Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

From Automation to Autonomy: Minimizing the Productivity J-Curve in Artificial Intelligence Adoption

Ashutosh Shanker

Palo Alto Networks Inc., USA

ARTICLE INFO

Received: 06 March 2026

Accepted: 10 March 2026

ABSTRACT

The productivity paradox is not unique to a single technology revolution, with the take-up of artificial intelligence having a similar pattern to that of previous information technology revolutions. Despite improvements in task-level productivity in knowledge work, macro-level productivity statistics for the economy are not increasing. There is a disconnect between the siloed growth of technology and macroeconomic growth. Organizations are also prone to replicate the fragmentation, creating "islands of automation" rather than autonomous applications that can orchestrate the execution of knowledge tasks. The productivity J-curve assumes that investment in capability development, redesign of workflows and processes, and alignment of architectures will have a negative productivity payback if the investment in the productivity gains has not returned. The length and depth of the productivity J-curve can be reduced through systematic reuse of the lessons learned in the evolution of IT (end-to-end workflow management, development of complementary organizational capabilities, and convergence of the pilots into architectural platforms). Cloud-native integration and distributed orchestration patterns help achieve this autonomy by enabling data-powered solutions that orchestrate and distribute autonomous decisions across enterprise systems to create micro-level efficiencies that lead to system-level productivity.

Keywords: Productivity Paradox, Autonomous Process Coordination, Intangible Investment Burden, Architectural Integration, Workflow Transformation

1. Introduction

Recent advances in artificial intelligence have led to visible and often measurable improvements in the execution of individual tasks. Across domains such as customer support, software development, and knowledge work, AI-based tools have reduced completion times, improved consistency, and lowered the cognitive burden associated with routine activities [4], [14], [15]. These gains have contributed to a widespread expectation that AI adoption would translate into substantial productivity growth at the organizational and economic level [1], [5].

Yet, despite the pace of technological progress and the proliferation of AI initiatives within enterprises, broad productivity improvements remain difficult to observe [3], [5]. In many cases, organizations report localized efficiency gains without corresponding improvements in end-to-end performance [3], [10]. This tension between apparent micro-level success and muted system-level outcomes echoes earlier episodes in which new general-purpose technologies failed, at least initially, to deliver the productivity benefits that were widely anticipated [1], [3].

One influential explanation for this pattern is the productivity J-curve, which describes a temporary decline or stagnation in productivity following the introduction of transformative technologies [3]. Prior work attributes this effect to intangible complements such as learning costs, organizational restructuring, process redesign, and changes in governance [3], [6]. While these explanations are well grounded, they often treat productivity delays as largely unavoidable and slow-moving, offering

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

limited insight into how adoption strategies might actively shape the depth or duration of the productivity trough [3], [6].

In the case of AI, an additional factor warrants closer attention. Much of contemporary AI adoption is oriented toward task-level automation, where intelligence is applied to discrete activities while the surrounding processes remain largely unchanged [7], [10]. As a result, organizations frequently accumulate a collection of pilots, tools, and point solutions that operate in parallel rather than as parts of a coherent system. Although such deployments can deliver meaningful local benefits, they often introduce new coordination costs, manual handoffs, and oversight requirements that constrain their ability to scale into system-level productivity gains [7], [12].

This paper argues that the persistence of the AI productivity J-curve is shaped not only by organizational adjustment costs but also by the architectural form through which AI is embedded into enterprise operations. In particular, productivity gains tend to stall when AI remains confined to isolated tasks rather than being integrated into processes capable of coordinated, policy-bound execution. We refer to this distinction as the difference between automation and autonomy. Automation accelerates individual steps, whereas autonomy enables multi-step processes to be executed across systems with managed exceptions and human oversight.

From this perspective, architecture is not a secondary implementation concern but a central determinant of whether AI capabilities translate into sustained productivity improvements. Platform-level integration, workflow orchestration, and governance-aware design influence how intelligence propagates across organizational boundaries [9], [10], [12]. When these elements are absent, even highly capable AI systems struggle to move beyond local optimization.

The objective of this paper is to develop a conceptual framework that clarifies the relationship between AI adoption patterns and productivity outcomes. Drawing on prior work in productivity economics, enterprise systems, and distributed architectures, the paper reframes the AI productivity J-curve through an autonomy-first lens. Rather than presenting new empirical results, the paper synthesizes existing evidence to articulate mechanisms, architectural choices, and adoption levers that shape how and when productivity gains emerge.

1.1 Contributions of This Paper

This paper offers several contributions to the ongoing discussion of AI adoption and productivity.

First, it advances a mechanism-based explanation for why task-level improvements enabled by AI often fail to aggregate into system-level productivity gains. Rather than treating this gap as an inevitable consequence of organizational inertia, the paper highlights the role of process autonomy and coordination as a missing link between localized intelligence and enterprise-wide performance.

Second, the paper refines the notion of the AI productivity J-curve by distinguishing between sources of productivity loss that are structural and those that are potentially mitigable. This distinction helps explain why some organizations experience prolonged productivity slowdowns [3] while others recover more quickly, despite adopting similar AI technologies.

Third, the paper contributes an architectural perspective that complements existing discussions of intangible investments. It argues that data integration, workflow orchestration, and governance mechanisms constitute practical instantiations of these complements in modern AI systems, and that their absence can materially delay productivity realization.

Finally, the paper translates its conceptual arguments into implications for both practice and research. It identifies adoption levers that organizations can use to reduce productivity delays and outlines testable propositions that can guide future empirical work on AI-enabled productivity.

