

Research Article

When to Take Authority Back: Dynamic Governance of Agentic AI Decision Rights in Automotive Supply Chains

Shuokai Wang¹ and Bo Zeng¹, *

¹ School of Economics and Management, Beijing Forestry University, Beijing, China 100091

* Corresponding author: Bo Zeng, zengbo@bjfu.edu.cn

ABSTRACT

Agentic artificial intelligence is being granted execution authority over supply chain decisions, but authority is granted once and then assumed to hold. Conditions do not hold: data pipelines degrade, disruption regimes begin, override channels fail, and agents start escalating exceptions they previously resolved. This study develops a lifecycle framework for governing agentic AI decision rights after delegation, addressing when authority must be reduced, on what evidence it may be restored, and how oscillation between authority states is prevented. Nineteen runtime trigger signals spanning readiness decay, oversight escalation, governance erosion and behavioural deviation are evaluated by four stakeholder groups using spherical fuzzy sets. Spherical fuzzy DEMATEL separates leading from lagging indicators, and the resulting structure is embedded in an authority state machine governed by an explicit hysteresis rule in which the downgrade threshold and the restoration threshold deliberately differ. Six disruption scenarios, parameterised against publicly documented events, are simulated over a 180-day horizon for seven automotive supply chain decision rights. Twelve of the nineteen signals are leading indicators and seven are lagging, and the split does not follow family lines: all four behavioural signals are lagging, as expected, but so are audit findings, override-channel unavailability and key-role vacancy. An organisation that monitors agent behaviour, or relies on internal audit, learns of authority failure only after it has occurred. Hysteresis suppresses authority thrashing from 1.69 to 1.02 changes per decision right, at the cost of 4.5 additional days at reduced authority, and a confirmation window trades a 31% reduction in thrashing against a mean exposure window of 2.95 days. The framework converts static delegation configurations into an auditable control loop and shows that the integration of enterprise, execution and security systems determines not only what an agent can do, but whether its authority can be withdrawn in time.

Keywords: Agentic AI; Decision Rights; Authority Revocation; Hysteresis; Spherical Fuzzy DEMATEL; Supply Chain Governance; Industrial Information Integration

1 Introduction

Agentic artificial intelligence (agentic AI) is entering industrial decision processes as an actor rather than an instrument. Where earlier systems predicted demand or recommended a reorder quantity, an agent now interprets an objective, plans a course of action, calls enterprise services and executes a transaction. In automotive supply chains this transition is already visible in replenishment, expediting and transportation rerouting, where the decision space is bounded, the data are structured, and the execution interface exists. The governance question that accompanies it has, until recently, been posed in a single form: which decisions may an agent be permitted to take, and how much authority should accompany that permission.

That question has been answered. Frameworks now exist for evaluating heterogeneous supply chain decision rights against readiness, criticality and control criteria, and for translating those evaluations into decision-specific autonomy levels ranging from human-led support to bounded autonomous execution [1], [2], [11]. Autonomy reference models describe progressions from manual operation to full autonomy [11], task-level instruments measure how much of a task has actually been delegated [3], [4], and runtime architectures classify individual agent actions as permitted, human-gated or blocked [13], [18]. Collectively, this literature establishes how authority should be allocated.

It does not establish what happens next. An authority configuration is a statement about conditions that held at the moment of assessment: the data pipeline was reliable, the override channel worked, the market was not in disruption, and the agent's behaviour was unremarkable. None of these conditions is stable. Interfaces fail during an enterprise resource planning migration; a semiconductor shortage converts a routine allocation decision into a contested one; a ransomware incident removes the very override channel on which the original authority grant

depended; and model drift degrades an agent's judgement long before any single decision looks wrong. In each case the configuration that was defensible at grant time has quietly ceased to be defensible, and nothing in the existing frameworks detects that or acts on it.

Practitioners resolve this informally, and the informality is the problem. Authority is withdrawn when someone senior notices something, which is late; or it is withdrawn reflexively at the first anomaly and restored as soon as the anomaly passes, which produces an agent whose permissions change weekly and which no operator can plan around. Neither behaviour is auditable. Both are consistent with the observation that governance failures in autonomous systems are rarely failures of the initial decision and usually failures of the control loop that was supposed to follow it.

This study therefore addresses the governance of decision rights after delegation. It treats an authority level not as a configuration but as a state, and asks what evidence should move that state, in which direction, and how quickly. Three design problems follow. The first is detection: among the many signals an integrated industrial information environment can produce, which ones lead an authority failure and which merely record it. The second is stability: a rule that reduces authority whenever a threshold is crossed and restores it whenever the threshold clears will oscillate under noise, producing what we term authority thrashing. The third is the trade-off that stability creates: every mechanism that suppresses oscillation also delays reaction, and delayed reaction is exposure. These problems are inseparable from industrial information integration, because a revocation rule can only be as timely as the systems that supply its inputs.

The study makes four contributions.

- Contribution 1: a decision-right lifecycle framework that extends delegation governance from grant to monitoring, adjustment, revocation and restoration, formalised as an authority state machine over the autonomy levels used in static allocation.
- Contribution 2: an empirically structured trigger architecture. Nineteen runtime signals are evaluated by four stakeholder groups and separated by spherical fuzzy DEMATEL into leading and lagging indicators, yielding a monitoring priority that differs sharply from what practitioners currently instrument.
- Contribution 3: an authority hysteresis rule in which downgrade and restoration thresholds deliberately differ, combined with a confirmation window and a minimum dwell time; simulation quantifies the resulting trade-off between thrashing, exposure and the governance cost of conservatism.
- Contribution 4: decision-specific dynamic governance profiles for seven automotive supply chain decision rights, specifying what to monitor, what triggers withdrawal and what evidence justifies restoration — an operational counterpart to static autonomy assignment.

2 Background and Related Work

Four literatures bear on the problem, and each stops short of it in a different way.

2.1.1 Delegation and human–agent authority

Information systems research frames agentic artefacts through delegation: the transfer of task-related rights and responsibilities from human to system [3]. Acceptance of that transfer depends strongly on perceived control, which implies that autonomy and control must be designed jointly rather than traded against each other [4], [6], and appropriate reliance on an automated system is itself a design outcome rather than a disposition [9]. Empirical work on human-in-the-loop configurations shows how unevenly that design is realised in practice [5]. Work on human–AI collaboration in supply chains reaches a compatible conclusion from the outcome side, showing that collaboration does not improve performance automatically and that responsible-AI mechanisms mediate the relationship [14], [20]. This literature is concerned almost exclusively with the act of granting. The reverse transfer — withdrawing rights already granted, and the organisational and technical conditions under which that withdrawal is legitimate — is largely unexamined, although organisational research on authority renegotiation suggests it is neither symmetric nor costless [10]. Organisational AI governance research describes the practices through which such authority is exercised [21] and the gap between ethical principles and deployable instruments [22].

