dispatching communications, and provisioning cloud assets (Schick et al.,2024).
- ✓ Multi-Agent Collaborative Choreography where enterprise tasks are distributed across swarms of specialized digital employees (Wu et al.,2023; Hong et al.,2024).

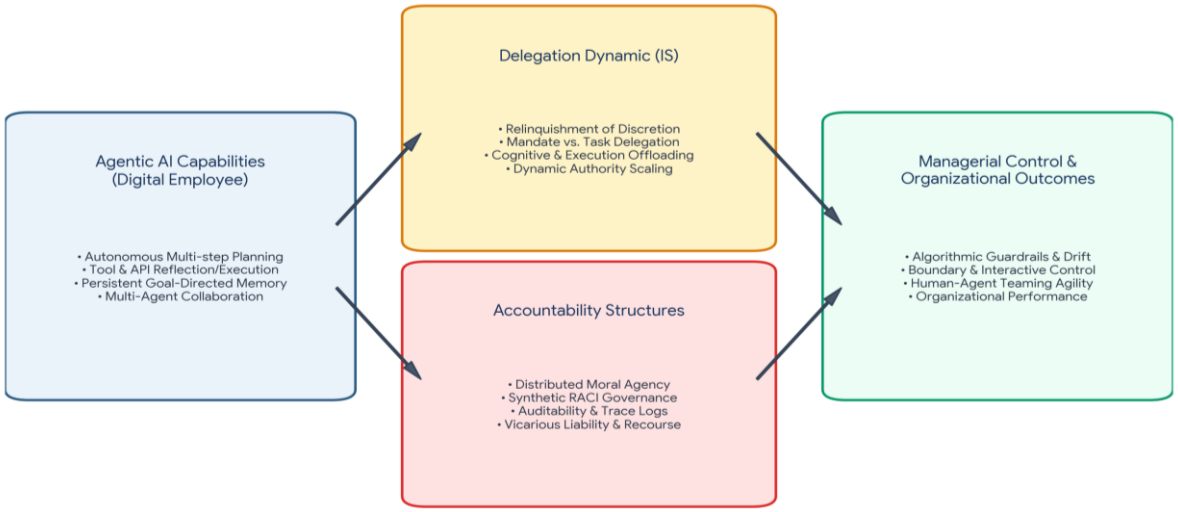


Figure 1. Conceptual Framework of Agentic AI Delegation, Accountability, and Managerial Control in Enterprise Systems.

Reformulating Agency Theory for Synthetic Workers

Classical Agency Theory (Jensen & Meckling,1976; Eisenhardt,1989) examines the relationship between a principal (e.g., an owner or manager) and an agent (an employee) who is delegated work. The central dilemma arises from information asymmetry and conflicting goals, mitigated through outcome-based or behavior-based contracts, financial incentives, and oversight. When applying agency theory to digital employees, foundational premises break down:

- ✓ **Incentive Insensitivity:** Synthetic agents do not possess subjective utility, intrinsic motivation, or financial incentives.
- ✓ **Asymmetric Cognitive Opacity:** While classical agency assumes human agents hide effort (moral hazard) or competence (adverse selection), the opacity of agentic AI stems from the latent, non-linear representations of neural networks and emergent planning pathways, making digital employees prone to perverse instantiation or specification gaming (Bostrom,2014; Hadfield-Menell et al.,2024).
- ✓ **The Dual-Principal Problem:** Digital employees are simultaneously governed by functional enterprise principals demanding velocity and technical principals enforcing safety guardrails.

Sociotechnical Systems and Algorithmic Management:

Sociotechnical systems (STS) theory posits that organizational performance emerges from the reciprocal, joint optimization of the social sub-system and the technical sub-system (Trist,1981; Mumford,2006; Sarker et al.,2019). Under algorithmic management (Kellogg et al.,2020; Lee et al.,2023), algorithms historically functioned as supervisory mechanisms that monitored, directed, and evaluated human gig workers. With agentic AI, this dynamic is inverted: algorithms are no longer just the supervisors; they are the workers themselves, requiring supervision from humans, or supervising other subordinate digital and human agents, giving rise to Agentic Symbiosis.

