

DOI: https://doi.org/10.48009/2_iis_2026_118

Who decides? Understanding the distribution of agency in human-AI decision systems

Cherie Noteboom, *Dakota State University*, cherie.noteboom@dsu.edu
Aravindh Sekar, *Dakota State University*, aravindh.sekar@trojans.dsu.edu

Abstract

As organizations increasingly embed artificial intelligence and advanced information systems into operational processes, decision-making is no longer solely a human activity. While prior research has extensively examined technology adoption, system performance, and user trust, less attention has been given to how decision-making authority, or agency, is distributed between human and machine actors. This study addresses this gap by examining agency in organizational decision systems. Drawing on qualitative data from sixteen semi-structured interviews with supply chain professionals, the study applies an agency lens to investigate how decision authority is enacted, shared, and contested in practice. The analysis identifies three empirically grounded modes of agency—human-centric, machine-supported, and machine-centric—alongside a fourth dynamic condition, contested agency, which emerges when system autonomy and human accountability become misaligned. Building on these findings, the study develops a conceptual framework organized around two dimensions—human control and machine autonomy—that explains how decision authority is distributed across organizational contexts. By empirically operationalizing distributed agency within information systems, this paper contributes a conceptual vocabulary for understanding decision authority in human–AI environments and offers practical guidance for designing governance structures that align system autonomy with organizational accountability.

Keywords: Human–AI Decision Systems, Distributed Agency, Artificial Intelligence, Sociotechnical Systems, Decision-Making, AI Governance

Introduction

When a supply chain planning system recommends shutting down a production line and the operations manager overrides it based on thirty years of experience, only for the system to prove correct, who bears responsibility for the resulting delay? This question, once hypothetical, is now increasingly common as organizations embed AI into operational decision processes. It highlights a gap in Information Systems research: not who uses AI systems, but who decides when AI systems actively participate in decision-making. As AI and advanced information systems become embedded in organizational processes, decision-making is no longer solely a human activity. In domains such as supply chain management, logistics, and operations, systems generate recommendations, automate workflows, and sometimes execute decisions independently, redistributing decision authority in ways that remain insufficiently understood.

Existing research in Information Systems has primarily focused on technology adoption, system performance, and user acceptance. Foundational models such as the Technology Acceptance Model emphasize perceived usefulness and ease of use as drivers of system adoption (Davis, 1989), while later

studies highlight the role of trust, explainability, and reliance on automated systems (Glikson & Woolley, 2020; Rai, 2020; Venkatesh et al., 2003). However, these perspectives generally assume that decision-making authority remains with human actors.

This assumption becomes increasingly problematic as AI-enabled systems move beyond supporting decisions to influencing or executing them. Algorithmic systems generate predictions, recommend actions, and automate operational decisions with minimal human intervention (Brynjolfsson & McAfee, 2017; Shrestha et al., 2019). More recently, IS scholarship identifies agentic AI—characterized by goal-directed behavior, multi-step planning, and minimal human oversight—as a distinct class of IS artifact that fundamentally redistributes decision authority within organizations (Kumar et al., 2026). Consequently, decision-making is increasingly distributed across human and technological actors, blurring traditional boundaries of authority and control.

From a theoretical perspective, this shift challenges conventional understandings of agency. Agency Theory conceptualizes agency as the capacity of actors to make decisions and act on behalf of organizational interests (Jensen & Meckling, 1976). Traditionally, agency is attributed to human actors, with technologies viewed as decision-support tools. However, as intelligent systems become increasingly autonomous and embedded within organizational processes, they influence, shape, and sometimes execute decisions, creating new tensions regarding accountability, governance, and control. Systems may generate recommendations that humans routinely accept or override based on contextual judgment (Dietvorst et al., 2015), while more advanced systems may execute decisions autonomously (Rai et al., 2019). These developments suggest that organizational decision-making is neither purely human-driven nor machine-driven but distributed across human and technological actors.

Despite these developments, there remains limited empirical understanding of how agency is enacted in organizational contexts. Specifically, little is known about how practitioners perceive decision authority, interact with system-generated decisions, and negotiate agency across varying levels of system autonomy. To address this gap, this study examines the distribution of agency in human–AI decision systems using qualitative data from sixteen semi-structured interviews with supply chain professionals. Applying an agency lens, the study examines how decision authority is enacted, shared, transferred, and contested in practice.

The study is guided by the following research question: *How is decision-making agency distributed between human and machine actors in organizational information systems?*

This paper makes three contributions. First, it reconceptualizes agency in information systems as a distributed and dynamic construct rather than a fixed human attribute. Second, it identifies distinct modes of human–AI agency in organizational decision-making. Third, it develops a conceptual framework explaining how agency shifts across varying levels of human control and machine autonomy. Together, these contributions advance understanding of decision-making in AI-enabled environments and provide practical insights for designing governance structures that align system autonomy with organizational accountability.

Theoretical Background

Agency and Decision-Making in Organization

The concept of agency has long been central to understanding decision-making within organizations. Rooted in Agency Theory, agency refers to the capacity of actors to make decisions and act on behalf of organizational objectives (Jensen & Meckling, 1976). Traditional perspectives assume that human actors—

such as managers and decision-makers—possess authority and control, while organizational structures are designed to align individual actions with collective goals.