2.1.2 Autonomy levels and their allocation

Reference models place supply chain autonomy on ordered scales [11], extending a tradition of automation-level taxonomies in human factors [7], [8], and task-level instruments measure delegation, consultation and collaboration to separate autonomy from capability [3], [4]. Allocation frameworks compare candidate decisions against readiness, criticality and governance criteria and assign autonomy non-compensatorily, so that strong technical

readiness cannot offset high oversight requirements [1], [2], [11]. These are static instruments by construction: they characterise a decision, not a running system, and the agent-based supply chain literature that precedes them shares the same framing [12], [16], [17]. A level assigned under one set of conditions carries no mechanism for detecting that those conditions have changed.

2.1.3 Runtime guardrails and agentic orchestration

A parallel line of work constrains agents at execution time. Risk-management architectures classify individual actions as automatic, human-gated or blocked according to governance rules [13]; feasibility masking and bounded exploration restrict the action space directly [18]; and policy-governed orchestration across sensor, execution and enterprise layers produces traceable control actions with audit trails [32], [33]. Runtime privacy and permission engineering for agentic systems addresses the same boundary from the legal side [19]. Cognitive digital twins pursue the same coupling between a decision layer and the systems that execute its decisions [15]. Cognitive collaboration research emphasises the same integration requirement from the architectural side [34], while interoperability and data governance are identified as structural preconditions for Industry 5.0 information integration [30], [31], [33]. These mechanisms operate on actions, one at a time. A blocked action does not change what the agent is permitted to attempt tomorrow; the authority itself is untouched, and an agent may go on being blocked indefinitely without any governance body being required to revisit its mandate.

2.1.4 Degradation, drift and rollback

Machine-learning operations research provides the missing vocabulary for deterioration. Concept and data drift degrade model performance silently and require explicit detection [26]; model governance practice defines rollback and re-approval procedures for deployed models [27], [29]; and disruption-regime detection in supply chains identifies when an environment has left the conditions under which a policy was validated [35], [39]. Data management in production pipelines is a further precondition, since a monitoring signal is only as reliable as the pipeline that supplies it [28]. This work monitors models. It does not monitor mandates: the object rolled back is a model version, not a decision right, and the decision to roll back is engineering, not governance.

2.2 Problem Definition and Research Gap

The unit of analysis is a decision right — a specific, recurring supply chain decision for which an agent holds a defined level of execution authority — observed over time within one automotive manufacturer's integrated information environment. Authority is represented on the five ordered levels used in static allocation: human-led (L0), AI-assisted (L1), human-in-the-loop (L2), human-on-the-loop (L3) and bounded autonomous (L4) [1], [2], [11]. We assume an initial allocation exists and is defensible; the question is what happens to it subsequently.

The gap is the absence of a governed transition rule. Delegation research explains the grant, autonomy models describe the levels, runtime guardrails police individual actions, and drift research detects deterioration in models. No framework connects a runtime signal to a change in a decision right, specifies the evidence required to reverse that change, or addresses the stability of the resulting control loop. In consequence, organisations deploying agentic AI possess a defensible answer to how much authority to grant and no defensible answer to when to take it back.

2.3 Design Objectives

The study is a design-oriented investigation and is organised around four objectives rather than hypotheses.

- DO1: identify the runtime signals that should govern authority adjustment, and determine their causal structure so that leading indicators can be distinguished from lagging ones.
- DO2: specify a transition rule that maps signal evidence to authority changes in both directions, with explicit conditions for restoration rather than automatic recovery.
- DO3: establish whether hysteresis suppresses authority thrashing, and quantify what it costs.
- DO4: derive decision-specific dynamic governance profiles for the seven automotive decision rights previously allocated static autonomy levels.

3 Materials and Methods

3.1 Industrial Context and System Boundary

The setting is a digitally integrated automotive supply chain in which an agentic AI layer is connected to enterprise resource planning, manufacturing execution, warehouse management and transportation management systems, together with the security monitoring and audit infrastructure that surrounds them. That infrastructure is itself an exposure: coverage of supply chain cyber security in critical-infrastructure frameworks is uneven [48], and documented compromises show how quickly a supplier-side incident removes the control channels on which an authority grant depends [49]. Seven decision rights are carried forward unchanged from the static allocation study to preserve comparability: inventory replenishment (A1), production rescheduling (A2), order expediting (A3), supplier substitution (A4), transportation rerouting (A5), component allocation (A6) and strategic sourcing (A7) [1], [2], [11]. Their initial authority levels are those the earlier analysis assigned: L4 for A1, L3 for A5, L2 for A2, A3, A4 and A6, and L1 for A7.

The system boundary is deliberately drawn at the point where authority can be enforced technically. A revocation rule is meaningless unless the platform can restrict what the agent is permitted to attempt, and a restoration rule is meaningless unless the evidence for restoration can be retrieved and audited. Both conditions are properties of the information architecture rather than of the agent, which is why the framework treats integration as constitutive of governance rather than as its precondition.




Figure 1. The decision-right lifecycle. Static allocation frameworks occupy the first stage only. This study addresses stages two to five: the monitoring of runtime signals, the confirmed adjustment of authority, its revocation under immediate or gradual rules, and its restoration under hysteresis, dwell and sign-off conditions. Restored authority re-enters monitoring, closing the loop.

3.2 Trigger Signals: Definition and Provenance

Nineteen signals were specified across four families. Readiness decay (T1) captures deterioration in the informational conditions that made delegation possible. Oversight escalation (T2) captures increases in the consequence, contestability or legal exposure of the decision itself. Governance erosion (T3) captures loss of the control mechanisms on which the authority grant depended. Behavioural deviation (T4) captures observed changes in how the agent is actually performing, and has no counterpart in static allocation frameworks, which evaluate decisions rather than deployed agents.

The set was assembled in two steps. An initial pool of [M] candidate signals was derived from the four literatures reviewed in Section 2 and from the source systems present in the reference architecture of the industrial context described below. That pool was put to the four groups in the first elicitation round with an instruction to add any signal their own function would act on, and to reject any signal that could not be obtained automatically. [X] signals were added by the panel and [Y] were rejected as requiring manual assessment, leaving the nineteen reported in Table 1.

Each signal is defined together with the system from which it is obtained, because a signal that cannot be retrieved automatically cannot govern anything in practice. Signals are normalised to a severity in [0,1] against organisation-specific reference ranges; the normalisation is a deployment parameter and not a finding of this study.

Table 1. Runtime trigger signals, their source systems and families. Every signal is defined so that it can be obtained from an existing enterprise, execution, security or audit system; signals requiring manual assessment were excluded during specification.