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

This paper is positioned as a conceptual, theory-building contribution that bridges research on productivity dynamics with scholarship on enterprise systems and organizational design. Rather than advancing new empirical estimates or proposing a specific technological implementation, the paper develops a mechanism-based explanation for why AI-enabled productivity gains often stall at the system level and identifies architectural and organizational conditions under which such gains are more likely to compound. By articulating testable propositions and clarifying boundary conditions, the paper aims to speak to researchers across information systems, management, and technology strategy, while remaining grounded in practical realities of AI adoption at scale.

2. The Productivity Paradox Across Technological Eras

The productivity paradox is the apparent failure of the investment in industrial (or transformative) technology to show up in observed aggregate productivity growth. It was prominent during the information technology revolution, when the American economist Robert Solow famously wrote in a 1987 book review that the computer age was visible everywhere except in the productivity statistics [1]. Much of the research on the digital economy emerged from the puzzle of how, despite exponential growth in computer processing power, the rise of the Internet, and high levels of investment in computer equipment, productivity seemed to be stagnating [5]. Computer processing power had grown at more than Moore's law from the 1970s until the mid-1990s. In that same period, productivity growth rates in advanced economies had slowed visibly compared to the post-war era [5].

The diffusion of general-purpose technologies, including artificial intelligence, exhibits many parallels with this historical episode, suggesting that the productivity paradox may be a general feature of general-purpose technology diffusion [3]. Empirical research has documented large micro-productivity gains for knowledge work, including faster and higher throughput in customer service calls, content generation, and software development tasks [4], [15]. These task-level improvements suggest that AI can increase the productivity of the individual worker for certain narrow tasks. However, higher-level measures of productivity have only been modestly affected by the adoption of AI. In spite of the rapid adoption of AI, productivity growth in the economy as a whole has remained relatively subdued [3], [5]. This decoupling between task-level productivity enhancement and stagnation of system-level economic productivity growth is analogous to earlier waves of information and communications technology (ICT) diffusion [3].

The fundamental principle linking the two technological revolutions is the shift from automation to autonomy. Information technology has automated individual tasks in customary processes, providing localized process productivity improvements [9]. More radical advancements for end-to-end business processes were limited [9], [10]. Artificial intelligence allows qualitatively different new capabilities, such as the autonomous coordination of complex, multi-step, interdependent decision flows across distributed organizational systems [7], [12]. Practically, organizations have thus far deployed artificial intelligence technologies to augment tasks, as opposed to reconfiguring entire processes. This pattern of implementation replicates the less-than-maximum productivity impact of IT during its first several decades of diffusion, characterized by focused attempts to automate tasks within established business processes [9], [10]. This reflects the contrast between the automation of individual activities in existing business processes and the end-to-end process orchestration of systems, which maximizes productivity at the systems level [9], [10].

3. The J-Curve Phenomenon and Intangible Investment Burden

The observation that major technological advances do not immediately translate into productivity gains has a long history in economic and organizational research. Early studies of electrification and

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

information technology documented extended periods in which productivity growth lagged behind technological capability, despite widespread diffusion of the underlying technologies. Similar concerns have re-emerged in recent years as artificial intelligence has moved from research laboratories into mainstream organizational use, with visible task-level improvements coexisting alongside muted aggregate productivity effects [1], [3].

This pattern is commonly described through the productivity J-curve, in which the introduction of a general-purpose technology is followed by a temporary decline or stagnation in measured productivity before longer-term gains materialize [3]. Prior work attributes this dynamic primarily to intangible investments, including organizational learning, process redesign, workforce adaptation, and changes in governance structures [3], [6]. These investments impose real short-term costs and delay the realization of productivity benefits, even when the underlying technology is demonstrably effective.

While this explanation is well established, much of the literature treats the intangible investment burden as an unavoidable transitional phase. Learning curves, coordination costs, and organizational inertia are often framed as slow-moving frictions that resolve largely with time. Yet empirical observations suggest substantial variation across firms and sectors in both the depth and duration of the productivity trough. Organizations adopting similar AI technologies frequently experience markedly different outcomes, indicating that adoption patterns and design choices play a significant role in shaping productivity trajectories [10].

One dimension of adoption that has received comparatively limited attention is the level at which AI is embedded within organizational processes. Much contemporary AI deployment emphasizes task-level automation, where intelligence is applied to discrete activities such as classification, drafting, prediction, or triage. These systems are often introduced as stand-alone tools or incremental enhancements to existing workflows. Although such deployments can deliver meaningful local efficiency gains, they typically leave the surrounding process architecture unchanged, relying on manual coordination to link automated tasks together.

It is in this context that the distinction between automation and autonomy becomes analytically important. In this paper, autonomy is used in a restricted and operational sense. It does not refer to fully self-governing or unconstrained artificial intelligence. Rather, autonomy denotes the extent to which AI-enabled systems can execute multi-step business processes across organizational and system boundaries within predefined policies, handling routine decisions automatically while escalating non-standard cases to human operators. This usage aligns with prior work on human–machine collaboration and decision delegation, which emphasizes bounded authority, escalation mechanisms, and oversight rather than full independence [12, 13].

Task automation and process autonomy differ in their implications for productivity. Automation improves the efficiency or accuracy of individual activities but leaves intact the coordination mechanisms that connect tasks into end-to-end processes. An automated task may still require human intervention for sequencing, approval, exception resolution, or reconciliation across systems. Autonomy, by contrast, concerns how decisions and actions propagate across an entire workflow. Autonomous processes are characterized by their ability to progress across multiple steps without continuous human mediation under normal operating conditions, while preserving accountability and governance.

Autonomy should therefore be understood as a matter of degree rather than as a binary property. At lower levels, AI systems provide recommendations while humans retain responsibility for execution and coordination. At intermediate levels, AI systems may execute predefined actions subject to policy constraints, approval thresholds, or confidence bounds. At higher levels, autonomy involves coordinated execution across multiple systems, with human involvement focused on supervision,

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

exception handling, and policy definition. This graduated view of autonomy is consistent with prior research on levels of automation and adaptive control in complex socio-technical systems [12].