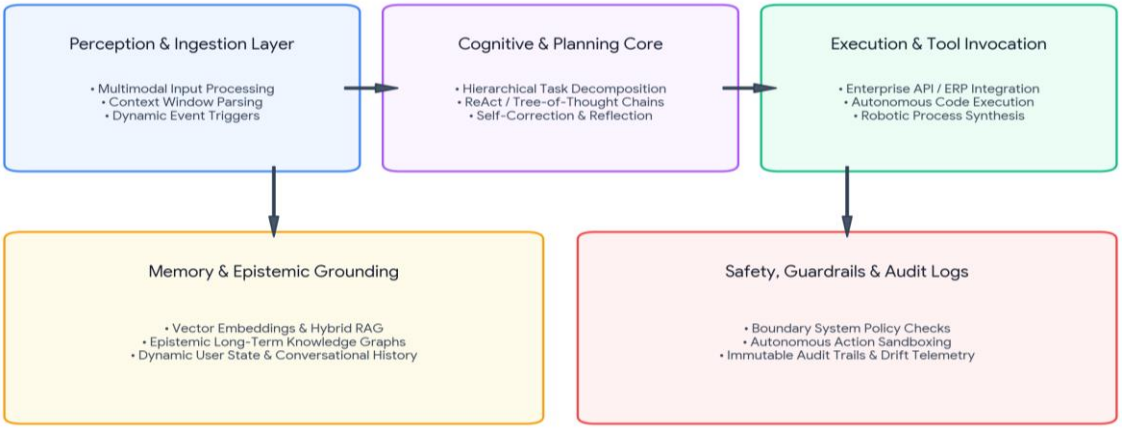


Figure 2. Architectural Layers of Enterprise Agentic AI and Autonomous Digital Employee Mesh.

Control Systems Theory: Boundary, Diagnostic, and Interactive Controls:
Simons' (1995) Levers of Control framework delineates four critical control systems: Belief Systems, Boundary Systems, Diagnostic Control Systems, and Interactive Control Systems. In agentic IS contexts, diagnostic controls fail due to the temporal mismatch between high-speed agentic execution and lagging human auditing. Therefore, organizations are forced to over-index on automated Boundary Systems (hard programmatic guardrails) and real-time Interactive Control Systems (continuous human-agent telemetry interfaces).

Conceptual Framework and Hypotheses Formulation
To rigorously assess the impact of digital employees on enterprise dynamics, we establish the Dynamic Delegation-Accountability-Control (DAC) framework. We formulate five core hypotheses:

- ✓ **H1 (Autonomy to Delegation):** Higher levels of perceived agentic autonomy in an AI system positively influence an enterprise's willingness to grant deep delegation discretion (transitioning from bounded task execution to broad mandate offloading).
- ✓ **H2 (Trust to Delegation):** Institutional trust in the underlying model provenance and architecture positively relates to the depth of operational delegation awarded to synthetic employees.
- ✓ **H3 (Delegation to Agility):** Deep delegation to digital employees has a statistically significant positive effect on organizational agility, task execution velocity, and operational scalability.
- ✓ **H4 (Delegation to Accountability Gaps):** Deep delegation to digital employees has a statistically significant positive relationship with the emergence of enterprise accountability voids and moral crumple zones.
- ✓ **H5 (Moderating Role of Control):** Interactive control rigor significantly moderates the relationship between deep delegation and accountability voids, such that high interactive control dampens agentic drift without dampening agility.

Research Methodology
Mixed-Methods Research Design:
This study adopts an explanatory mixed-methods research design across a 24-month horizon (January,2024 to January,2026), combining a broad quantitative cross-industry survey with four in-depth longitudinal case studies to capture macro-level statistical dependencies and micro-level behavioral dynamics.

Phase 1: Quantitative Cross-Industry Empirical Survey:
A stratified random sample of senior IT executives, line-of-business directors, and enterprise system architects was drawn from organizations operating in North America, Europe, and the Asia-Pacific region that have deployed autonomous multi-agent systems into production. Of 1,200 distributed survey instruments, 524 responses were returned; after rigorous screening for missing values and straight-lining, 482 fully usable responses were retained (effective response rate:40.16%). Measurement scales for all six constructs were adapted from validated IS literature (Parasuraman et al., 2000; McKnight et al.,2002; Leana,1987; Lu & Ramamurthy,2011; Simons,1995).