Code	Signal	Source system	Family
T1.1	Data timeliness decay	ERP / MES latency monitors	Readiness decay

Code	Signal	Source system	Family
T1.2	Interface or system outage	Integration middleware health	Readiness decay
T1.3	Input distribution drift	Model monitoring layer	Readiness decay
T1.4	Outcome feedback delay	Closed-loop KPI reporting	Readiness decay
T1.5	Master-data inconsistency	Master-data management	Readiness decay
T2.1	Disruption regime onset	Supply risk monitoring	Oversight escalation
T2.2	Regulatory or contractual change	Legal and contract management	Oversight escalation
T2.3	Safety or recall linkage	Quality and recall systems	Oversight escalation
T2.4	Customer priority conflict	Order management / allocation	Oversight escalation
T2.5	Scarcity escalation	Inventory and supply signals	Oversight escalation
T3.1	Audit finding	Internal audit management	Governance erosion
T3.2	Override channel unavailable	Workflow and escalation platform	Governance erosion
T3.3	Authority boundary failure	Identity and access management	Governance erosion
T3.4	Cybersecurity incident	Security monitoring (SIEM)	Governance erosion
T3.5	Key-role vacancy	Human resources / role registry	Governance erosion
T4.1	Exception escalation rate	Agent execution logs	Behavioural deviation
T4.2	Human override rate	Agent execution logs	Behavioural deviation
T4.3	Post-hoc correction rate	Transaction reversal records	Behavioural deviation
T4.4	Near-miss occurrence	Incident reporting	Behavioural deviation

3.3 Stakeholder Design and Elicitation

Twenty experts were organised into four groups of five. The first three groups mirror the composition of the static allocation study, and the selection criteria are restated here rather than carried by cross-reference — academics and researchers (G1), supply chain and operations professionals (G2), and digital technology and AI professionals (G3) — so that the two studies address comparable populations. Each participant had [N] or more years of experience in [scope] and direct responsibility for, or audit of, an automated decision process in an automotive supply chain; G4 additionally required a recognised audit or compliance qualification. [K] individuals were approached and twenty accepted. No participant had been involved in the design of the framework under assessment. A fourth group, IT governance, internal audit and compliance professionals (G4), was added for this study because the withdrawal and restoration of authority are executed by audit and compliance functions rather than by the operational or technical functions that specify it. Assessing decision attributes does not require that perspective; assessing the enforcement of controls does.

Each group deliberated and returned a single consensus linguistic judgement per item, producing four group-level matrices rather than twenty individual ones. Groups carry equal importance, $\lambda_g = 1/4$. Two rounds were used: an initial independent assessment followed by a reconciliation round in which items with unresolved disagreement were revisited and the reasoning recorded. Unresolved disagreement was operationalised as a spread of more than [X] levels on the nine-point linguistic scale within a group, or a difference of more than [Y] levels between group medians. [Z] of the 171 pairwise influence items met that criterion and were returned for a second round; convergence across the two rounds was measured by [metric], which moved from [value] to [value]. Each group completed four tasks: pairwise influence judgements among the nineteen signals; severity-to-downgrade mappings; minimum dwell time and clear-period requirements; and the hysteresis band separating the restoration threshold from the downgrade threshold.

Data provenance: expert elicitation, four consensus groups of five, two rounds, 2026; simulation inputs derived from publicly documented disruption timelines; no proprietary enterprise data were used.

3.4 Spherical Fuzzy Representation and SF-DEMATEL

Linguistic judgements are represented as spherical fuzzy numbers $\tilde{A} = (\mu, \nu, \pi)$ denoting membership, non-membership and hesitancy, subject to $\mu^2 + \nu^2 + \pi^2 \leq 1$ [50], and aggregated across groups by the spherical weighted arithmetic mean with $\lambda_g = 1/4$. Aggregated numbers are defuzzified by the score function

$S(\tilde{A}) = (2 + \mu^2 - \nu^2 - \pi^2)/3$, which preserves the ordering of the nine-level linguistic scale used for elicitation. The operator definitions and the aggregation rule follow the spherical fuzzy literature, in which the same representation supports ranking, weighting and interval-valued extensions [51], [52], [53].

The defuzzified direct-influence matrix A is normalised by $s = \max(\max_i \sum_j a_{ij}, \max_j \sum_i a_{ij})$ to give $X = A/s$, and the total-influence matrix is obtained as $T = X(I - X)^{-1}$. Writing R for row sums and C for column sums of T , prominence is $R + C$ and relation is $R - C$. Signals with $R - C > 0$ form the cause group and are interpreted here as leading indicators: they influence other signals more than they are influenced. Signals with $R - C < 0$ form the effect group and are lagging indicators, registering deterioration that has already propagated. The sign rule is the conventional DEMATEL criterion, and has been applied under intuitionistic, triangular and spherical fuzzy representations across domains including the automotive sector [54], [55], [56]. The distinction has a direct operational meaning. Monitoring budgets are finite, and an organisation that instruments only the effect group will detect every authority failure reliably and always too late.

3.5 Authority State Machine and the Hysteresis Rule

For decision right d , the composite pressure at time t aggregates the severities of the signals relevant to that right, $P_d(t) = \sum_j w_{dj} s_j(t)$, where w_{dj} is the relevance weight of signal j for right d and $\sum_j w_{dj} = 1$. Relevance weights are decision-specific: a cybersecurity incident bears on rerouting authority differently from the way it bears on strategic sourcing, which is already human-authorised.

Two downgrade thresholds are defined, θ_1 for a one-level reduction and $\theta_2 > \theta_1$ for a two-level reduction, together with an immediate-revocation set comprising safety and recall linkage (T2.3), cybersecurity incidents (T3.4) and loss of the override channel (T3.2). Signals in that set reduce authority to L1 without confirmation, on the reasoning that an agent whose override channel has failed cannot be supervised at all, whatever the state of the remaining evidence.

All other downgrades require confirmation over m consecutive periods, which prevents a single noisy observation from moving a governance state. Restoration is deliberately harder than withdrawal and requires three conditions to hold jointly: the pressure must fall below $\theta_1 - \delta$, where $\delta > 0$ is the hysteresis band; it must remain below that line for k consecutive clear periods; and a minimum dwell time τ must have elapsed since the last authority change. Restoration proceeds one level at a time, so an agent reduced by two levels must earn each level separately.

The asymmetry is the point. A symmetric rule — one that restores authority at the same threshold that withdrew it — oscillates whenever the pressure hovers near that threshold, which is precisely the condition under which governance attention is most warranted. Hysteresis converts a boundary into a band and makes the authority state a function of history as well as of current evidence.




Figure 2. The hysteresis rule. Authority is reduced when composite pressure $P(t)$ crosses the downgrade line θ and the crossing persists for m confirmation periods. It is restored only when $P(t)$ falls below the restore line $\theta - \delta$, remains clear for k periods, and the dwell time τ has elapsed. Within the shaded band the authority state depends on the direction from which the band was entered.

A disruption proviso completes the rule in the opposite direction. Where a decision right combines high response value with enforceable boundaries — transportation rerouting is the clearest case — authority may be temporarily raised by one level during a confirmed disruption regime, subject to a sunset clause that expires automatically and a mandatory post-hoc audit of every action taken under the elevated mandate. Governance that can only tighten

will be circumvented in the situations where speed matters most, and a framework that cannot accommodate that will be overridden informally rather than followed.

3.6 Simulation Design

Six scenarios were simulated over a 180-day horizon at daily resolution. Signal trajectories are constructed as ramp, hold and decay profiles with additive Gaussian noise ($\sigma = 0.035\text{--}0.045$), parameterised against publicly documented disruption timelines rather than proprietary data. The scenario profiles follow the disruption-propagation literature: ripple-effect simulation supplies the shape of S1 and S2 [36], [37], [38], [40], epidemic-outbreak modelling the duration and recovery of sustained scarcity [41], [43], and viability analysis the long-horizon framing of S4 [42]. Five scenarios represent distinct deterioration mechanisms; the sixth is diagnostic, constructed so that composite pressure oscillates around the downgrade threshold, which is the condition under which the hysteresis rule can be evaluated at all.