This distinction helps explain why task-level AI gains often fail to aggregate into system-level productivity improvements. A substantial portion of productivity loss arises not from task execution itself, but from the coordination work required to connect tasks into coherent processes. Manual handoffs, fragmented data flows, duplicated oversight, and exception management introduce delays and cognitive overhead that can offset local efficiency gains. When AI adoption increases the number of automated tasks without reducing these coordination costs, the result can be a widening gap between micro-level improvements and overall performance.

From this perspective, the intangible investment burden associated with AI adoption is shaped not only by learning and adjustment costs, but also by architectural and organizational choices that determine how intelligence is distributed and coordinated. Where AI remains confined to isolated tasks, coordination overhead persists and the productivity trough is prolonged. Where higher degrees of process autonomy are achieved, coordination costs can be reduced, allowing task-level gains to compound rather than dissipate. In this sense, autonomy functions as a mediating mechanism within the productivity J-curve, influencing both its depth and duration.

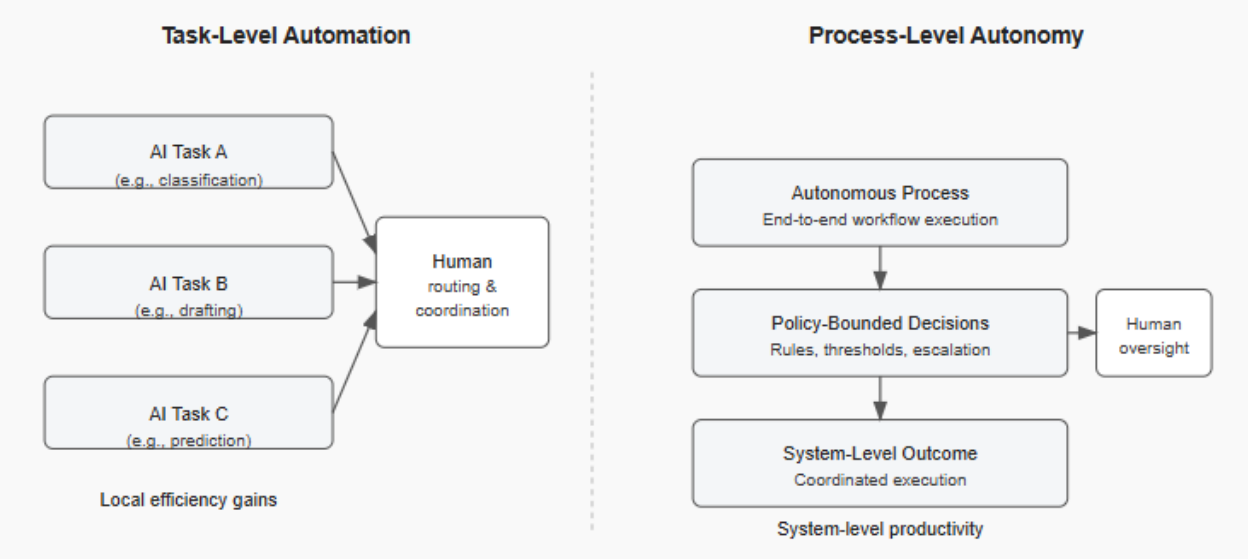

Fig 1: Conceptual Contrast: Task-Level Automation versus Process-Level Autonomy

Task automation and process autonomy differ in their implications for productivity. Automation improves the efficiency or accuracy of individual activities but leaves intact the coordination mechanisms that connect tasks into end-to-end processes. An automated task may still require human intervention for sequencing, approval, exception resolution, or reconciliation across systems. Autonomy, by contrast, concerns how decisions and actions propagate across an entire workflow. Autonomous processes are characterized by their ability to progress across multiple steps without continuous human mediation under normal operating conditions, while preserving accountability and governance.

Autonomy should therefore be understood as a matter of degree rather than as a binary property. At lower levels, AI systems provide recommendations while humans retain responsibility for execution and coordination. At intermediate levels, AI systems may execute predefined actions subject to policy constraints, approval thresholds, or confidence bounds. At higher levels, autonomy involves coordinated execution across multiple systems, with human involvement focused on supervision,

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

exception handling, and policy definition. This graduated view of autonomy is consistent with prior research on levels of automation and adaptive control in complex socio-technical systems [12].

This distinction helps explain why task-level AI gains often fail to aggregate into system-level productivity improvements. A substantial portion of productivity loss arises not from task execution itself, but from the coordination work required to connect tasks into coherent processes. Manual handoffs, fragmented data flows, duplicated oversight, and exception management introduce delays and cognitive overhead that can offset local efficiency gains. When AI adoption increases the number of automated tasks without reducing these coordination costs, the result can be a widening gap between micro-level improvements and overall performance.

From this perspective, the intangible investment burden associated with AI adoption is shaped not only by learning and adjustment costs, but also by architectural and organizational choices that determine how intelligence is distributed and coordinated. Where AI remains confined to isolated tasks, coordination overhead persists and the productivity trough is prolonged. Where higher degrees of process autonomy are achieved, coordination costs can be reduced, allowing task-level gains to compound rather than dissipate. In this sense, autonomy functions as a mediating mechanism within the productivity J-curve, influencing both its depth and duration.