Within this view, decision-making is inherently human-centric, and technologies are conceptualized as tools that support human judgment. Information systems provide data, enhance analytical capabilities, and improve efficiency, but decision authority remains firmly with human actors. This assumption has underpinned much of the research on organizational decision-making and governance. However, as organizations increasingly integrate advanced information systems into their operations, this human-centric view of agency becomes less tenable. Systems are no longer passive tools; they are capable of influencing, shaping, and even executing decisions. This shift challenges the traditional assumption that agency resides solely with human actors and suggests the need to reconsider agency as a more distributed construct.

Information Systems and Technology-Mediated Decision Making

Research in Information Systems has extensively explored how technologies influence organizational processes, particularly decision-making. Early studies on decision support systems emphasized the role of technology in augmenting human decision-making by providing relevant information and analytical capabilities (Keen & Scott Morton, 1978). In such systems, technology enhances human judgment but does not replace it. Subsequent developments in enterprise systems and analytics further expanded the role of information systems in organizational decision-making. These systems integrate data across organizational functions, enabling more informed and coordinated decisions (Davenport, 1998). Despite these advancements, decision authority is still generally assumed to reside with human actors, with systems serving as enablers of decision-making processes.

More recently, advances in artificial intelligence and machine learning have significantly altered this landscape. AI-enabled systems can generate predictions, recommend actions, and automate routine decisions, often with minimal human intervention (Brynjolfsson & McAfee, 2017; Shrestha et al., 2019). These capabilities enable systems to move beyond support roles and actively participate in decision processes. As a result, decision-making is increasingly mediated by interactions between human and technological actors.

However, much of the existing literature continues to focus on adoption, usability, and performance outcomes, rather than examining how decision authority itself is distributed. This creates a gap in understanding how systems influence not just what decisions are made, but who makes them.

Human-AI Interaction and Decision Authority

Emerging research on human–AI interaction has begun to explore how individuals engage with intelligent systems in decision-making contexts. Studies have highlighted the importance of trust, explainability, and transparency in shaping user reliance on AI systems (Glikson & Woolley, 2020; Rai, 2020). For example, users may rely on AI-generated recommendations when they perceive systems to be accurate and trustworthy, but may override them when confidence is low (Dietvorst et al., 2015).

While this stream of research provides valuable insights into user behavior, it often frames AI as a decision-support tool rather than as an actor with potential agency. The focus is typically on whether users accept or reject system recommendations, rather than on how decision authority is shared or transferred between human and machine actors.

Recent work has begun to recognize the more active role of AI systems in organizational decision-making. For instance, algorithmic decision systems can shape organizational processes by influencing how information is interpreted and how actions are selected (Kellogg et al., 2020). Similarly, research on digital platforms suggests that human and AI actors increasingly operate in hybrid configurations, where decision-making responsibilities are distributed across multiple entities (Rai et al., 2019).

Despite these advances, there remains limited conceptual clarity regarding how agency is enacted in such hybrid systems. Specifically, existing research does not fully explain how decision authority can be distributed differently across different levels of system autonomy or how tensions arise when human accountability is misaligned with system-driven decisions.

Toward a Distributed View of Agency

The increasing autonomy of AI-enabled systems necessitates a reconceptualization of agency as a distributed phenomenon. Traditional organizational perspectives have largely treated agency as residing

within human actors who exercise judgment, make decisions, and enact organizational goals (Eisenhardt, 1989; Jensen & Meckling, 1976). Within this view, technologies function primarily as tools that support human decision-making rather than actively shaping or participating in organizational action.

However, advances in artificial intelligence and digitally embedded systems increasingly challenge this assumption by enabling technologies to influence, guide, and in some cases autonomously execute decisions.

This perspective is consistent with a long tradition of Information Systems research demonstrating that technologies are not passive instruments but active participants in shaping organizational structures and work practices. Barley (1986) argued that technologies act as “occasions for structuring,” reorganizing workflows, communication patterns, and role relationships within organizations. Similarly, Orlikowski (1992) proposed that technologies and organizational structures are recursively constituted through human action, such that organizational outcomes emerge from the ongoing interaction between people and technological systems rather than from either in isolation. Later work further emphasized that technologies are enacted through use, with their organizational effects shaped by interpretation, context, and practice (Orlikowski, 2000, 2007).

Building on this sociotechnical tradition, sociomaterial perspectives suggest that human and technological agencies are fundamentally intertwined. Leonardi (2011) argued that human and material agencies become imbricated through ongoing organizational practice, making clear distinctions between human and technological action increasingly difficult to sustain. Likewise, Boudreau and Robey (2005) demonstrated that the implementation and use of integrated information systems simultaneously reflects human agency while restructuring the conditions under which future organizational actions occur. Together, these perspectives suggest that technologies do not merely support organizational decision-making but actively shape how decision authority is enacted and distributed.