Table 2. Simulation scenarios. Five represent distinct deterioration mechanisms; S6 is a diagnostic case constructed to test stability rather than to represent a specific event.

Scenario	Dominant signals	Mechanism and anchor
S1 Semiconductor shortage	T2.1, T2.5, T4.1, T1.1	Sustained scarcity regime; profile follows the publicly reported 2021 automotive semiconductor shortage
S2 Carrier failure / port congestion	T1.2, T2.1, T4.2	Abrupt logistics interruption with partial recovery; anchored on documented port congestion episodes
S3 ERP/MES migration window	T1.5, T3.3, T1.2, T4.2	Planned system migration degrading master data and authority enforcement simultaneously
S4 Model drift	T1.3, T4.1, T4.3, T1.4	Slow silent deterioration over 110 days with no external disruption
S5 Ransomware incident	T3.4, T3.2, T1.2, T3.1	Acute security event removing the override channel; audit findings follow
S6 Borderline oscillation	T1.1, T2.1, T4.1, T1.4	Diagnostic case: pressure hovers at the downgrade threshold under noise

Three outcome measures are recorded. Thrashing is the number of authority changes per decision right over the horizon. The exposure window is the interval between the first threshold crossing and the executed downgrade, and measures the risk created by requiring confirmation. Governance cost is the number of days spent at reduced authority, which represents forgone automation value. The threshold, dwell and band parameters are not free. θ_1 and θ_2 follow from the severity-to-downgrade mappings elicited in task 2, τ and k from the dwell-time and clear-period requirements in task 3, and δ from the hysteresis band in task 4; the confirmation length m was set to the shortest value that suppresses single-observation transitions at the elicited noise level. Baseline parameters are $\theta_1 = 0.20$, $\theta_2 = 0.32$, $\delta = 0.05$, $m = 2$, $\tau = 7$ days and $k = 4$ periods; each is varied in the sensitivity analysis.

3.7 Reproducibility and Open Methods

Digital supply chain twins pursue a comparable objective with observed rather than constructed trajectories, and would be the natural vehicle for the deployment study proposed in Section 6 [45], [46]; blockchain-based traceability has been proposed for the audit evidence that restoration requires [47]. The simulation, the DEMATEL computation and all figures are produced by a single deterministic script with a fixed random seed. Signal trajectories, relevance weights, threshold values and the elicitation instrument are reported in full so that the analysis can be reproduced without access to the expert panel. Because the empirical inputs are expert judgements and constructed scenarios rather than observed operations, the study is a design and simulation contribution, and its numerical results should be read as properties of the proposed rule under stated conditions rather than as measurements of any installed system.

4 Results

4.1 Causal Structure of the Trigger Signals

The SF-DEMATEL analysis separates the nineteen signals into twelve leading and seven lagging indicators, and the split does not follow family lines exactly. Every readiness-decay signal (T1.1–T1.5) and every oversight-escalation

signal (T2.1–T2.5) falls in the cause group. The governance-erosion family divides: authority-boundary failure (T3.3) and cybersecurity incidents (T3.4) are leading, whereas audit findings (T3.1), override-channel unavailability (T3.2) and key-role vacancy (T3.5) are lagging. All four behavioural-deviation signals are lagging, and by a wide margin: their relation values range from -1.19 to -1.04, against -0.061 for the least negative non-behavioural effect signal.

The behavioural signals nonetheless carry the highest prominence values in the system — exception escalation (6.86), human override (6.83) and post-hoc correction (6.71) — which is precisely why they are misleading as governance triggers. They are the most connected signals in the network because everything else flows into them. An organisation that instruments agent behaviour, which is the cheapest and most obvious thing to instrument, will observe the strongest correlation with authority failure and will consistently observe it after the failure has propagated.

The position of audit findings is the study's least expected result and its most uncomfortable. Internal audit is the function that formally owns control effectiveness, yet T3.1 sits in the effect group ($R - C = -0.061$), meaning that an audit finding predominantly records deterioration that other signals have already registered. The same holds for override-channel unavailability (T3.2), which is lagging in the causal network while remaining in the immediate-revocation set: it is a poor early warning and an unarguable reason to act at once. Lagging status is therefore an argument about monitoring priority, not about severity.

Within the cause group, input distribution drift (T1.3), data timeliness decay (T1.1) and outcome feedback delay (T1.4) combine the highest prominence with positive relation, making them the natural targets for continuous monitoring investment. Scarcity escalation (T2.5) and regulatory change (T2.2) show the strongest relation values but lower prominence: they drive the system forcefully when they occur, yet they are episodic rather than continuous, and are better handled by event subscription than by continuous measurement.

One qualification belongs with this result rather than after it. The behavioural signals are lagging by a wide margin, and that finding is robust to any plausible perturbation of the elicited judgements: their relation values lie an order of magnitude below zero. The governance-erosion family is not. Its five signals span $R - C$ from -0.116 to +0.181, so their assignment to the cause or the effect group rests on differences far smaller than those separating the behavioural signals from zero. We therefore treat the behavioural result as a finding, and the governance-erosion split as an indication that warrants confirmation on a larger panel.




Figure 3. SF-DEMATEL causal structure of the nineteen trigger signals. Signals above the horizontal axis form the cause group (leading indicators); those below form the effect group (lagging indicators). Marker size is proportional to prominence $R + C$. All behavioural-deviation signals are lagging despite having the highest prominence in the system, while three governance-erosion signals — audit findings, override-channel availability and key-role vacancy — are also lagging.

Table 3. SF-DEMATEL prominence and relation values for all trigger signals.

Code	Signal	$R + C$	$R - C$	Role
T1.3	Input distribution drift	6.260	+0.397	Cause
T1.1	Data timeliness decay	6.238	+0.338	Cause
T1.4	Outcome feedback delay	6.218	+0.207	Cause
T1.5	Master-data inconsistency	6.017	+0.364	Cause
T1.2	Interface or system outage	6.003	+0.395	Cause
T2.5	Scarcity escalation	6.067	+0.494	Cause
T2.4	Customer priority conflict	6.018	+0.555	Cause
T2.1	Disruption regime onset	6.017	+0.539	Cause
T2.2	Regulatory or contractual change	5.969	+0.619	Cause
T2.3	Safety or recall linkage	5.965	+0.472	Cause
T3.4	Cybersecurity incident	6.061	+0.181	Cause
T3.2	Override channel unavailable	5.936	-0.083	Effect
T3.1	Audit finding	5.914	-0.061	Effect
T3.3	Authority boundary failure	5.910	+0.140	Cause
T3.5	Key-role vacancy	5.872	-0.116	Effect
T4.1	Exception escalation rate	6.855	-1.167	Effect
T4.2	Human override rate	6.828	-1.189	Effect
T4.3	Post-hoc correction rate	6.707	-1.045	Effect
T4.4	Near-miss occurrence	6.603	-1.039	Effect

Note: R and C are row and column sums of the total-influence matrix $T = X(I - X)^{-1}$. Cause group: $R - C > 0$. Prominence $R + C$ indicates total involvement in the signal network.