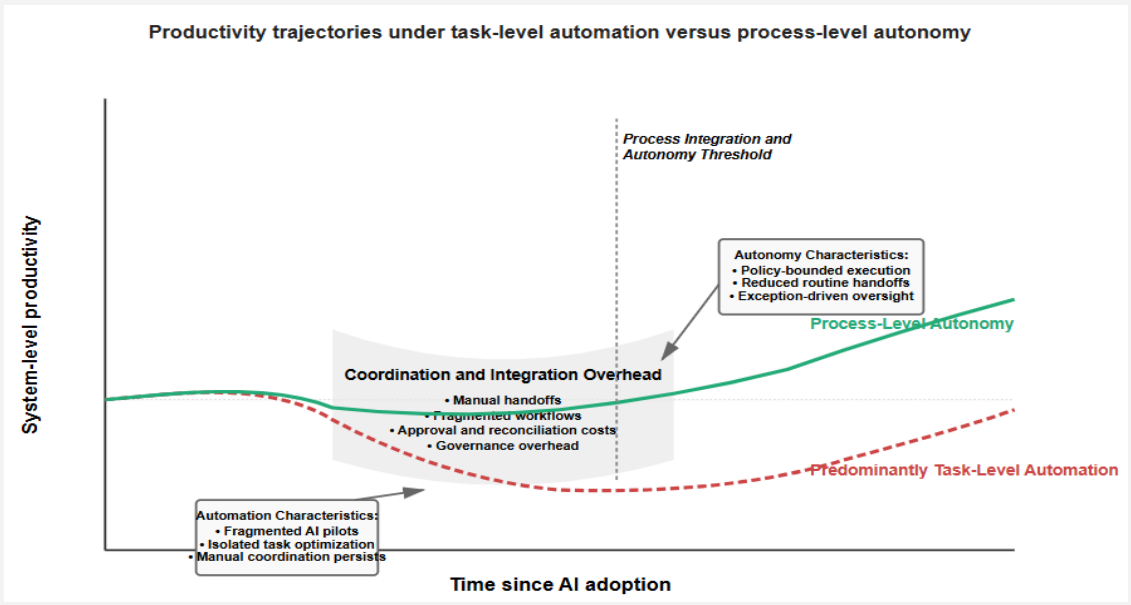

Figure 2. Conceptual illustration of productivity trajectories under task-level automation and process-level autonomy. Coordination and integration overhead act as a mediating mechanism that influences the depth and duration of the productivity J-curve following AI adoption.

Figure 2 illustrates the mechanism through which different AI adoption patterns give rise to divergent productivity trajectories over time. When AI is deployed primarily at the level of individual tasks, local efficiency gains are often offset by persistent coordination and integration costs, including manual handoffs, fragmented workflows, and increased oversight requirements. These frictions deepen and prolong the productivity trough, limiting the extent to which task-level improvements translate into system-level performance gains [3]. By contrast, when AI is embedded within more autonomous, policy-bounded processes, routine coordination overhead is reduced and human involvement shifts toward exception handling and supervision [12]. Under these conditions, task-level gains are more likely to compound, resulting in a shallower productivity trough and a faster recovery toward sustained productivity growth [3]. The figure highlights coordination overhead as a mediating mechanism that shapes both the depth and duration of the AI productivity J-curve.

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Investment Component	Structural Necessity	Strategic Mitigability	Productivity Impact Phase	Primary Cost Driver
Workforce Capability Development	High	Low	Initial Decline	Training Time Allocation
Organizational Learning Cycles	High	Medium		Experimentation Iterations
Workflow Redesign	High	Medium		Process Reengineering
Data Infrastructure	High	Low		Pipeline Construction
Governance Frameworks	Medium	High		Quality Assurance Systems
Integration Architecture	High	Medium	Transition Phase	Cross-System Coordination
Pilot Program Consolidation	Low	High		Fragmented Investments
Architectural Planning	Low	High	Recovery Phase	Rework Prevention

Table 1: Intangible Investment Components and Their Impact on J-Curve Dynamics [3]

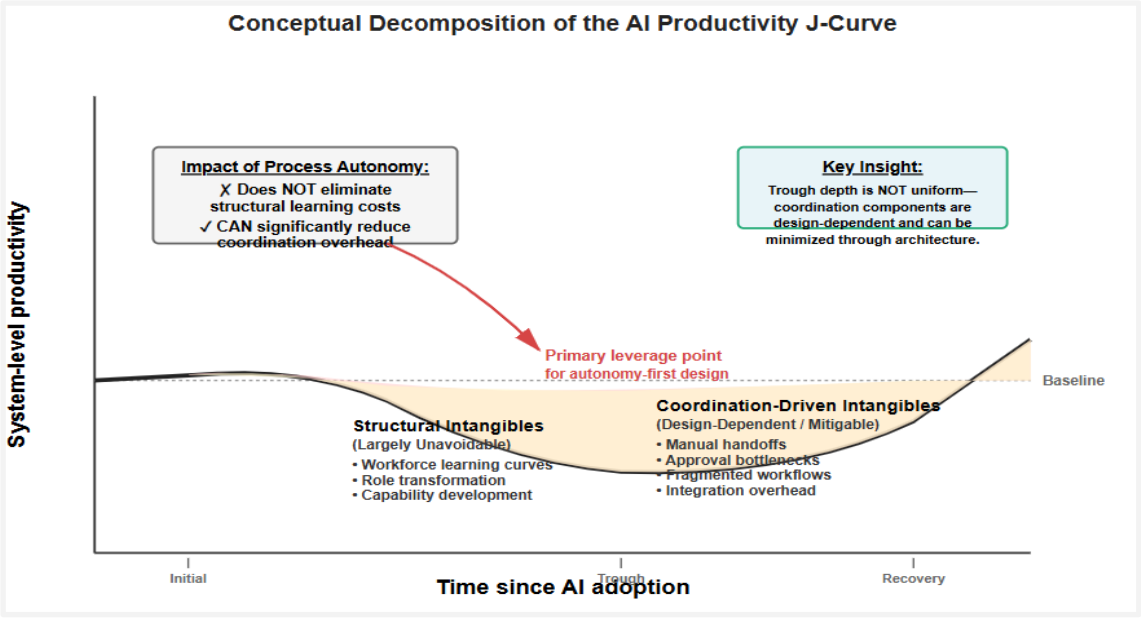

Fig 3: Decomposition of the AI Productivity J-Curve: Structural versus Mitigable Components