Recent organizational research has extended these ideas into the context of intelligent technologies and AI-enabled environments. Murray, Rhymer, and Sirmon (2021) introduced the concept of conjoined agency, arguing that organizational action increasingly emerges through collaborative configurations between humans and technologies. Their work identifies multiple forms of human–technology agency, including assisting, augmenting, arresting, and automating, demonstrating that the relationship between humans and intelligent systems is neither fixed nor binary. Instead, agency varies depending on the degree of technological autonomy and the nature of organizational routines.

Scott and Orlikowski (2025) extend this sociomaterial tradition into the age of AI, arguing that agency should be understood not as a property residing in human or technological actors but as an ongoing enactment constituted through sociomaterial practices. Their work calls attention to the performative effects of AI phenomena and the accountability challenges that arise when agency is distributed across multiple,

shifting configurations of human and technological action, a concern that directly motivates the present study.

Taken together, these perspectives suggest that agency in AI-enabled environments should be understood as dynamic, distributed, and context dependent. In some situations, humans retain primary control while systems provide analytical support. In others, systems generate recommendations that strongly shape organizational decisions, even when humans formally retain authority. In more advanced contexts, systems may autonomously execute operational decisions with limited human oversight. Additionally, tensions may emerge when human judgment conflicts with system-generated outputs, leading to contested forms of agency.

Although prior research has established that technologies shape organizational action, there remains limited empirical understanding of how agency is distributed and experienced in AI-enabled decision environments. In particular, little is known about how practitioners negotiate decision authority between human and machine actors in operational settings. This study addresses this gap by examining how agency is enacted, shared, and contested within technology-mediated decision-making processes.

Research Gap and Positioning

While prior Information Systems research has established that agency is distributed across human and technological actors (Leonardi, 2011; Orlikowski & Scott, 2008), comparatively little empirical research has examined how organizations experience and enact distinct forms of distributed agency in AI-enabled decision environments. This study extends prior work by developing an empirically grounded typology of agency configurations that explains how decision authority is enacted, shared, and contested in organizational practice. In particular, existing research provides a limited understanding of how agency shifts as system autonomy increases and how tensions emerge when human accountability conflicts with system-driven decisions. This study addresses this gap by applying an agency lens to qualitative data from organizational practitioners to examine how decision authority is enacted, negotiated, and contested in practice. In doing so, the paper extends existing research by focusing on the distribution of agency within AI-enabled organizational decision environments.

Methodology

Research Design

This study adopts a qualitative research design to examine how decision-making agency is distributed between human and machine actors in organizational contexts. Sixteen semi-structured interviews were conducted with supply chain professionals directly involved in the use, implementation, or oversight of advanced information systems within their organizations. The interview data were rich in descriptions of operational workflows, system interactions, and decision-making practices, providing a strong empirical foundation for examining how decision authority is enacted in technology-mediated environments. The analytical approach applies an agency lens to interpret participants' accounts of human-system interactions. While the data were collected to examine technology-mediated decision-making broadly, the theoretical framework applied here focuses specifically on the distribution of decision authority between human and machine actors, enabling new conceptual insights to emerge from the empirical material. This approach is consistent with established interpretive qualitative research practice, which supports iterative theoretical development from rich empirical datasets (Corbin & Strauss, 2015; Gioia et al., 2013). By examining these interactions through the lens of agency, the study develops new theoretical insights regarding the distribution of decision authority in human–AI environments.

Data Collection

The dataset consists of sixteen semi-structured interviews conducted with professionals in supply chain and logistics organizations. Participants held roles such as system architects, planners, managers, and operations specialists, and were directly involved in the use, implementation, or oversight of advanced information systems within their organizations.

Semi-structured interviews were selected to capture rich, contextually grounded insights into organizational processes and decision-making practices. This method allows participants to articulate their experiences while enabling the researcher to probe deeper understanding of emerging themes (Kvale & Brinkmann, 2009). The interview protocol focused on system implementation, operational workflows, governance mechanisms, and user interactions with information systems.

Each interview was recorded and transcribed, resulting in a detailed qualitative dataset. The transcripts contain extensive accounts of how decisions are made in practice, how systems influence those decisions, and how participants perceive the role of technology in organizational processes. These data provide a strong empirical basis for examining the distribution of agency in decision-making contexts. Data collection was conducted under Institutional Review Board approval, with participants providing informed consent prior to participation. Detailed participant background information, including years of consulting and blockchain experience and the industries they represent, is presented in Figure 1.

Participant	Consulting Experience (In Years)	Blockchain Implementation Experience (In Years)	Industry
Participant 1	12	7	Waste Management, FinTech
Participant 2	25	4	Pharmaceutical Industry
Participant 3	10	7	Multiple Industries - Natural Gas, Food Supply chain
Participant 4	7	5	Retail Industry
Participant 5	5	5	Automotive
Participant 6	10	5	Shipping Industry
Participant 7	10	7	Energy
Participant 8	5	5	Food and Beverage
Participant 9	15	7	Retail and Consumer Goods
Participant 10	6	6	Pharmaceutical Industry
Participant 11	10	7	Shipping Industry
Participant 12	15	7	Manufacturing Industry
Participant 13	12	5	Natural Gas
Participant 14	5	5	Insurance Industry
Participant 15	15	10	Luxury Goods
Participant 16	5	5	Food Industry

Figure 1. Participant Profile Overview

Data Analysis: Agency – Oriented Interpretation

The analysis follows an interpretive qualitative approach, focusing on identifying patterns related to the distribution of decision-making agency. The study employs a selective coding strategy, targeting excerpts that explicitly or implicitly reflect decision-making processes and human-system interactions.