4.2 Authority Trajectories under Disruption

Under baseline parameters, the state machine produces differentiated and interpretable behaviour across the six scenarios. The semiconductor shortage (S1) affects six of the seven decision rights, with order expediting (A3) reduced by two levels to L0 and spending 92 of 180 days at reduced authority — the most severe outcome in the study, reflecting the combination of scarcity, customer priority conflict and rising exception escalation that a shortage produces simultaneously. Strategic sourcing (A7) is unaffected in every scenario, which is the expected result: at L1 the agent holds no execution authority to withdraw.

The scenarios differ in structure as well as severity. The ransomware incident (S5) produces the fastest response, since loss of the override channel triggers immediate revocation without confirmation; inventory replenishment falls from L4 to L1 and accumulates 28 days at reduced authority. Model drift (S4) produces the slowest, with a nine-day exposure window for production rescheduling: no external event marks its onset, and the composite pressure rises so gradually that confirmation is repeatedly reset by noise. The ERP migration scenario (S3) is the most instructive for information integration, because authority is withdrawn from three decision rights not by any external disruption but by the organisation's own planned system work degrading master data and authority enforcement at the same time.

Table 4. Authority outcomes by scenario under baseline parameters ($\theta_1 = 0.20$, $\theta_2 = 0.32$, $\delta = 0.05$, $m = 2$, $\tau = 7$, $k = 4$), over a 180-day horizon.

Scenario	Rights affected	Which	Deepest reduction	Max days reduced	Max exposure (days)
S1	5	A1, A2, A3, A5, A6	A3 → L0	92	1
S2	4	A2, A3, A4, A5	A2 → L0	55	3

Scenario	Rights affected	Which	Deepest reduction	Max days reduced	Max exposure (days)
S3	2	A2, A5	A2 → L1	51	3
S4	2	A1, A2	A2 → L1	48	9
S5	2	A1, A5	A1 → L1	28	0
S6	2	A3, A5	A3 → L1	82	13




Figure 4. Semiconductor shortage scenario (S1). Upper panel: severity trajectories of the four dominant trigger signals. Lower panel: composite pressure for three decision rights against the one-level and two-level downgrade thresholds, with the resulting authority trajectory for order expediting (A3) on the right axis. A3 is reduced to L0 as pressure exceeds θ_2 and is restored in two separate steps as the scarcity regime decays.

4.3 Stability, Exposure and the Cost of Conservatism

The diagnostic scenario (S6) confirms that authority thrashing is a real failure mode rather than a hypothetical one. Without hysteresis or confirmation ($\delta = 0$, $m = 1$), the mean number of authority changes per decision right across all scenarios is 1.69, and transportation rerouting alone changes state six times in the borderline scenario. Introducing both mechanisms at their strongest tested settings ($\delta = 0.08$, $m = 5$) reduces this to 1.02, a 39% reduction.

Neither mechanism is free, and they are costly in different currencies. The confirmation window buys stability with time: raising m from 1 to 5 lengthens the mean exposure window from 0 to 2.95 days, during which the agent continues to act under an authority level the evidence no longer supports. The hysteresis band buys stability with forgone automation: raising δ from 0 to 0.08 at $m = 2$ increases the mean time at reduced authority from 19.5 to 24.1 days, because authority that has been withdrawn is harder to earn back. The two costs fall on different parties, which is why the choice between them is a governance decision rather than a tuning exercise.

Table 5. Full parameter grid: mean outcomes over six scenarios $\times$ seven decision rights (42 right-scenario pairs) for each combination of hysteresis band δ and confirmation window m , with $\tau = 7$ and $k = 4$ held constant.

δ	m	Changes per right	Exposure (days)	Days reduced
0.00	1	1.690	0.00	22.95
0.00	2	1.214	0.93	19.55
0.00	3	1.167	1.95	18.76

δ	m	Changes per right	Exposure (days)	Days reduced
0.00	5	1.167	2.95	17.81
0.02	1	1.357	0.00	23.45
0.02	2	1.214	0.93	19.93
0.02	3	1.167	1.95	19.14
0.02	5	1.119	2.95	18.33
0.05	1	1.310	0.00	25.14
0.05	2	1.214	0.93	21.48
0.05	3	1.167	1.95	20.69
0.05	5	1.119	2.95	19.88
0.08	1	1.119	0.00	27.45
0.08	2	1.024	0.93	24.14
0.08	3	1.024	1.95	23.12
0.08	5	1.024	2.95	22.12

Note: ↓ lower is better for changes and exposure; days reduced represents forgone automation and is minimised subject to acceptable exposure. Baseline configuration is $\delta = 0.05$, $m = 2$.

Dwell time and clear-period requirements behave as expected and were varied separately. Extending the dwell time τ from 1 to 14 days raises the mean days at reduced authority from 20.9 to 22.9, and extending the clear-period requirement k from 2 to 8 raises it from 21.4 to 22.9. Neither materially affects thrashing once δ and m are set, which indicates that the two restoration delays are substitutes rather than complements: an organisation should choose one and calibrate it, not impose both conservatively.

4.4 Dynamic Governance Profiles

The analysis yields a dynamic profile for each decision right, specifying what to monitor, what withdraws authority and what restores it. Table 6 is the operational counterpart to a static autonomy assignment: the static level states where authority begins, and the profile states the conditions under which it does not remain there.

Table 6. Dynamic governance profiles for the seven automotive supply chain decision rights. Static levels are carried forward from the prior allocation study; the remaining columns are outputs of this study.

Right	Static level	Leading signals to monitor	Immediate revocation	Restoration evidence
A1 Inventory replenishment	L4	T1.1, T1.3, T2.5, T3.3	T3.2, T3.4	4 clear periods; reconciliation of open commitments; audit sign-off
A2 Production rescheduling	L2	T1.2, T1.3, T2.1, T3.3	T2.3, T3.4	4 clear periods; schedule stability confirmed over one cycle
A3 Order expediting	L2	T1.1, T2.1, T2.5, T3.3	T3.2, T3.4	4 clear periods; supplier confirmation latency within range
A4 Supplier substitution	L2	T2.1, T2.2, T3.1, T1.5	T2.3	4 clear periods; approved-vendor list revalidated
A5 Transportation rerouting	L3	T1.1, T1.2, T2.1, T3.3	T3.2, T3.4	4 clear periods; carrier interface health restored
A6 Component allocation	L2	T2.4, T2.5, T2.3, T3.1	T2.3	4 clear periods; allocation policy reconfirmed with customer commitments
A7 Strategic sourcing	L1	T2.2, T3.1, T3.5	—	not applicable: no execution authority held




Figure 5. Parameter trade-offs across all six scenarios and seven decision rights. (a) Thrashing against exposure as the confirmation window m varies, for four hysteresis bands δ . (b) Governance cost, measured as mean days at reduced authority, as a function of δ for each value of m .

Note: leading signals are drawn from the cause group in Table 3 and weighted by decision-specific relevance. Restoration additionally requires the hysteresis and dwell conditions of the state machine in all cases. A5 is eligible for the temporary elevation proviso during a confirmed disruption regime.