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

4. Fragmented Adoption Patterns and Organizational Mechanisms

The organizational bottlenecks that create lag periods during shifts in technology are remarkably similar across historical and modern waves of adoption. In response to the introduction of IT, early adopters organized "islands of automation" within their organizations that isolated the new technology within departmental boundaries without system-wide or cross-organization integration of processes and technology [1], [9]. While these partial implementations improved the efficiency of individual activities within a functional unit, they nevertheless failed to produce productivity gains at the system level [3], [5]. The mechanization methods applied were unable to change the underlying process architectures. Individual business functions used them to automate individual activities using stand-alone systems, separately from other functions, without improving cross-process workflows, information flow, and decision-making [9], [10]. This pattern of isolated deployment resulted in a continued drag on aggregate productivity, as organizations made technology investments without process transformation [3].

This can also be observed in the contemporary era in the proliferation of disconnected pilot programs across enterprise organizations, where enterprises can rapidly initiate an array of AI deployments for different use cases, departments, or pieces of the workflow without thorough architectural planning or integration strategy [10]. Each pilot can yield task-level efficiency gains in its bounded domain, but the projects collectively do not yield improved systemic productivity: these pilot projects reinforce the existing boundaries and silos of the workflow and avoid the coordination problems associated with more ambitious projects that would redesign end-to-end processes [3], [10]. They continue to invest in technology acquisition, pilot projects, new capabilities, and features, deferring the harder work of architectural integration and end-to-end process and workflow transformation [6], [9].

The barriers to adoption of artificial intelligence due to siloed adoption are not only the inefficiencies of communication between siloed systems, but rather the violations of the architectural principles for achieving autonomy. Siloed systems do not allow the data integration, cross-functional orchestration, and distributed decision coordination needed to achieve autonomy across business processes [9], [12]. Applied only within departments, though, they do not supply the necessary information flows, integration points, and processes spanning multiple systems that would allow these technologies to autonomously coordinate complex processes. This architectural fragmentation is the primary organizational mechanism anchoring organizations to the left-hand side of the productivity curve, resulting in intangible investments that do not serve as building blocks for moving from task automation to process autonomy [3], [6]. It suggests that firms remain stuck in long-run low-productivity traps, not because of the technology choices open to them, but because of the adoption patterns that prevent architectural integration and process tasks necessary to increase productivity.

Adoption Pattern	Historical IT Era	Contemporary AI Era	Architectural Barrier	System-Level Impact
Isolated Departmental Deployment	Islands of Automation	Disconnected Pilot Programs	Siloed Data Access	Prevented Cross-Functional Integration
Task-Level Optimization	Standalone Systems	Task Augmentation Tools	Legacy Workflow Preservation	Limited Process Transformation
Sequential Implementation	Manual Coordination	Human Approval Gates	Fragmented Decision Flows	Constrained Autonomous

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

	Points			Operation
Technology-First Approach	Infrastructure Retrofit	Point-Solution Deployment	Inadequate Integration Planning	High Coordination Costs
Departmental Boundaries	Functional Silos	Use Case Silos	Incomplete Information Flows	Blocked End-to-End Visibility
Incremental Automation	Individual Task Focus	Individual Activity Enhancement	Process Architecture Inertia	Localized Efficiency Gains

Table 3: Organizational Mechanisms Anchoring Enterprises in Low-Productivity Phases [2]

5. Task-Level Gains Versus System-Level Transformation

However, experiments have shown that the use of AI tools in specific knowledge work areas may lead to large productivity gains at the micro level. For example, in customer support, studies showed that generative AI tools improved the number of issues resolved by support agents, with differing effects based on the worker's skill and prior experience [4]. Findings indicate that the use of conversational assistants to provide workers with real-time recommendations during customer interactions increases worker capacity (i.e., the ability to resolve customer interactions in a period of time), particularly for newer employees and those with lower baseline productivity. The findings suggest that generative AI technologies can improve human performance in structured interactions by providing contextually relevant recommendations, suggesting actions, and reducing the cognitive demands involved in real-time information search during interactions with customers [4].

Similar, large task-level effects have also been documented across other domains of knowledge work, with large reductions in writing and editing times for professional writing via generative AI, as well as for the generation and debugging of code via AI coding assistants [15], [14]. These micro-level findings are in line with an otherwise broad literature on the ability of AI to increase productivity across a range of cognitive tasks in which information is synthesized, patterns are discerned, and structured output is produced. The productivity gains are non-negligible and lead to observable time savings and throughput gains.

However, these localized productivity gains are at odds with macroeconomic data, which has shown that productivity growth has stagnated [3], [5]. AI tools are commonly employed in enterprise work environments, and evidence at the task level indicates that productivity gains have occurred. However, these gains have not been reflected in economy-wide data [3]. This articulates the fundamental paradox of early investing in IT: the costs and benefits at the task level did not automatically translate to higher productivity at the organizational level [1], [3]. This suggests that task augmentation leads to more productive employees at the task level, but not necessarily to higher productivity at the organizational level, unless work processes are restructured and workflow architectures redesigned [9], [10].

One potential explanation for the persistence of the gap is task automation and the structural barrier separating it from system transformation. When organizations deploy AI technologies to augment tasks within existing workflows, they obtain efficiency gains without modifying their legacy process architectures and coordinative mechanisms. While an individual query might be answered more quickly through agent-assisted AI, the support process itself is bound by time-consuming manual escalation processes, information silos, and stage-gate workflows, preventing process-level automation [9], [10]. Organizations are still limited to adopting automation in a piecemeal fashion to optimize individual use cases without addressing the architectural foundations that are required for

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

process-level automation end-to-end, preventing productivity from being realized across the organization [3], [6].