This approach is aligned with established qualitative analysis methods that emphasize iterative coding and conceptual abstraction (Corbin & Strauss, 2015; Gioia et al., 2013). The analysis proceeded in three stages:

Stage 1: Identification of Decision-Relevant Data

Interview transcripts were reviewed to identify segments related to decision-making activities, including planning, execution, monitoring, and system interaction. These excerpts included descriptions of how decisions were made, how system outputs were used, and how participants responded to system-generated recommendations.

Stage 2: Development of Agency Categories

Initial coding focused on identifying interview excerpts describing decision-making activities, system recommendations, automated execution, human overrides, accountability, and operational judgment. These excerpts were first coded using descriptive labels reflecting the locus of decision authority. Through iterative comparison across transcripts, similar codes were grouped into higher-order conceptual categories representing recurring patterns of agency. Multiple rounds of comparison and refinement resulted in four recurring modes of agency: human-centric, machine-supported, machine-centric, and contested agency. These categories emerged through continuous movement between the empirical data and the theoretical lens of agency, consistent with interpretive qualitative analysis (Corbin & Strauss, 2015; Gioia et al., 2013). Using an agency lens informed by prior literature, the identified excerpts were coded based on the locus of decision authority. Table 1 shows the four categories of agency were developed inductively through iterative analysis:

Table 2. Example Coding Structure

First-Order Concepts	Second-Order Theme	Agency Mode
Human overrides recommendation	Human judgment	Human-Centric
Review system output	Shared authority	Machine-Supported
Automated processing	Embedded decision authority	Machine-Centric
Conflict with system output	Decision tension	Contested

These categories reflect patterns observed across the dataset and are consistent with qualitative coding practices that balance inductive insights with theoretical framing (Gioia et al., 2013).

Stage 3: Thematic Synthesis and Pattern Identification

Following coding, patterns across the four agency categories were analyzed to understand how agency is distributed and how it shifts across different contexts. The analysis focused on identifying recurring themes, transitions between agency modes, and conditions under which tensions arise. Attention was given to how participants describe trust, control, and accountability in relation to system-driven decisions.

Validity and Reliability

To ensure the rigor and credibility of the analysis, several qualitative research principles were followed. First, the use of verbatim interview excerpts enhances transparency and supports the validity of interpretations by grounding findings in participant experiences (Lincoln & Guba, 1985). Second, the iterative coding process allows for the refinement of categories and ensures conceptual clarity (Corbin & Strauss, 2015). Third, the application of a clearly defined analytical lens provides consistency in how data are interpreted across the dataset.

The study further strengthens rigor by building upon a rich qualitative dataset collected through semi-structured interviews with organizational practitioners directly involved in technology-enabled decision-making processes.

Results

Distribution of Agency in Human-AI Decision Systems

The analysis reveals that decision-making in contemporary organizational settings cannot be characterized as solely human-driven or machine-driven. Instead, agency is distributed across human and technological actors, with decision authority shifting depending on system capabilities, organizational practices, and user interactions.

Consistent with prior research on technology-mediated work, decision-making emerges as a sociotechnical process, where outcomes are shaped through interactions between human judgment and system-generated inputs (Orlikowski & Scott, 2008; Rai et al., 2019). However, rather than treating systems as passive tools, the findings indicate that systems actively participate in decision processes, influencing, shaping, and in some cases executing decisions.

To capture these dynamics, through iterative coding and constant comparison of participant accounts, four recurring modes of agency emerged that describe how decision authority is distributed in practice. These empirically derived categories form the basis of the conceptual framework developed later in the paper.

Human-Centric Agency

Human-centric agency refers to situations in which decision authority remains primarily with human actors, with systems serving a supportive role. In these contexts, systems provide data, visibility, or recommendations, but human decision-makers retain control over the final decision.

Participants frequently emphasized the importance of operational experience, contextual understanding, and human judgment when interacting with technology-enabled systems. One participant explained that *“the system gives us visibility into the process, but operationally we still make the final decision because there are situations the system cannot fully account for.”*

Another participant similarly noted that *“technology helps organize the information, but experience and judgment are still critical when unexpected issues happen.”* These observations suggest that systems are viewed as tools that enhance efficiency and decision visibility without displacing human authority. Consistent with traditional decision-support perspectives, technology augments human judgment while leaving ultimate decision responsibility with organizational actors (Keen & Scott Morton, 1978). Human-centric agency was particularly evident in situations involving uncertainty, exceptions, or operational complexity where contextual interpretation remained essential.

Machine-Supported Agency

Machine-supported agency represents a hybrid configuration in which systems generate recommendations that are reviewed, validated, or approved by human actors. In this mode, decision-making authority is shared, with systems significantly influencing decisions while humans retain formal oversight.