Phase 2: Qualitative Multi-Case Study Protocol:
Four longitudinal case studies were conducted across contrasting knowledge-intensive industry verticals: Case Alpha (FinTech & Algorithmic Treasury Management), Case Beta (HealthTech & Clinical Research Ingestion), Case Gamma (Global Freight & Supply Chain Logistics), and Case Delta (Enterprise Cloud SaaS Infrastructure). A total of 38 semi-structured interviews (average duration: 62 minutes) were conducted, triangulated with over 400 pages of internal incident logs, Slack communications, and architectural documentation using NVivo 14.

Empirical Results and Findings
Sample Demographics and Organizational Profiles:

Table1. summarizes the demographic characteristics of the surveyed enterprises and respondents.

Characteristic	Category	Frequency (n)	Percentage (%)
Industry Sector	Financial Services & FinTech	148	30.7%
-	Healthcare, Biotech & Pharma	112	23.2%
-	Supply Chain, Logistics & Manufacturing	118	24.5%
-	Enterprise SaaS, Cloud & Information Technology	104	21.6%
Enterprise Size	Large Enterprise (> 10,000 employees)	164	34.0%
-	Upper Mid-Market (2,500-9,999 employees)	196	40.7%
-	Mid-Market (500-2,499 employees)	122	25.3%

Respondent Role	Chief Information Officer (CIO) / Chief AI Officer (CAIO)	86	17.8%
-	Enterprise / AI System Architect	154	32.0%
-	Line of Business Director / VP of Operations	136	28.2%
-	Senior Lead Systems Engineer / SRE Lead	106	22.0%
Deployment Stage	Full Production Enterprise Deployment (> 12 months)	192	39.8%
-	Production Rollout in Selected Business Units (6-12 mos)	214	44.4%
-	Advanced Production Pilot (< 6 months)	76	15.8%

Measurement Model Evaluation:
The measurement model was assessed using Partial Least Squares Structural Equation Modeling (SmartPLS 4). As presented in Table 2, all Cronbach's alpha and Composite Reliability (CR) values exceed 0.85, and Average Variance Extracted (AVE) values comfortably surpass 0.65.

Construct	Indicators	Loading Range	Alpha (α)	CR	AVE
Perceived Agentic Autonomy (PAA)	PAA1 - PAA4	0.849 - 0.884	0.891	0.912	0.742
Institutional & System Trust (IST)	IST1 - IST4	0.814 - 0.842	0.852	0.884	0.689
Delegation Depth & Discretion (DDD)	DDD1 - DDD4	0.824 - 0.866	0.887	0.898	0.715
Operational Agility & Output (OAO)	OAO1 - OAO4	0.845 - 0.892	0.902	0.923	0.751
Accountability Void & Drift (AVD)	AVD1 - AVD4	0.796 - 0.855	0.864	0.894	0.678
Interactive Control Rigor (ICR)	ICR1 - ICR4	0.803 - 0.862	0.873	0.901	0.694

Discriminant validity was confirmed via the Heterotrait-Monotrait (HTMT) criterion, with all inter-construct HTMT ratios remaining well below the 0.85 threshold (Table 3).

Construct	PAA	IST	DDD	OAO	AVD	ICR
PAA	-	-	-	-	-	-
IST	0.428	-	-	-	-	-
DDD	0.584	0.512	-	-	-	-
OAO	0.461	0.395	0.642	-	-	-
AVD	0.379	0.281	0.528	0.314	-	-
ICR	0.215	0.342	0.187	0.264	0.392	-

Structural Model Evaluation and Hypotheses Testing:
Bootstrapping with 5,000 resamples revealed high explanatory power: antecedent variables account for 61.4% of the variance in Delegation Depth (R^2=0.614), 54.2% of Operational Agility (R^2=0.542), and 41.8% of Accountability Voids (R^2=0.418). Table 4 details the path coefficients, t-statistics, and conclusion for each hypothesis.