Knowledge Work Domain	Task-Level Improvement	Measurement Context	Structural Constraint	System-Level Barrier
Customer Support Operations	Increased Resolution Rate	Individual Agent Throughput	Manual Escalation Procedures	Sequential Workflow Stages
Professional Writing	Reduced Completion Time	Content Generation Speed	Editorial Review Gates	Multi-Stage Approval Processes
Software Development	Accelerated Code Generation	Individual Developer Output	Integration Dependencies	Deployment Coordination Requirements
Content Editing	Enhanced Efficiency	Task Completion Velocity	Quality Assurance Checkpoints	Cross-Functional Review Cycles
Customer Interaction	Lower Cognitive Burden	Real-Time Recommendation Utility	Information Silos	Disconnected System Architecture
Code Debugging	Faster Issue Resolution	Development Workflow Speed	Testing Protocols	Release Management Complexity

Table 4: Micro-Level Productivity Gains Versus Macro-Level Stagnation Dynamics [2, 4]

6. Architectural Requirements for Autonomy-First Process Design

The preceding discussion suggests that productivity outcomes following AI adoption are shaped not only by improvements in task-level performance but by how intelligence is embedded within organizational processes. From this perspective, architecture plays a central role, not as a collection of tools or platforms, but as the set of design choices that determine where decisions are made, how work is coordinated, and how responsibility is distributed between humans and systems. Prior research on enterprise systems and organizational complements emphasizes that such architectural choices influence whether coordination costs are internalized by systems or externalized to human actors, with direct implications for productivity outcomes [9], [6].

Much of the current discourse on AI architecture focuses on scalability, deployment efficiency, and model integration. While these concerns are important, they are insufficient for explaining productivity dynamics at the system level. Research on digital transformation and process innovation consistently shows that productivity gains depend less on the presence of advanced technology than on whether technologies are integrated into end-to-end processes that reduce fragmentation and manual coordination [9], [10], [12]. In this sense, architecture functions as the mechanism through which autonomy is operationalized.

When AI is introduced primarily as a set of task-level capabilities, architectural complexity often increases rather than decreases. Discrete AI components are layered onto existing systems, while coordination across tasks remains manual or loosely specified. Humans continue to manage

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

sequencing, approvals, and exception handling, even as individual tasks become faster or more accurate. Studies of enterprise automation and decision-support systems have shown that such partial automation can inadvertently increase coordination overhead by introducing new dependencies and oversight requirements [7, 12].

By contrast, architectures that support higher degrees of process autonomy reassign routine coordination work from humans to systems. This does not eliminate human involvement, but changes its role. Rather than managing every handoff, humans define policies, constraints, and escalation conditions, while systems execute routine process flows and surface exceptions. This view aligns with research on levels of automation and human–machine collaboration, which emphasizes that effective autonomy depends on clear boundaries, adaptive control, and well-designed escalation mechanisms [12], [13].

To illustrate this distinction, consider the example of incident response in an enterprise IT or security context. Under a task-level automation approach, AI tools may assist with individual activities such as alert classification, log analysis, or remediation recommendation. Each task may be performed more efficiently, yet humans remain responsible for triaging incidents, determining next steps, coordinating across teams, and tracking resolution. Empirical studies of security operations and complex incident management indicate that, as event volumes grow, the coordination burden associated with such fragmented workflows can offset gains achieved through automation [7].

Under a more autonomous, process-oriented architecture, the same AI capabilities are embedded within an end-to-end incident response workflow. Events trigger predefined response paths governed by policy. Routine incidents are investigated, enriched, and remediated automatically within bounded conditions, while atypical cases are escalated for human review. Human operators focus on policy definition, threshold tuning, and exception handling rather than on manual orchestration. Prior work on workflow automation and business process management suggests that such integration is a key determinant of cycle-time reduction and operational scalability [9], [11].

Importantly, such architectures do not eliminate the need for organizational learning, data quality improvement, or governance mechanisms. Learning costs remain, particularly as roles and responsibilities shift. Data integration challenges persist, especially in heterogeneous and legacy environments. Governance remains essential to ensure safety, compliance, and accountability. Research on enterprise system implementation consistently shows that these intangible investments are unavoidable, but that their productivity impact depends on how tightly processes are integrated and coordinated [6], [10].

Architectural choices also introduce new failure modes that must be acknowledged. Poorly designed autonomous processes can amplify errors if feedback loops are not well controlled. Over-automation can increase exception load, shifting cognitive burden back onto humans in less predictable ways. Governance mechanisms that are overly rigid can negate the benefits of autonomy by reintroducing manual approvals at scale. These risks echo long-standing concerns in automation research regarding brittleness, loss of situational awareness, and the need for calibrated autonomy [7], [12].

Viewed through the lens of the productivity J-curve, architecture functions as a mediating layer between AI capability and realized productivity. Architectures that confine AI to isolated tasks tend to preserve coordination costs, deepening and prolonging the productivity trough illustrated in Figure 1. Architectures that embed AI within autonomous, policy-bounded processes reduce routine coordination effort and allow task-level gains to accumulate more rapidly. In this way, architectural design choices shape not only how AI systems operate, but how organizations experience the transition from experimentation to sustained productivity growth [3].