Participants described workflows in which system-generated outputs routinely guide operational decisions due to their perceived efficiency and consistency. One participant stated that *“most of the time we rely on the recommendations generated by the system because it speeds up the process and improves consistency.”* Another participant explained that *“the planners usually review what the system suggests and approve it unless there is an operational concern.”*

These findings indicate that while humans remain formally responsible for decisions, system recommendations strongly shape organizational action. This reflects a gradual transition from purely human-centered decision-making toward technologically influenced decision processes, where agency becomes increasingly distributed between human and machine actors. Consistent with prior research on algorithmic decision-making, system recommendations become embedded within operational routines and influence how organizational decisions are enacted (Shrestha et al., 2019).

Machine-Centric Agency

Machine-centric agency occurs when systems execute decisions with minimal or no direct human intervention. In these contexts, operational workflows become highly automated, with systems acting on predefined rules, integrated logic, or learned patterns.

Participants described organizational processes in which technology independently managed transactions and operational activities once configurations and workflows had been established. One participant

explained that *“once the rules and workflows are configured, the system automatically processes transactions without requiring manual intervention.”* Another participant noted that *“a lot of operational activities are system-driven now, and people only step in when exceptions occur.”* These observations demonstrate a significant shift in decision authority from human actors toward technological systems.

While humans continue to maintain oversight responsibilities, their direct participation in routine operational decisions becomes substantially reduced. Machine-centric agency therefore reflects the increasing autonomy of intelligent systems within organizational environments and raises important questions regarding accountability, governance, and control. Prior research similarly suggests that highly automated systems can restructure organizational decision-making by embedding authority within technological processes themselves (Rai et al., 2019).

Contested Agency

Contested agency captures situations in which decision authority becomes unclear, disputed, or negotiated between human actors and technological systems. This mode is characterized by tension between system-generated recommendations and human judgment, particularly when participants perceive that systems fail to account for contextual or operational realities. Several participants described deliberately overriding system outputs because of concerns regarding reliability, practicality, or downstream consequences. One participant explained that *“there are times when the system recommendation does not match operational reality, so we override it based on experience,”* while another noted that *“the technology may suggest the most efficient option, but sometimes we know from experience that it will create downstream issues.”* These examples illustrate that agency is actively negotiated when human expertise conflicts with system-generated decisions.

Participants also expressed concerns regarding trust in automated recommendations, with one observing that *“people do not always trust the automated recommendations, especially when the system lacks the context behind the decision.”* Consistent with research on algorithm aversion and human–AI interaction (Dietvorst et al., 2015; Glikson & Woolley, 2020), contested agency emerges when system outputs conflict with human expectations, contextual understanding, or operational realities. More broadly, these tensions reflect one of the defining governance challenges of AI-enabled organizations, where increasing system autonomy raises concerns regarding accountability, explainability, and the preservation of human oversight (Engin, 2025; Goldsmith & Yang, 2025). Scott and Orlikowski (2025) similarly argue that AI shifts accountability from the traditional “problem of many hands” toward a “problem of many things,” where responsibility becomes increasingly difficult to attribute across distributed human and technological actors.

Leonardi (2025) extends this perspective by arguing that agency attribution in the age of agentic AI is a recursive process of power redistribution rather than a permanent transfer of authority. As humans attribute greater agency to AI systems, perceptions of control and responsibility shift, even when formal accountability remains with organizational actors. The contested agency observed in this study reflects this dynamic, revealing that agency is neither fixed nor exclusively human or machine-based, but continuously negotiated as organizations balance system autonomy with human responsibility. Together with the other three agency modes, contested agency demonstrates that distributed agency is both observable in practice and sufficiently dynamic to adapt as technologies, organizational routines, and governance structures evolve.

A Framework of Distributed Agency

Building on these findings, this study proposes a conceptual framework that explains how agency is distributed between human and technological actors within AI-enabled organizational environments. The

framework is structured around two dimensions: 1) Human Control (low to high) and 2) Machine Autonomy (low to high)

Together, these dimensions explain how decision authority can be distributed differently across varying levels of human involvement and system-driven decision-making.

While four modes of agency emerged empirically from the data, the analysis indicates that human-centric, machine-supported, and machine-centric agency represent relatively stable configurations of distributed decision authority, whereas contested agency reflects a cross-cutting condition that may emerge when human judgment conflicts with system-generated outputs.

The framework identifies three primary configurations of distributed agency:

1. **Human-Centric Agency** – characterized by high human control and low machine autonomy, where systems primarily support human decision-making.
2. **Machine-Supported Agency** – characterized by high human control and high machine autonomy, where systems strongly influence organizational decisions while humans retain formal oversight responsibilities.
3. **Machine-Centric Agency** – characterized by low human control and high machine autonomy, where systems execute decisions with minimal human intervention.

The framework also contains a fourth conceptual quadrant labeled Transitional or Minimal Agency, representing situations in which both human involvement and machine autonomy remain limited. Rather than reflecting a stable organizational configuration, this condition represents weak or underdeveloped decision structures where agency remains fragmented or minimally enacted.

