EUROPEAN CONFERENCE

Conference Proceedings

XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

August 24–26, 2026
Prague, Czech Republic

PRIORITY AREAS OF SCIENTIFIC RESEARCH IN THE MODERN WORLD

Abstracts of XXXIV International Scientific and Practical Conference

Prague, Czech Republic
(August 24–26, 2026)

ORGANIZATIONAL RESILIENCE IN THE TIMES OF ALGORITHMIC DECISION SUPPORT

Prokopenko Serhii Olexandrovych
Ph.D., Senior Lecturer
Simon Kuznets Kharkiv National University of Economics,
Department of Marketing

Organizations across virtually every sector are embedding algorithmic decision-support systems into their operational core, from demand forecasting and inventory optimization to credit scoring, workforce scheduling, and dynamic pricing. This shift is frequently framed in terms of efficiency: algorithms process more data faster and with fewer errors than human decision-makers operating under time pressure.

What is far less frequently examined is what happens to an organization's capacity to adapt, absorb shocks, and recover from disruption once a substantial share of its operational decisions has been delegated, wholly or partially, to systems whose internal logic is opaque even to the managers who rely on them. This paper addresses that gap by applying resilience theory, and in particular the frameworks developed by Woods (2015) and by Walker et al. (2004), to the specific problem of organizational resilience under conditions of algorithmic decision support.

Resilience, as developed in socio-technical systems research, is not merely the ability to withstand disruption without changing, but the capacity to anticipate, absorb, adapt to, and, where necessary, transform in response to disturbance. Woods identifies several essential characteristics of resilient systems, including the capacity to recognize when established models of the operating environment no longer fit changing conditions, and the ability to reallocate resources and adjust priorities before a disruption becomes a crisis. Walker and colleagues, working within ecological and social-ecological systems theory, further distinguish resilience from adaptability and transformability: the former refers to a system's capacity to absorb disturbance while remaining within existing structures and functions; the latter refers to the capacity to create a fundamentally different system when existing structures become untenable. Applied to organizations, this distinction matters considerably, since many firms that appear resilient in the narrow sense of continuing operations without visible disruption may in fact be accumulating latent fragility that only becomes apparent when a sufficiently large shock exposes the limits of their adaptive capacity.

Algorithmic decision support introduces a distinctive form of this latent fragility. When operational decisions are governed by models trained on historical data, the resulting systems are structurally biased toward the assumption that future conditions will resemble the conditions under which the model was trained and validated. This is not a flaw specific to any particular algorithm but a general property of statistical learning (Quiñonero-Candela et al., 2009): models generalize from observed patterns, and their reliability necessarily degrades as operating conditions diverge from the training distribution. In stable environments, this property is largely inconsequential. Under conditions of genuine disruption, however, precisely the conditions in which organizational resilience matters most, algorithmic systems may continue to produce confident recommendations even as the assumptions underlying those recommendations quietly fail. Woods describes this pattern as brittleness: a system that performs well within an anticipated range of conditions but degrades sharply, rather than gracefully, once conditions move outside that range. The central argument of this paper is that algorithmic decision support systems, absent deliberate organizational design to counteract this tendency, systematically push organizations toward brittleness rather than resilience, even as they improve day-to-day operational efficiency.

This dynamic is compounded by a second effect that resilience scholars describe as the erosion of requisite variety (Weick & Sutcliffe, 2015), the principle that a system's capacity to respond to disturbance depends on maintaining a repertoire of responses at least as diverse as the range of disturbances it may face. Human decision-makers operating without heavy algorithmic mediation retain a wide, if inconsistent, repertoire of judgment strategies, informal knowledge, and improvisational capacity that can be drawn upon when standard operating procedures fail to fit an unanticipated situation. As algorithmic systems take over routine decisions, this repertoire does not disappear immediately, but it atrophies through disuse. Employees who previously exercised judgment over scheduling, pricing, or resource allocation may retain formal authority to override the algorithm, but the tacit skill required to exercise that authority well, calibrated through years of direct engagement with the underlying problem, degrades when decisions are habitually deferred to the system. The organization thus loses precisely the distributed, improvisational capacity that resilience theory identifies as central to adaptive response, at the moment when algorithmic confidence is least warranted.

A further complication concerns the visibility of failure. In traditional operational settings, a manager who makes a poor decision typically receives relatively direct and interpretable feedback, and the reasoning behind the decision remains available for retrospective examination and organizational learning. Algorithmic decision-support systems complicate this feedback loop in two ways. First, many such systems operate as effective black boxes even to the personnel who use them daily, meaning that when outcomes diverge from expectations, it is often unclear whether the fault lies in the model, the data, the way outputs were interpreted, or some interaction among these factors. Second, organizations frequently lack established protocols for systematically capturing and analyzing near-misses and small failures in algorithmic recommendations, in contrast to the relatively mature incident-reporting cultures that exist in domains such as aviation and healthcare, from which much of the empirical resilience literature originates. Without this feedback infrastructure, organizations risk repeating the same latent errors indefinitely, since the signals that would ordinarily prompt corrective adaptation are either invisible or are not routed to anyone with the authority and expertise to act on them.