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Architectural Element	Traditional Automation	Autonomy-First Design	Technical Capability	Organizational Requirement
Integration Pattern	Departmental Systems	Cloud-Native Architecture	Dynamic Resource Allocation	Cross-Functional Coordination
Orchestration Layer	Sequential Workflows	Distributed Orchestration	Autonomous Decision Logic	Real-Time Process State Management
Data Architecture	Task-Specific Views	Data-Centric Workflows	Comprehensive Information Access	Unified Data Environments
Execution Environment	Monolithic Deployment	Containerized Services	Microservices Isolation	Component Independence
Communicatio n Model	Synchronous Handoffs	Event-Driven Patterns	Asynchronous Coordination	Flexible Process Composition
Decision Coordination	Human Approval Gates	Autonomous Routing	Context Evaluation Capability	Distributed Decision Authority
Infrastructure Investment	Application-First	Platform-First	Shared Integration Capabilities	Systematic Architectural Planning
Process Visibility	Departmental Scope	End-to-End Transparency	Multi-System Integration	Cross-Domain Information Flow

Table 5: Architectural Components Enabling Autonomy-First Process Design [5, 6]

7. Testable Propositions and Evaluation Implications

The arguments developed in this paper suggest that productivity outcomes following AI adoption depend not only on improvements in task-level performance, but also on how intelligence is embedded within organizational processes. While this work is conceptual in nature and does not present new empirical tests, the proposed mechanism yields a set of propositions that can be evaluated in future research using observable organizational outcomes [3], [6].

Proposition 1. Organizations that deploy AI primarily through task-level automation are likely to experience longer productivity troughs than organizations that embed AI within more autonomous, process-level workflows, due to the persistence of coordination and integration costs [3].

This proposition reflects the distinction between localized efficiency gains and system-level productivity outcomes. Where AI adoption increases the number of automated tasks without reducing manual handoffs or approval cycles, coordination overhead is expected to delay productivity recovery.

Proposition 2. The relationship between AI capability adoption and realized productivity gains is mediated by the degree of process autonomy achieved, such that higher levels of autonomy enable task-level improvements to aggregate into system-level performance [3].

This proposition formalizes autonomy as a mediating mechanism rather than a complementary feature. It implies that advances in AI capability alone are insufficient to generate productivity gains unless accompanied by changes in how processes are executed and coordinated [6], [9].

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Proposition 3. Reductions in routine coordination overhead, including manual handoffs, reconciliation effort, and approval cycles, are more strongly associated with productivity recovery than improvements in task-level AI accuracy or speed alone [3].

This proposition emphasizes that productivity bottlenecks often arise from process friction rather than from task execution inefficiency. It suggests that investments targeting coordination costs may yield greater productivity returns than further optimization of isolated AI tasks.

Proposition 4. Governance mechanisms that bound autonomous execution through explicit policies and escalation rules moderate the relationship between process autonomy and productivity outcomes, enabling productivity gains while maintaining organizational control [12], [13].

This proposition acknowledges that autonomy operates within organizational and regulatory constraints. It highlights governance not as a barrier to autonomy, but as a condition that shapes whether autonomous execution can be scaled without increasing oversight burden.

Proposition 5. Organizations with higher baseline levels of process integration will realize productivity gains from AI adoption more rapidly than organizations with fragmented workflows, even when deploying similar AI technologies [3], [10].

This proposition helps explain observed variation in productivity outcomes across organizations. It implies that pre-existing process structure conditions the speed with which AI-enabled autonomy can be realized.

Although these propositions are not empirically tested in this paper, they can be examined using observable indicators available in organizational and operational data. Relevant measures may include the time required to move from pilot deployment to scaled operation, the proportion of workflows executed end-to-end without manual intervention, the frequency and cost of exception handling, changes in cycle time across functions, and shifts in coordination effort per transaction. Together, such measures provide a basis for evaluating how different AI adoption patterns influence productivity trajectories over time.

8. Strategic Pathways for Minimizing the Productivity Dip

To minimize the temporary productivity loss of adopting artificial intelligence, general principles from previous technology transitions should be applied, especially the eventual productivity gains of the information technology era. There is evidence that these productivity gains only occurred after incremental process automation had given way to end-to-end process and architectural integration [3], [9], [10]. The history of computing suggests that the depth and duration of the J-curve can be managed using organizational redesign tools rather than just investments in infrastructure. It is not so much a question of technology sophistication, but of how to redesign processes, invest in complementary capabilities, and build integrated infrastructures that support system-level rather than task-level optimization [3], [6].

Organizations that want to capture productivity gains quickly have to start with workflow redesign, instead of technology rollout. Otherwise, they will repeat the pattern of piecemeal rollout over several years that caused information technology productivity to lag for decades [1], [3]. Planned decisions should start with an end-to-end business process analysis, aiming for autonomous operation on a process level rather than for piecemeal support of each business operation. This process-first approach requires deciding systematically which decisions are based on human judgment and which can be automated, and rethinking the information flows to support continuous end-to-end process execution in a cross-process way [9], [11].

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Investments in intangible complements are a calculated consideration. The acquisition of technology is estimated to account for only a fraction of the total investment in transformation. For instance, in information technology adoption, the costs of investing in skills and systems, organizational learning, knowledge, and processes very often exceed the costs of acquisition by a factor of some multiples [5], [6]. Underinvestment in these critical intangibles can prolong productivity slumps, since firms' technologies are operating in capability-constrained environments. Investment plans for building related capabilities should be contained in calculated plans, with the underlying capabilities such as skills, data, integration architectures, and coordination systems being scaled in synchronization with growing technologies, thus raising the overall value [6].

Integrated architectural platforms on which experimental projects are developed address the challenge of fragmentation, which traps organizations in a long phase of low productivity [3], [9], [10]. Instead of a buildup of independent pilot programs across departments and business units, strategies use integrated platforms with shared orchestration and integration interfaces [9], [10]. Furthermore, it allows for systematic learning transfer across implementation teams, economies of scale for infrastructure investments, acceleration from task automation to process autonomy, and earlier realization of transformative productivity gains [3].