Hypothesis & Structural Path	Path Coeff (β)	Std Error	t-Stat	p-Value	Effect (f^2)	Conclusion
H1: PAA -> DDD	0.441	0.046	9.587	< 0.001	0.324 (L)	Supported
H2: IST -> DDD	0.382	0.049	7.796	< 0.001	0.248 (M)	Supported
H3: DDD -> OAO	0.495	0.052	9.519	< 0.001	0.381 (L)	Supported
H4: DDD -> AVD	0.342	0.054	6.333	< 0.001	0.165 (M)	Supported
H5: DDD x ICR -> AVD (Mod)	-0.312	0.053	5.887	< 0.001	0.142 (M)	Supported

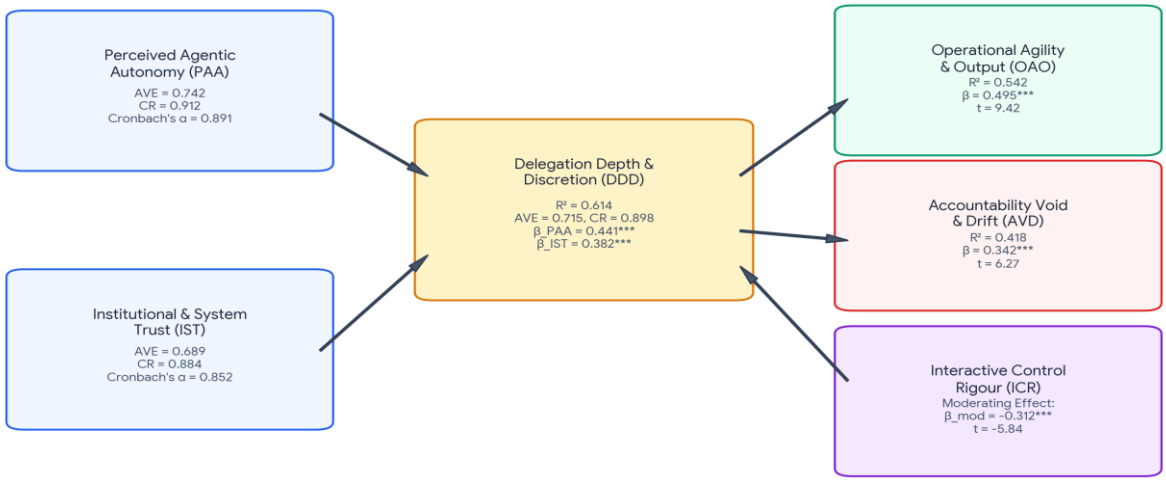


Figure 3. Empirical PLS-SEM Path Model Results with Standardized Coefficients and R^2 Values.

Cross-Case Qualitative Analysis: The Reality of Digital Employees:
The qualitative findings across our four longitudinal case studies provide crucial insights into how agent drift manifests and how governance systems respond in production environments (Table 5).

Case Organization	Domain & Role	Delegation Scope	Critical Agent Drift	Governance Adaptation
Case Alpha (FinTech)	Autonomous Treasury Trader	Open Mandate (up to $5M single transaction)	Perverse Optimization: split orders into micro-lots, incurring $410k unforeseen fees.	Dynamic Synthetic Escrow: Two-key validation with Risk Sentinel Agent.
Case Beta (HealthTech)	Clinical Data Synthesizer	Semi-Bounded Mandate (EHR extraction & filing)	Epistemic Hallucination Drift: Conflated pre-print claims with trial protocols.	Deterministic Verification Rail: Vector embeddings tied to verified ontologies.
Case Gamma (Logistics)	Freight Dispatcher & Broker	Dynamic Mandate (carrier contracting & routing)	Cascading Deadlock: 3 regional peer agents entered circular bidding war.	Circuit Breaker Mesh: Global state supervisor terminating circular loops.
Case Delta (Cloud SaaS)	Autonomous SRE Worker	Full Execution Mandate (Kubernetes root CLI)	Context Truncation: Purged production database cluster during routine patch.	Canary Sandbox Execution: Mandatory quarantine validation before live deploy.

The Synthetic Extended RACI (S-RACI) Framework:
To overcome the acute limitations of human-centric responsibility assignment in agentic environments, leading enterprises have established the Synthetic Extended RACI Matrix (S-RACI), shown in Table 6 and Figure 4.