5 Discussion

5.1 Scientific and Technical Implications

The separation between leading and lagging signals is the study's most consequential empirical result, and it contradicts common monitoring practice on two fronts. Organisations deploying agentic systems instrument agent behaviour first, because execution logs are available, cheap and unambiguous. The DEMATEL structure shows why that instinct produces late governance: escalation rates, override rates and correction rates are the most prominent signals in the network precisely because they are downstream of everything else. They are excellent evidence that authority has already been misplaced and poor evidence that it is about to be. The behavioural signals are also the ones most exposed to the overseer's own dispositions: reliance on an algorithm falls sharply after an observed error [23] and recovers when the overseer can modify rather than merely accept its output [24], while in other conditions algorithmic advice is preferred to human advice outright [25]. Override and correction rates therefore measure the overseer as much as the agent. Effective monitoring inverts the intuition, placing investment on data timeliness, distribution drift and feedback latency — signals that are less obviously about governance and are typically owned by engineering rather than by risk functions. The position of audit findings sharpens the point: the instrument most closely associated with control assurance is itself a lagging indicator, which is an argument for continuous signal-based monitoring alongside periodic audit rather than in place of it.

The hysteresis result contributes to the delegation literature in a more specific way. Research on human–AI delegation treats perceived control as central to acceptance [4], [6], and control is usually operationalised as the ability to intervene. The simulations suggest a second dimension: predictability. An authority configuration that changes 1.69 times per right under noise offers formal intervention rights while being operationally unusable, because no operator can build a routine around a mandate that may differ next week. Hysteresis makes the authority state a function of history rather than of instantaneous evidence, and in doing so trades a small, quantified increase in exposure for a large gain in predictability. This is a design principle borrowed from control engineering rather than from information systems, and it appears to transfer; control-theoretic formulations have been applied to supply chain operations before, though to flows rather than to mandates [44].

Third, the framework demonstrates why authority governance is inseparable from information integration [30], [31], [33]. Every element of the rule depends on the architecture: the leading indicators originate in middleware and model-monitoring layers; immediate revocation depends on identity and access management being able to reduce an agent's permissions within one cycle; and restoration depends on audit evidence being retrievable and attributable. The ERP migration scenario makes the dependency concrete in the opposite direction, showing an organisation withdrawing authority from three decision rights because of its own planned system work. Integration is not merely an enabler of delegation; the same integration determines whether delegation can be reversed.

5.2 Industrial and Managerial Implications

Four implications follow for organisations deploying agentic AI in supply chain operations. First, an authority grant should be issued with its revocation conditions attached, in the same document and at the same time; a delegation decision that does not specify what would reverse it is incomplete. Second, monitoring investment should follow the cause group, which means that the business case for data-quality monitoring is partly a governance case and should be argued as one. Third, the hysteresis parameters are policy choices with distributional consequences: a longer confirmation window transfers risk to operations, whereas a wider hysteresis band transfers cost to the automation business case, and the choice should be made explicitly by the body that owns the risk rather than defaulted by the engineering team that implements it. Fourth, restoration deserves as much procedural attention as withdrawal. In practice, authority tends to be withdrawn decisively and restored informally, which produces a governance record that documents every reduction and no recovery.

The temporary elevation proviso addresses a failure mode observed in adjacent domains. Governance frameworks that can only tighten are circumvented when speed matters, and circumvention is worse than a documented, time-limited, audited elevation. A rule that permits transportation rerouting to operate at bounded autonomy during a confirmed disruption, with automatic expiry, keeps the exception inside the governance system instead of driving it outside.

5.3 Limitations and Threats to Validity

The study's principal limitation is that its empirical inputs are expert judgements and constructed scenarios rather than observations of a running agentic system. Scenario trajectories are parameterised against publicly documented disruption timelines, but they remain stylised, and the reported values are properties of the rule under stated conditions rather than measurements of any installation. Internal validity is supported by the deterministic and fully reported computation; external validity is not established and requires deployment evidence.

Three further threats deserve statement. The relevance weights linking signals to decision rights are specified by the panel and not derived from incident data, so a different weighting could alter which rights are affected in a given scenario, although it would not change the causal structure of the signals themselves. The consensus elicitation design preserves distinct professional perspectives at the group level but suppresses within-group disagreement and permits no inter-rater reliability statistic. Finally, the framework assumes that authority can be enforced technically and revoked within one monitoring cycle; in organisations where permissions are changed through a manual change-management process, the exposure windows reported here are lower bounds.

6 Conclusion and Future Work

Delegating decision authority to an agentic AI system is a decision that ages. This study asked when that authority should be taken back and on what evidence it may be returned, and answered by formalising the decision right as a governed state rather than a configuration. Nineteen runtime signals were structured by spherical fuzzy DEMATEL into leading and lagging indicators; the resulting evidence base was embedded in an authority state machine with an explicit hysteresis rule; and six disruption scenarios were simulated across seven automotive supply chain decision rights.

Three findings stand out. The signals organisations most readily instrument — exception escalation, human override, post-hoc correction and audit findings — are all lagging indicators, so behaviour-based and audit-based monitoring both guarantee late governance. Authority thrashing is a genuine failure mode, reaching 1.69 changes per decision right under a symmetric rule, and hysteresis suppresses it to 1.02 at a cost of 4.5 additional days at reduced authority. And the mechanisms that stabilise authority impose costs in different currencies — exposure for confirmation delay, forgone automation for hysteresis — which makes their calibration a governance decision rather than a technical one.

Three directions follow. The framework requires validation against a running deployment, where actual incident records would replace expert-specified relevance weights and observed authority changes would test the simulated trade-off curves. The single-organisation boundary should be extended, since decisions such as rerouting and supplier substitution cross legal boundaries and the authority to revoke may not rest with the party bearing the consequence. And the oversight load created by reduced authority states deserves study in its own right: withdrawing authority from three decision rights simultaneously, as the ERP migration scenario does, returns a substantial approval burden to human staff at precisely the moment when the systems supporting them are least reliable.

Author Contributions (CRedit)

Conceptualization, S.W. and B.Z.; methodology, S.W. and B.Z.; software, S.W.; validation, B.Z.; formal analysis, S.W.; investigation, S.W.; data curation, S.W.; writing — original draft, S.W.; writing — review and editing, B.Z.; visualization, S.W.; supervision, B.Z. Both authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Use of Generative AI or AI-Assisted Technologies

The authors used Claude (Anthropic) to assemble and verify the reference list against publisher records, to check internal consistency between the abstract and the reported results, and to draft the methodological disclosures on participant selection, elicitation convergence and signal provenance. The authors reviewed every reference, every figure and every claim, conducted the expert elicitation and the simulation themselves, and take full responsibility for the content of this manuscript.