In addition to these structural configurations, the analysis identifies Contested Agency as a dynamic condition that may emerge across any configuration when system-generated outputs conflict with human expertise, contextual judgment, or operational realities. Unlike the other agency modes, contested agency does not represent a stable organizational state but instead reflects tensions surrounding authority, accountability, trust, explainability, and operational legitimacy in AI-enabled environments. Figure 2 illustrates the framework of distributed agency and demonstrates how decision authority can be distributed differently across varying levels of human control and machine autonomy.

This study contributes to research in Information Systems by reconceptualizing agency as a distributed and dynamic construct within human–AI decision systems. Traditional perspectives, grounded in Agency Theory, assume that decision authority resides with human actors, with technologies serving as tools that support or enhance decision-making (Jensen & Meckling, 1976). However, the findings of this study challenge this assumption by demonstrating that decision-making authority is increasingly shared, transferred, and negotiated between human and technological actors.

Rather than viewing agency as a fixed attribute of individuals, this study shows that agency emerges through interactions between humans and systems. The identification of four modes of agency—human-centric, machine-supported, machine-centric, and contested—illustrates that decision authority is not binary but exists along a continuum shaped by system capabilities, organizational practices, and user perceptions.

This perspective extends prior research on sociotechnical systems, which emphasizes the interdependence of human and technological elements in shaping organizational outcomes (Orlikowski & Scott, 2008). By

focusing specifically on decision-making authority, this study adds a new dimension to understanding how technologies influence organizational processes.




Figure 2: Framework of Distributed Agency in Human-AI Decision Systems
Discussion

From Decision Support to Decision Authority

A key insight from this study is the shift from decision support to decision authority. Much of the existing literature conceptualizes AI and information systems as tools that assist human decision-makers by providing information and recommendations (Keen & Scott Morton, 1978; Shrestha et al., 2019). However, the findings suggest that in many organizational contexts, systems do more than support decisions; they actively shape or enact them.

In machine-supported agency, system-generated recommendations are frequently accepted with minimal scrutiny, effectively positioning the system as a primary driver of decision outcomes. In machine-centric agency, systems execute decisions autonomously, reducing the direct involvement of human actors. These dynamics indicate that decision authority is increasingly embedded within technological systems, even when formal accountability remains with humans.

This shift has important implications for how decision-making is conceptualized in information systems research. It suggests that scholars must move beyond models of technology adoption and usage to examine how technologies redistribute decision authority within organizations.

Tensions in Human-AI Decision Making

The identification of contested agency highlights the presence of tensions and misalignments in human-AI decision systems. These tensions arise when system outputs conflict with human judgment or when users lack trust in system recommendations. Consistent with prior research on algorithm aversion, individuals may override system-generated decisions when they perceive them to be inaccurate or unreliable (Dietvorst et al., 2015).

At the same time, increasing system autonomy introduces challenges related to accountability. When systems execute decisions, it becomes less clear who is responsible for outcomes—particularly when decisions are based on complex or opaque algorithms (Rai et al., 2019). This misalignment between system autonomy and human accountability creates uncertainty and can undermine trust in both the system and the broader decision-making process.

These findings suggest that the distribution of agency is not always stable or seamless. Instead, it involves ongoing negotiation between human and system actors, with tensions emerging when expectations regarding control, trust, and responsibility are not aligned.

Theoretical Implications

This study makes several contributions to theory.

First, it extends Agency Theory by demonstrating that agency in contemporary organizational contexts is not exclusively human but distributed across human and technological actors. This challenges traditional assumptions about the locus of decision authority and calls for a broader conceptualization of agency in technology-mediated environments.

Second, the study contributes to research on human–AI interaction by shifting the focus from adoption and trust to decision authority and control. While prior studies have examined how users interact with AI systems, this study highlights how these interactions reshape the underlying structure of decision-making.

Third, the proposed framework of distributed agency provides a conceptual foundation for future research. By identifying distinct modes of agency and the conditions under which they emerge, the study offers a basis for examining how the findings suggest that agency may vary across different organizational contexts and levels of system autonomy. The findings of this study generate several theoretical claims about how agency is distributed, negotiated, and contested in human-AI decision systems. We present the following propositions to stimulate empirical investigation and advance a research agenda focused on how decision authority is enacted and governed as AI systems become increasingly autonomous in organizational environments. Based on the analysis, the study develops the following propositions:

P1: As machine autonomy increases, human actors shift from decision-makers to exception handlers, retaining authority primarily in conditions of high uncertainty, system failure, or ethical sensitivity.

P2: Contested agency emerges when system autonomy outpaces human actors' ability to interpret, challenge, or take responsibility for system-generated decisions, creating accountability gaps.

P3: The distribution of agency in human-AI decision systems is inherently dynamic, shifting as AI capabilities expand, organizational routines adapt, and human actors renegotiate their relationship with increasingly autonomous systems.

The framework developed in this study treats agency modes as relatively stable configurations, an approach that proves useful for organizing empirical observations and generating theoretical propositions. However, more radical sociomaterial perspectives suggest that agency is never fully stable but is instead continuously enacted through practice. Scott and Orlikowski (2025) argue that agency in AI-enabled environments should be understood as an ongoing effect of sociomaterial practices rather than a property located in human or technological actors.