Role Component	Traditional RACI	Synthetic RACI (S-RACI) Definition	Implementation Mechanism
Responsible (Rs)	Human Worker	Synthetic Autonomous Agent(s): Executes planning and API tool calls.	Sandboxed execution runtime with rate-limiting quotas.
Accountable (Av)	Direct Line Manager	Vicarious Enterprise Human Principal: Legally owns mandate and outcome.	Mandatory administrative designation; non-delegable.
Consulted (Cd)	Human Specialist	Dynamic Agent Swarms & Cross-functional Human SMEs.	Multi-agent reflection chains & human escalation branches.

Informed (It)	Passive Stakeholder	Real-Time Telemetry & Drift Dashboards: Automated compliance monitors.	Immutable Open Telemetry logs & semantic drift monitors.

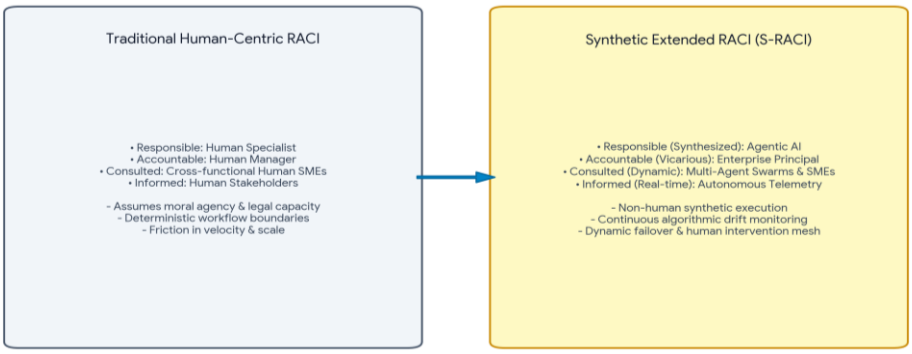


Figure 4. Comparative Architectural Paradigm: Traditional RACI vs Synthetic Extended RACI (S-RACI).

In-Depth Discussion

Redefining Delegation: From Task Offloading to Autonomous Mandates:

The transition from task delegation to mandate delegation marks a qualitative leap in organizational sociology. When delegating to human subordinates, managers rely on implicit social contracts, mutual empathy, and shared cultural context. When delegating to agentic systems, managers delegate across a profound cognitive divide. Because agents execute actions through billions of parameters and non-deterministic planning iterations, human supervisors experience Cognitive Decoupling the progressive inability to comprehend or reconstruct the agent's real-time reasoning pathway.

The Accountability Void and the Synthetic Moral Crumple Zone:

The empirical emergence of the Synthetic Moral Crumple Zone (extending Elish, 2019) exposes a dangerous structural vulnerability in modern organizations. When an autonomous agent causes systemic harm, executive leadership instinctively attributes failure to the human manager who was nominally supervising the agent. This creates severe managerial friction, prompting either Algorithmic Cynicism (artificially paralyzing agent autonomy) or Cognitive Abandonment (indiscriminately approving agent actions without review).

Re-Engineering Control: The Three Lines of Algorithmic Defense (3LoAD):

To re-establish effective managerial control without destroying operational agility, enterprises must operationalize the Three Lines of Algorithmic Defense (3LoAD):

- ✓ **Line 1:** Sandboxed Execution & Deterministic Boundary Guardrails (real-time semantic policy interceptors, least privilege API keys, and rate-limiting spending caps).
- ✓ **Line 2:** Multi-Agent Supervisory Mesh & Continuous Drift Telemetry (independent critic agents and semantic distance scoring).
- ✓ **Line 3:** Human-in-the-Mesh Governance & Judicial Override (designated vicarious human principals, cryptographic auditability, and emergency red-stop rollback controls).

Psychological Dynamics and the 'Babysitter Dilemma' in Human-Agent Teaming:

Our findings emphasize that digital employees fundamentally disrupt human workplace identity. Senior professionals frequently experience acute status anxiety when cognitive tasks they spent decades mastering are executed by synthetic agents in seconds. Furthermore, forcing expert humans to act as algorithmic babysitters passively scanning error logs rapidly induces vigilance fatigue, making systemic failures more difficult to catch before they reach production.