Acknowledgments

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] K. Heinrich, C. Janiesch, O. Krancher, P. Stahmann, J. Wanner, and P. Zschech, "Decision factors for the selection of AI-based decision support systems: the case of task delegation in prognostics," *PLOS ONE*, vol. 20, no. 7, art. e0328411, 2025. doi: 10.1371/journal.pone.0328411.
- [2] Y. Jia, H. Evans, Z. Porter, S. Graham, J. McDermid, T. Lawton, D. Snead, and I. Habli, "The case for delegated AI autonomy for human-AI teaming in healthcare (preprint)," *arXiv, preprint 2503.18778*, 2025. doi: 10.48550/arXiv.2503.18778.
- [3] A. Baird, and L. M. Maruping, "The next generation of research on IS use: a theoretical framework of delegation to and from agentic IS artifacts," *MIS Quarterly*, vol. 45, no. 1, pp. 315–341, 2021. doi: 10.25300/MISQ/2021/15882.
- [4] A. Fügener, J. Grahl, A. Gupta, and W. Ketter, "Cognitive challenges in human-artificial intelligence collaboration: investigating the path toward productive delegation," *Information Systems Research*, vol. 33, no. 2, pp. 678–696, 2022. doi: 10.1287/isre.2021.1079.
- [5] T. Grønsund, and M. Aanestad, "Augmenting the algorithm: emerging human-in-the-loop work configurations," *Journal of Strategic Information Systems*, vol. 29, no. 2, art. 101614, 2020. doi: 10.1016/j.jsis.2020.101614.
- [6] D. Sele, and M. Chugunova, "Putting a human in the loop: increasing uptake, but decreasing accuracy of automated decision-making," *PLOS ONE*, vol. 19, no. 2, art. e0298037, 2024. doi: 10.1371/journal.pone.0298037.
- [7] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Trans. Systems, Man, and Cybernetics - Part A*, vol. 30, no. 3, pp. 286–297, 2000. doi: 10.1109/3468.844354.
- [8] T. B. Sheridan, and R. Parasuraman, "Human-automation interaction," *Reviews of Human Factors and Ergonomics*, vol. 1, no. 1, pp. 89–129, 2005. doi: 10.1518/155723405783703082.
- [9] J. D. Lee, and K. A. See, "Trust in automation: designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004. doi: 10.1518/hfes.46.1.50_30392.
- [10] J. Laux, "Institutionalised distrust and human oversight of artificial intelligence: towards a democratic design of AI governance under the European Union AI Act," *AI & Society*, vol. 39, pp. 2853–2866, 2024. doi: 10.1007/s00146-023-01777-z.
- [11] L. Xu, S. Mak, Y. Proselkov, and A. Brintrup, "Towards autonomous supply chains: definition, characteristics, conceptual framework, and autonomy levels," *Journal of Industrial Information Integration*, vol. 42, art. 100698, 2024. doi: 10.1016/j.jii.2024.100698.
- [12] L. Xu, S. Mak, and A. Brintrup, "Will bots take over the supply chain? Revisiting agent-based supply chain automation," *International Journal of Production Economics*, vol. 241, art. 108279, 2021. doi: 10.1016/j.ijpe.2021.108279.
- [13] V. Jannelli, S. Schoepf, M. Bickel, T. Netland, and A. Brintrup, "Agentic LLMs in the supply chain: towards autonomous multi-agent consensus-seeking," *International Journal of Production Research*, in press, 2025. doi: 10.1080/00207543.2025.2604311.
- [14] A. Samuels, "Examining the integration of artificial intelligence in supply chain management from Industry 4.0 to 6.0," *Frontiers in Artificial Intelligence*, vol. 7, art. 1477044, 2025. doi: 10.3389/frai.2024.1477044.
- [15] K. Kalaboukas, J. Rožanec, A. Košmerlj, D. Kiritsis, and G. Arampatzis, "Implementation of cognitive digital twins in connected and agile supply networks: an operational model," *Applied Sciences*, vol. 11, no. 9, art. 4103, 2021. doi: 10.3390/app11094103.
- [16] S. Murugesan, "The rise of agentic AI: implications, concerns, and the path forward," *IEEE Intelligent Systems*, vol. 40, no. 2, pp. 8–14, 2025. doi: 10.1109/MIS.2025.3544940.
- [17] L. Hughes, Y. K. Dwivedi, et al., "AI agents and agentic systems: a multi-expert analysis," *Journal of Computer Information Systems*, vol. 65, no. 4, pp. 489–517, 2025. doi: 10.1080/08874417.2025.2483832.
- [18] N. Kolt, "Governing AI agents," *SSRN Working Paper*, 2025. doi: 10.2139/ssrn.4772956.
- [19] K. Navaie, "From rights to runtime: privacy engineering for agentic AI," *AI Magazine*, vol. 46, no. 4, art. e70036, 2025. doi: 10.1002/aaai.70036.
- [20] E. Papagiannidis, P. Mikalef, and K. Conboy, "Responsible artificial intelligence governance: a review and research framework," *Journal of Strategic Information Systems*, vol. 34, no. 2, art. 101885, 2025. doi: 10.1016/j.jsis.2024.101885.
- [21] E. Papagiannidis, I. M. Enholm, C. Dremel, P. Mikalef, and J. Krogstie, "Toward AI governance: identifying best practices and potential barriers and outcomes," *Information Systems Frontiers*, vol. 25, pp. 123–141, 2023. doi: 10.1007/s10796-022-10251-y.
- [22] E. Prem, "From ethical AI frameworks to tools: a review of approaches," *AI and Ethics*, vol. 3, no. 1, pp. 699–716, 2023. doi: 10.1007/s43681-023-00258-9.