From this perspective, the modes identified here capture recurring patterns in how decision authority is experienced and negotiated in practice, while leaving open the question of how these patterns emerge, shift, and dissolve over time. Future research examining how specific configurations of human-AI agency develop within organizational settings, and with what consequences for accountability and governance, would deepen and extend the contributions of this study.

Practical Implications

The findings of this study have several important implications for organizational practice. As organizations increasingly integrate AI-enabled systems into operational decision-making, they must recognize that implementation involves not only technological change but also a reconfiguration of decision authority between human actors and technological systems. Managers should therefore establish clear governance structures that explicitly define when decisions should remain under human control, when AI systems may operate autonomously, and how responsibility for system-driven outcomes is assigned.

The identification of contested agency further highlights the need for organizations to promote trust, transparency, and explainability through user training, explainable AI capabilities, and clearly documented decision processes that enable employees to understand, evaluate, and appropriately challenge system-generated recommendations (Glikson & Woolley, 2020). Organizations should also implement governance mechanisms that address accountability gaps associated with increasing system autonomy, including human oversight for high-impact decisions, periodic audits of AI-supported decision processes, and ethical review procedures designed to identify potential bias, unfairness, or unintended discriminatory outcomes.

Finally, AI governance should be aligned with emerging regulatory and ethical frameworks that emphasize fairness, transparency, accountability, and responsible AI deployment. Collectively, these practices can help organizations balance the benefits of AI-enabled decision-making with the safeguards necessary to maintain stakeholder trust, organizational legitimacy, and effective governance as decision authority becomes increasingly distributed between humans and intelligent systems.

Summary of Contributions, Limitations, and Future Research

This study makes three contributions to Information Systems research. First, it reconceptualizes agency as a distributed and dynamic construct, challenging the long-held assumption that decision authority resides exclusively with human actors in technology-mediated environments. Second, it identifies four distinct modes of human–AI agency—human-centric, machine-supported, machine-centric, and contested—providing a conceptual vocabulary for examining how decision authority is enacted, shared, and negotiated in practice. Third, it advances three theoretical propositions that establish a research agenda for investigating how decision authority diverges from formal accountability, how contested agency emerges, and how agency distribution evolves as AI systems become increasingly autonomous. Together, these contributions offer a theoretically grounded foundation for future empirical investigation into human–AI decision systems.

This study should also be interpreted in light of several limitations. First, the findings are based on qualitative interview data collected within supply chain and logistics environments, which may limit transferability to other organizational or industry settings. Future research could examine whether similar forms of distributed and contested agency emerge in domains such as healthcare, finance, manufacturing, or public administration. Second, the study relies on practitioner interpretations and experiences related to human–AI decision interactions rather than direct observation of decision processes. Future studies may strengthen these insights through ethnographic observation, longitudinal analysis, system usage analytics, or experimental approaches examining human–AI decision behavior in practice.

Third, while the study identifies distinct modes of agency, organizational decision environments remain dynamic, and agency configurations may evolve over time as technologies mature and governance structures adapt. Additional longitudinal research may therefore provide deeper insight into how distributed agency changes across different stages of AI integration. Finally, this study primarily examines agency from an organizational and operational perspective. Future research may further explore broader issues

related to explainability, algorithmic accountability, ethical governance, and regulatory oversight, particularly as organizations increasingly deploy generative and autonomous AI systems.

Conclusion

The increasing integration of artificial intelligence and advanced information systems into organizational processes is fundamentally transforming how decisions are made. While prior research has focused on technology adoption, system performance, and user interaction, this study shifts attention to a more fundamental question: who holds decision-making authority in technology-mediated environments?

Drawing on qualitative data from sixteen interviews with supply chain professionals, this study reconceptualizes agency as a distributed and dynamic construct that emerges through interactions between human and machine actors. The findings identify four distinct modes of agency—human-centric, machine-supported, machine-centric, and contested—demonstrating that decision-making is neither purely human-driven nor machine-driven, but rather negotiated across multiple actors.

By developing a framework of distributed agency, the study contributes to Information Systems research by moving beyond traditional views of technology as a decision-support tool and highlighting how systems increasingly shape and enact decision authority. This perspective also extends Agency Theory by illustrating how agency is reconfigured in AI-enabled organizational contexts. The question is no longer whether AI systems will make organizational decisions. They already do. The more consequential question is whether organizations will recognize that agency has shifted before accountability gaps become governance failures. This study offers the conceptual vocabulary to ask that question rigorously, the framework to begin answering it, and the empirical grounding to make both credible.