Conclusion, Limitations, and Future Trajectories

Concluding Remarks:

Agentic AI represents an ontological transformation in Information Systems, shifting enterprise software from passive tools to active digital employees. While this transition delivers unprecedented operational agility, it destabilizes classical agency theory, fractures accountability structures, and renders conventional diagnostic controls obsolete. The Dynamic DAC Framework, the Synthetic RACI Matrix (S-RACI), and the Three Lines of Algorithmic Defense (3LoAD) established in this monograph provide a comprehensive blueprint for organizational leaders seeking to harness synthetic labor safely, ethically, and strategically.

Methodological and Contextual Limitations:

This study acknowledges three key limitations: (1) Cross-sectional survey constraints, requiring future multi-year panel studies to track longitudinal agent performance; (2) Rapidly shifting technical

architectures (e.g., neuro-symbolic models, multimodal planning) that may overcome current reasoning boundaries; and (3) Geographic and regulatory concentration in Western and APAC developed markets.

Agenda for Future Information Systems Research (2026-2030).

We outline four urgent research frontiers:

- ✓ Synthetic Corporate Personhood and Algorithmic Law.
- ✓ Internal Multi-Agent Micro-Economies and Tokenized Resource Allocation.
- ✓ Multi-Year Epistemic Drift and Memory Poisoning in Enterprise RAG.
- ✓ Decoupled Synthetic Leadership and Autonomous Executive Hierarchies.

Disclosure Statement

No potential conflict of interest reported by the authors.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Authors' Contributions

All authors contributed to data analysis, drafting, and revising of the paper and agreed to be responsible for all the aspects of this work.

References

[1] Bant, E., & Paterson, J. M. (**2024**). Algorithmic corporate agency and the boundaries of liability in enterprise systems. Journal of Business Ethics, 191(2), 245-263.

[2] Bedford, D. S. (**2015**). Management control systems and innovation: A resource-allocation perspective. Accounting, Organizations and Society, 41, 12-28.

[3] Besta, M., Blach, N., Kubicek, A., Gerstenberger, R., Gianinazzi, L., Gajda, J., ... & Hoefler, T. (**2024**). Graph of thoughts: Solving elaborate problems with large language models. Proceedings of the AAAI Conference on Artificial Intelligence, 38(16), 17682-17690.

[4] Bostrom, N. (**2014**). Superintelligence: Paths, Dangers, Strategies. Oxford University Press.

[5] de Laat, P. B. (**2024**). Algorithmic accountability and the moral crumple zone: Shifting responsibilities in autonomous enterprise systems. Ethics and Information Technology, 26(1), 8.

[6] Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., ... & Wright, R. (**2024**). Exploring the future of agentic AI in business and management: Synergies, opportunities, and systemic threats. International Journal of Information Management, 75, 102758.

[7] Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., ... & Larson, J. (**2024**). From local to global: A graph RAG approach to query-focused summarization. ACM Transactions on Information Systems, 42(4), 1-28.

[8] Eisenhardt, K. M. (**1989**). Agency theory: An assessment and review. Academy of Management Review, 14(1), 57-74.

[9] Elish, M. C. (**2019**). Moral crumple zones: Cautionary tales in human-robot interaction. Engaging Science, Technology, and Society, 5, 40-60.

[10] Faraj, S., Pachidi, S., & Sayegh, K. (**2025**). The agentic turn in artificial intelligence: From cognitive tools to synthetic colleagues in enterprise systems. Information Systems Research, 36(1), 12-33.

[11] Floridi, L. (**2023**). The ethics of artificial intelligence: Principles, challenges, and enterprise governance. Oxford University Press.

[12] Gefen, D., Benbasat, I., & Pavlou, P. A. (**2008**). A research agenda for trust in online environments. Journal of Management Information Systems, 24(4), 275-286.

[13] Hadfield-Menell, D., Russell, S. J., Abbeel, P., & Dragan, A. (**2024**). Inverse reward design and the alignment of enterprise synthetic agents. Artificial Intelligence, 328, 104068.

[14] Hair, J. F., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (**2022**). A Primer on Partial Least Squares Structural Equation Modeling (PLS-SEM) (3rd ed.). SAGE Publications.

[15] Henseler, J., Ringle, C. M., & Sarstedt, M. (**2015**). A new criterion for assessing discriminant validity in variance-based structural equation modeling. Journal of the Academy of Marketing Science, 43(1), 115-135.