- [23] B. J. Dietvorst, J. P. Simmons, and C. Massey, "Algorithm aversion: people erroneously avoid algorithms after seeing them err," *Journal of Experimental Psychology: General*, vol. 144, no. 1, pp. 114–126, 2015. doi: 10.1037/xge0000033.
- [24] B. J. Dietvorst, J. P. Simmons, and C. Massey, "Overcoming algorithm aversion: people will use imperfect algorithms if they can (even slightly) modify them," *Management Science*, vol. 64, no. 3, pp. 1155–1170, 2018. doi: 10.1287/mnsc.2016.2643.
- [25] J. M. Logg, J. A. Minson, and D. A. Moore, "Algorithm appreciation: people prefer algorithmic to human judgment," *Organizational Behavior and Human Decision Processes*, vol. 151, pp. 90–103, 2019. doi: 10.1016/j.obhdp.2018.12.005.
- [26] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, art. 44, 2014. doi: 10.1145/2523813.
- [27] A. Paleyes, R. G. Urma, and N. D. Lawrence, "Challenges in deploying machine learning: a survey of case studies," *ACM Computing Surveys*, vol. 55, no. 6, art. 114, 2022. doi: 10.1145/3533378.
- [28] N. Polyzotis, S. Roy, S. E. Whang, and M. Zinkevich, "Data management challenges in production machine learning," *Proc. ACM SIGMOD*, 1723–1726, 2017. doi: 10.1145/3035918.3054782.
- [29] L. E. Lwakatare, A. Raj, I. Crnkovic, J. Bosch, and H. H. Olsson, "Large-scale machine learning systems in real-world industrial settings: a review of challenges and solutions," *Information and Software Technology*, vol. 127, art. 106368, 2020. doi: 10.1016/j.infsof.2020.106368.
- [30] L. D. Xu, "Enterprise systems: state-of-the-art and future trends," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 4, pp. 630–640, 2011. doi: 10.1109/TII.2011.2167156.
- [31] L. D. Xu, W. He, and S. Li, "Internet of Things in industries: a survey," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2233–2248, 2014. doi: 10.1109/TII.2014.2300753.
- [32] A. Bagozi, D. Bianchini, M. Garda, M. Melchiori, and A. Rula, "A multi-layered data service model for cyber-physical production networks," *Journal of Industrial Information Integration*, vol. 46, art. 101043, 2025. doi: 10.1016/j.jii.2025.101043.
- [33] Y. Chen, P. Jian, Y. Zhang, et al., "A systematic solution of distributed and trusted chain-network integration," *Journal of Industrial Information Integration*, vol. 41, art. 100664, 2024. doi: 10.1016/j.jii.2024.100664.
- [34] N. Hakam, and K. Benfriha, "Enhancement of industrial information systems through AI models," *Journal of Industrial Information Integration*, vol. 43, art. 100744, 2025. doi: 10.1016/j.jii.2024.100744.
- [35] A. Dolgui, D. Ivanov, and B. Sokolov, "Ripple effect in the supply chain: an analysis and recent literature," *International Journal of Production Research*, 56(1–2) 414–430, 2018. doi: 10.1080/00207543.2017.1387680.
- [36] D. Ivanov, "Simulation-based ripple effect modelling in the supply chain," *International Journal of Production Research*, vol. 55, no. 7, pp. 2083–2101, 2017. doi: 10.1080/00207543.2016.1275873.
- [37] D. Ivanov, B. Sokolov, and A. Dolgui, "The ripple effect in supply chains: trade-off efficiency-flexibility-resilience in disruption management," *International Journal of Production Research*, vol. 52, no. 7, pp. 2154–2172, 2014. doi: 10.1080/00207543.2013.858836.
- [38] D. Ivanov, A. Dolgui, and B. Sokolov, "The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics," *International Journal of Production Research*, vol. 57, no. 3, pp. 829–846, 2019. doi: 10.1080/00207543.2018.1488086.
- [39] A. Dolgui, and D. Ivanov, "Ripple effect and supply chain disruption management: new trends and research directions," *International Journal of Production Research*, vol. 59, no. 1, pp. 102–109, 2021. doi: 10.1080/00207543.2021.1840148.
- [40] S. Hosseini, D. Ivanov, and A. Dolgui, "Ripple effect modelling of supplier disruption: integrated Markov chain and dynamic Bayesian network approach," *International Journal of Production Research*, vol. 58, no. 11, pp. 3284–3303, 2020. doi: 10.1080/00207543.2019.1661538.
- [41] D. Ivanov, "Predicting the impacts of epidemic outbreaks on global supply chains: a simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case," *Transportation Research Part E*, vol. 136, art. 101922, 2020. doi: 10.1016/j.tre.2020.101922.
- [42] D. Ivanov, "Viable supply chain model: integrating agility, resilience and sustainability perspectives," *Annals of Operations Research*, vol. 319, no. 1, pp. 1411–1431, 2022. doi: 10.1007/s10479-020-03640-6.
- [43] M. M. Queiroz, D. Ivanov, A. Dolgui, and S. Fosso Wamba, "Impacts of epidemic outbreaks on supply chains: mapping a research agenda amid the COVID-19 pandemic through a structured literature review," *Annals of Operations Research*, vol. 319, no. 1, pp. 1159–1196, 2022. doi: 10.1007/s10479-020-03685-7.
- [44] D. Ivanov, S. Sethi, A. Dolgui, and B. Sokolov, "A survey on control theory applications to operational systems, supply chain management, and Industry 4.0," *Annual Reviews in Control*, vol. 46, pp. 134–147, 2018. doi: 10.1016/j.arcontro1.2018.10.014.

- [45] D. Ivanov, and A. Dolgui, "A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0," *Production Planning & Control*, vol. 32, no. 9, pp. 775–788, 2021. doi: 10.1080/09537287.2020.1768450.
- [46] D. Ivanov, A. Dolgui, A. Das, and B. Sokolov, "Digital supply chain twins: managing the ripple effect, resilience, and disruption risks," *Handbook of Ripple Effects in the Supply Chain*, Springer, 309–332, 2019. doi: 10.1007/978-3-030-14302-2_15.
- [47] S. Saberi, M. Kouhizadeh, J. Sarkis, and L. Shen, "Blockchain technology and its relationships to sustainable supply chain management," *International Journal of Production Research*, vol. 57, no. 7, pp. 2117–2135, 2019. doi: 10.1080/00207543.2018.1533261.
- [48] C. Topping, A. Dwyer, O. Michalec, B. Craggs, and A. Rashid, "Beware suppliers bearing gifts! Analysing coverage of supply chain cyber security in critical national infrastructure frameworks," *Computers & Security*, vol. 108, art. 102324, 2021. doi: 10.1016/j.cose.2021.102324.
- [49] L. Lazarovitz, "Deconstructing the SolarWinds breach," *Computer Fraud & Security*, vol. 2021, no. 6, pp. 17–19, 2021. doi: 10.1016/S1361-3723(21)00065-8.
- [50] F. Kutlu Gündoğdu, and C. Kahraman, "Spherical fuzzy sets and spherical fuzzy TOPSIS method," *J. Intelligent & Fuzzy Systems*, vol. 36, no. 1, pp. 337–352, 2019. doi: 10.3233/JIFS-181401.
- [51] F. Kutlu Gündoğdu, and C. Kahraman, "A novel VIKOR method using spherical fuzzy sets and its application to warehouse site selection," *J. Intelligent & Fuzzy Systems*, vol. 37, no. 1, pp. 1197–1211, 2019. doi: 10.3233/JIFS-182651.
- [52] F. Kutlu Gündoğdu, and C. Kahraman, "A novel fuzzy TOPSIS method using emerging interval-valued spherical fuzzy sets," *Engineering Applications of Artificial Intelligence*, vol. 85, pp. 307–323, 2019. doi: 10.1016/j.engappai.2019.06.003.
- [53] F. Kutlu Gündoğdu, and C. Kahraman, "A novel spherical fuzzy analytic hierarchy process and its renewable energy application," *Soft Computing*, vol. 24, no. 6, pp. 4607–4621, 2020. doi: 10.1007/s00500-019-04222-w.
- [54] H. S. Kilic, A. E. Demirci, and D. Delen, "An integrated decision analysis methodology based on IF-DEMATEL and IF-ELECTRE for personnel selection," *Decision Support Systems*, vol. 137, art. 113360, 2020. doi: 10.1016/j.dss.2020.113360.
- [55] V. Suzan, and H. Yavuzer, "A Fuzzy DEMATEL method to evaluate the most common diseases in internal medicine," *International Journal of Fuzzy Systems*, vol. 22, pp. 2385–2395, 2020. doi: 10.1007/s40815-020-00921-x.
- [56] J. Li, C. H. Wu, C. W. Chen, Y. F. Huang, and C. T. Lin, "Apply Fuzzy DEMATEL to explore the decisive factors of the auto lighting aftermarket industry in Taiwan," *Mathematics*, vol. 8, no. 7, art. 1187, 2020. doi: 10.3390/math8071187.

INDUSTRICS · ADVANCING INDUSTRY. SHAPING SYSTEMS.