References

- Barley, S. R. (1986). Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly*, 31(1), 78–108. <https://doi.org/10.2307/2392767>
- Boudreau, M.-C., & Robey, D. (2005). Enacting Integrated Information Technology: A Human Agency Perspective. *Organization Science*, 16(1), 3–18. <https://doi.org/10.1287/orsc.1040.0103>
- Brynjolfsson, E., & McAfee, A. (2017, July 18). The Business of Artificial Intelligence. *Harvard Business Review*. <https://hbr.org/2017/07/the-business-of-artificial-intelligence>
- Corbin, J., & Strauss, A. (2015). *Basics of Qualitative Research*. SAGE Publications Ltd. <https://uk.sagepub.com/en-gb/eur/basics-of-qualitative-research/book235578>
- Davenport, T. (1998). *Putting the Enterprise into the Enterprise System*. Retrieved April 26, 2026, from <https://hbr.org/1998/07/putting-the-enterprise-into-the-enterprise-system>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q.*, 13(3), 319–340. <https://doi.org/10.2307/249008>

- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). *Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err*.
- Eisenhardt, K. M. (1989). Agency Theory: An Assessment and Review. *The Academy of Management Review*, 14(1), 57–74. <https://doi.org/10.2307/258191>
- Engin, Z. (2026). *Human-AI Governance (HAIG): A Trust-Utility Approach* (arXiv:2505.01651). arXiv. <https://doi.org/10.48550/arXiv.2505.01651>
- Engin, Z., & Hand, D. (2025). *Toward Adaptive Categories: Dimensional Governance for Agentic AI* (arXiv:2505.11579). arXiv. <https://doi.org/10.48550/arXiv.2505.11579>
- Gioia, D. A., Corley, K., & Hamilton, A. (n.d.). *Seeking Qualitative Rigor in Inductive Research*. Retrieved April 26, 2026, from https://www.researchgate.net/publication/258174099_Seeking_Qualitative_Rigor_in_Inductive_Research
- Glikson, E., & Woolley, A. (n.d.). *Human Trust in Artificial Intelligence: Review of Empirical Research*. Retrieved April 26, 2026, from <https://www.semanticscholar.org/paper/Human-Trust-in-Artificial-Intelligence:-Review-of-Glikson-Woolley/ef0c62ff070a476f216fe478cc190c773f12a1f6>
- Goldsmith, S., & Yang, J. “Tony.” (2026). AI and the Transformation of Accountability and Discretion in Urban Governance. *Urban Governance*, S2664328626000069. <https://doi.org/10.1016/j.ugj.2026.02.005>
- Jensen, M. C., & Meckling, W. (1976). *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*.
- Keen, P. G. W., & Scott Morton, M. S. (with Internet Archive). (1978). *Decision support systems: An organizational perspective*. Reading, Mass.: Addison-Wesley Pub. Co. http://archive.org/details/decisionsupports00keen_0
- 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. <https://doi.org/10.5465/annals.2018.0174>
- Kumar, N., Wei, X., & Zhang, H. (2026). Agentic artificial intelligence as a new frontier in information systems: Promise, peril, and research opportunities. *Information & Management*, 63(3), 104317. <https://doi.org/10.1016/j.im.2026.104317>
- Kvale, S., & Brinkmann, S. (2009). *InterViews: Learning the craft of qualitative research interviewing*, 2nd ed (pp. xviii, 354). Sage Publications, Inc.
- Leonardi, P. M. (2011). When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Q.*, 35(1), 147–168.
- Leonardi, P. M. (2025). *Homo agenticus* in the age of agentic AI: Agency loops, power displacement, and the circulation of responsibility. *Information and Organization*, 35(1), 100582. <https://doi.org/10.1016/j.infoandorg.2025.100582>

- Murray, A., Rhymer, J., & Sirmon, D. G. (2021). Humans and Technology: Forms of Conjoined Agency in Organizations. *Academy of Management Review*, 46(3), 552–571. <https://doi.org/10.5465/amr.2019.0186>
- Orlikowski, W. J. (1992). The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science*, 3(3), 398–427. <https://doi.org/10.1287/orsc.3.3.398>
- Orlikowski, W. J. (2000). Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations. *Organization Science*, 11(4), 404–428. <https://doi.org/10.1287/orsc.11.4.404.14600>
- Orlikowski, W. J. (2007). Sociomaterial Practices: Exploring Technology at Work. *Organization Studies*, 28(9), 1435–1448. <https://doi.org/10.1177/0170840607081138>
- Orlikowski, W. J., & Scott, S. V. (2008). Sociomateriality: Challenging the separation of technology, work and organization. *The Academy of Management Annals*, 2(1), 433–474. <https://doi.org/10.1080/19416520802211644>
- Rai, A. (2020). *Explainable AI: from black box to glass box*. Retrieved April 26, 2026, from <https://link.springer.com/article/10.1007/s11747-019-00710-5>
- Rai, A., Constantinides, P., & Sarker, S. (2019). *Next Generation Digital Platforms: Toward Human-AI Hybrids*. Retrieved April 26, 2026, from <https://research.manchester.ac.uk/en/publications/next-generation-digital-platforms-toward-human-ai-hybrids/>
- Scott, S. V., & Orlikowski, W. J. (2025). Exploring AI-in-the-making: Sociomaterial genealogies of AI performativity. *Information and Organization*, 35(1), 100558. <https://doi.org/10.1016/j.infoandorg.2025.100558>
- Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational Decision-Making Structures in the Age of Artificial Intelligence. *California Management Review*, 61(4), 66–83. <https://doi.org/10.1177/0008125619862257>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). *User Acceptance of Information Technology: Toward a Unified View* (SSRN Scholarly Paper No. 3375136). Social Science Research Network. <https://papers.ssrn.com/abstract=3375136>