9. Limitations and Scope

This paper develops a conceptual framework to explain how AI adoption patterns influence productivity trajectories, with a particular focus on the role of process-level autonomy. As such, the arguments advanced here are not empirically tested within the scope of this study. While the framework is grounded in established literatures on productivity, enterprise systems, and human–machine interaction, empirical validation remains necessary to assess the magnitude and generalizability of the proposed relationships across different organizational settings [3], [6].

A second limitation concerns the measurement of productivity and coordination costs in complex organizations. System-level productivity effects often manifest with substantial delays and are difficult to attribute to specific technologies or design choices [3], [5]. Coordination overhead, which plays a central role in the proposed mechanism, is particularly challenging to observe directly and is frequently embedded in informal practices, managerial effort, and exception handling. As a result, empirical studies evaluating the propositions outlined in this paper will need to rely on proxies and mixed methods to capture these dynamics [6].

The framework also abstracts from important sources of organizational heterogeneity. Differences in industry context, regulatory constraints, legacy system complexity, and workforce composition are likely to influence both the feasibility and the impact of process-level autonomy. Organizations with tightly coupled processes or high regulatory oversight may experience different productivity trajectories than those operating in less constrained environments [10]. Accordingly, the framework should be interpreted as identifying general mechanisms rather than predicting uniform outcomes across all contexts.

Finally, the discussion of autonomy assumes the presence of governance mechanisms that bound system behavior through policies, escalation rules, and human oversight. In practice, poorly designed governance structures can either constrain autonomy excessively or allow it to operate without adequate safeguards. The framework does not attempt to resolve these design trade-offs, but instead highlights autonomy as a contingent organizational capability whose productivity effects depend on careful calibration rather than unrestricted deployment [12], [13].

Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Conclusion

The productivity paradox is a recurrent pattern of technology adoption, and it occurs in other eras of information technology and AI [1], [3]. While productivity improvements are possible, the continued stagnation indicates that important organizational obstacles lie between potential technological impact and actual productivity improvements [3], [5]. This J-curve shape is largely driven by time lags, as the intangible investments in capabilities, work processes, and process architecture must precede payoffs [3], [6]. The historical fragmentation of organizations is preserved by their modern disconnected pilot projects, which maintain the process architectures of the legacy processes that cannot coordinate autonomously[9], [10].

To avoid low productivity, purposeful learning from information technology evolution calls for a full redesign of business processes, investments in complementary capabilities, and the integration of implementations as end-to-end architectural platforms [3], [6], [9]. The transition from task automation to process autonomy takes cloud-native integration patterns, distributed orchestration frameworks for connected services, and data-oriented architectural styles to achieve end-to-end process autonomy. Those organizations able to navigate the transition and copy success will accelerate the downward slope of the productivity curve, transforming local productivity gains into solutions that spread and deliver the economic promise of artificial intelligence [3].

References

[1] R. M. Solow, "We'd better watch out," *New York Times Book Review*, Jul. 12, 1987.
[2] E. Brynjolfsson, T. Malone, V. Gurbaxani, and A. Kambil, "Does information technology lead to smaller firms?" *Manage. Sci.*, vol. 40, no. 12, pp. 1628–1644, 1994.
[3] E. Brynjolfsson, D. Rock, and C. Syverson, "The productivity J-curve: How intangibles complement general-purpose technologies," *Amer. Econ. J.: Macroecon.*, vol. 13, no. 1, pp. 333–372, 2021, doi: 10.1257/mac.20180386.
[4] E. Brynjolfsson, D. Li, and L. Raymond, "Generative AI at work," *Nat. Bur. Econ. Res.*, Working Paper 31161, 2023. [Online]. Available: https://www.nber.org/papers/w31161 *(Used as illustrative micro-level evidence)*
[5] D. W. Jorgenson, M. S. Ho, and K. J. Stiroh, "A retrospective look at the U.S. productivity growth resurgence," *J. Econ. Perspect.*, vol. 22, no. 1, pp. 3–24, 2008.
[6] E. Brynjolfsson, L. M. Hitt, and S. Yang, "Intangible assets: Computers and organizational capital," *Brookings Papers Econ. Activity*, vol. 2002, no. 1, pp. 137–198, 2002, doi: 10.1353/eca.2002.0016.
[7] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983, doi: 10.1016/0005-1098(83)90046-8.
[8] E. Brynjolfsson and K. McElheran, "The rapid adoption of data-driven decision-making," *Amer. Econ. Rev. Papers Proc.*, vol. 106, no. 5, pp. 133–139, 2016, doi: 10.1257/aer.p20161016.
[9] T. H. Davenport and J. E. Short, "The new industrial engineering: Information technology and business process redesign," *Sloan Manage. Rev.*, vol. 31, no. 4, pp. 11–27, 1990.
[10] M. L. Markus and C. Tanis, "The enterprise systems experience—From adoption to success," in *Framing the Domains of IT Management*, R. W. Zmud, Ed. Cincinnati, OH, USA: Pinnaflex, 2000, pp. 173–207.
[11] J. Mendling, H. A. Reijers, and W. M. P. van der Aalst, "Seven process modeling guidelines (so far)," *Inf. Softw. Technol.*, vol. 94, pp. 191–201, 2018, doi: 10.1016/j.infsof.2017.09.003.
[12] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Trans. Syst., Man, Cybern.*, vol. 30, no. 3, pp. 286–297, 2000, doi: 10.1109/3468.844354.

[13] B. Shneiderman, "Human-centered artificial intelligence: Reliable, safe, and trustworthy," *Int. J. Hum.-Comput. Interact.*, vol. 36, no. 6, pp. 495–504, 2020, doi: 10.1080/10447318.2020.1741118.
[14] GitHub, "Quantifying the impact of GitHub Copilot on developer productivity," GitHub Blog, 2022. [Online]. Available: https://github.blog *(Illustrative field evidence only)*
[15] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023, doi: 10.1126/science.adh2586.