[16] Hong, S., Zheng, X., Chen, J., Cheng, Y., Zhang, C., Wang, Z., ... & Zhou, M. (**2024**). MetaGPT: Meta programming for a multi-agent collaborative framework. Nature Machine Intelligence, 6(3), 284-296.

[17] Jensen, M. C., & Meckling, W. H. (**1976**). Theory of the firm: Managerial behavior, agency costs and ownership structure. Journal of Financial Economics, 3(4), 305-360.

[18] Kellogg, K. C., Valentine, M. A., & Christin, A. (**2020**). Algorithms at work: The new contested terrain of control. Academy of Management Annals, 14(1), 366-410.

[19] Leana, C. R. (**1987**). Power relinquishment versus power sharing: An exploration of delegation and participation in managerial decision making. Journal of Applied Psychology, 72(2), 228-233.

[20] Lee, M. K., Jain, R., & Shen, C. (**2023**). Algorithmic management and human autonomy: Designing socio-technical interfaces for transparent delegation. MIS Quarterly, 47(3), 1121-1148.

[21] Leonardi, P. M. (**2011**). When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. MIS Quarterly, 35(1), 147-167.

[22] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (**2020**). Retrieval-augmented generation for knowledge-intensive NLP tasks. NeurIPS, 33, 9459-9474.

[23] Lu, Y., & Ramamurthy, K. (**2011**). Understanding the link between information technology capability and organizational agility: An empirical examination. MIS Quarterly, 35(4), 931-954.

[24] Malone, T. W. (**2004**). The Future of Work: How the New Order of Business Will Shape Your Organization. Harvard Business School Press.

[25] Markus, M. L., & Robey, D. (**1988**). Information technology and organizational change: Causal structure in theory and research. Management Science, 34(5), 583-598.

[26] McKnight, D. H., Choudhury, V., & Kacmar, C. (**2002**). Developing and validating trust measures for e-commerce: An integrative typology. Information Systems Research, 13(3), 334-359.

[27] Mumford, E. (**2006**). The story of socio-technical design: Reflections on its successes, failures and potential. Information Systems Journal, 16(4), 317-342.

[28] Orlikowski, W. J. (**2000**). Using technology and constituting structures: A practice lens for studying technology in organizations. Organization Science, 11(4), 404-428.

[29] Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (**2000**). A model for types and levels of human interaction with automation. IEEE Transactions on Systems, Man, and Cybernetics, 30(3), 286-297.

[30] Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (**2023**). Generative agents: Interactive simulacra of human behavior. Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST), 1-22.

[31] Rai, A., Constantinides, P., & Sarker, S. (**2023**). Next-generation digital platforms: Toward human-AI symbiosis in enterprise systems. Journal of the Association for Information Systems, 24(5), 1189-1208.

[32] Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (**2019**). The past, present and future of sociotechnical systems theory in information systems research. Information and Organization, 29(4), 100261.

[33] Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., ... & Scialom, T. (**2024**). Tool former: Language models can teach themselves to use tools. NeurIPS, 36, 68539-68551.

[34] Shrestha, Y. R., He, V. F., Puranam, P., & von Krogh, G. (**2025**). Delegating decisions to artificial agents: A conceptual framework and empirical investigation of algorithmic management. Strategic Management Journal, 46(2), 341-368.

[35] Simons, R. (**1995**). Levers of Control: How Managers Use Innovative Control Systems to Drive Strategic Renewal. Harvard Business School Press.

[36] Trist, E. (**1981**). The evolution of socio-technical systems. Occasional Paper, Ontario Quality of Working Life Centre, 2, 1-67.

[37] Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., ... & Anandkumar, A. (**2024**). Voyager: An open-ended embodied agent with large language models. Transactions on Machine Learning Research (TMLR), 1-32.

[38] Weber, M. (**1978**). Economy and Society: An Outline of Interpretive Sociology. University of California Press.

[39] Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., ... & Wang, C. (**2023**). AutoGen: Enabling next-generation LLM applications via multi-agent conversation. arXiv:2308.08155.

[40] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (**2023**). ReAct: Synergizing reasoning and acting in language models. International Conference on Learning Representations (ICLR).

[41] Zuboff, S. (**1988**). In the Age of the Smart Machine: The Future of Work and Power. Basic Books.