The paper argues that preserving organizational resilience under conditions of extensive algorithmic decision support requires a set of deliberate structural interventions rather than reliance on the natural adaptive tendencies of the organization.

First, organizations should maintain what can be termed a manual override capacity, not merely as a formal policy permitting human intervention, but as an actively exercised capability, periodically tested and rehearsed, so that the judgment required to exercise it effectively does not atrophy through disuse. This is analogous to practices already established in highly automated technical domains, where operators are required to periodically operate systems manually, specifically to maintain the skill required to intervene when automation fails. Second, organizations should institutionalize systematic monitoring for distributional drift, the gradual divergence between current operating conditions and the conditions represented in the data on which decision-support models were trained, treating early signs of drift as a trigger for structured review rather than waiting for visible operational failure. Third, organizations should cultivate what resilience scholars term requisite diversity of perspective within the teams responsible for overseeing algorithmic systems, ensuring that model outputs are subject to genuine scrutiny from personnel whose expertise and incentives differ from those of the system's designers, rather than being rubber-stamped by teams whose primary incentive is to demonstrate the system's continued validity.

These interventions carry organizational costs that are worth acknowledging directly, since much of the appeal of algorithmic decision support lies precisely in the efficiency gains achieved by removing redundant human oversight from routine decisions. Maintaining manual override capacity, monitoring for drift, and preserving diverse oversight all require resources that a narrowly efficiency-focused evaluation of algorithmic adoption would tend to eliminate. This paper argues that this tension is not a design flaw to be optimized away but reflects a genuine trade-off between short-run operational efficiency and long-run adaptive capacity, a trade-off that resilience theory has long identified in other domains, and one that requires that organizations treat resilience-preserving redundancy as a deliberate investment rather than as waste to be minimized.

The implications extend beyond individual firms to sector-level and economy-wide considerations. As algorithmic decision-support systems become more standardized across an industry, whether through shared vendor platforms, common training data sources, or convergent competitive pressures toward similar architectures, the brittleness identified at the level of individual organizations risks becoming correlated across firms, such that a shock that exposes the limits of one organization's algorithmic assumptions may simultaneously expose the same limits across much of the sector. This mirrors concerns already well established in financial stability research regarding correlated risk models across institutions, and suggests that organizational resilience in the age of algorithmic decision support cannot be addressed purely at the level of individual firm strategy but also requires attention to the systemic homogenization of decision-support architectures across an industry.

These dynamics play out somewhat differently across functional domains, and a brief comparative sketch illustrates the breadth of the problem. In supply chain and inventory management, algorithmic forecasting systems have become the default mechanism for replenishment and allocation decisions across large retail and logistics operations. These systems perform well under conditions of demand stability, but the early phase of the COVID-19 pandemic offered a clear illustration of the brittleness this paper describes: forecasting models trained on years of relatively stable consumption patterns produced systematically misleading recommendations precisely when accurate forecasting mattered most, and organizations that had allowed manual override capacity and buffer-stock judgment to atrophy over the preceding years of algorithmic optimization were demonstrably slower to recover appropriate ordering behavior than organizations that had preserved parallel human forecasting capability as a matter of policy.

In human resources and workforce scheduling, algorithmic systems optimize shift assignments against historical demand patterns and labor cost constraints, but these same systems have limited capacity to anticipate discontinuous shifts in workforce availability, whether due to public health emergencies, sudden regulatory change, or localized disruption (Kim, 2025). Organizations relying exclusively on algorithmic scheduling may therefore be slower to restore adequate staffing once historical assumptions cease to hold, a delay attributable in part to the loss of the tacit scheduling expertise that had previously resided with human managers before being displaced by automated systems.

In credit risk and lending, algorithmic underwriting models are trained on historical repayment behavior under specific macroeconomic conditions, and resilience concerns here are compounded by the systemic homogenization risk discussed above, since many financial institutions rely on a comparatively small number of shared modeling approaches and third-party risk-scoring vendors, raising the possibility that a macroeconomic shock invalidating one institution's model assumptions may simultaneously invalidate the assumptions embedded across a substantial share of the sector, precisely the correlated-fragility mechanism that financial stability regulators have already identified as a systemic risk in other contexts.

These sectoral illustrations point toward a broader observation about the temporal structure of the resilience problem this paper identifies. The costs of preserving manual judgment capacity, rehearsing override procedures, and maintaining diverse oversight are incurred continuously, in the form of reduced short-run efficiency, while the benefits are realized only intermittently, at the moment of disruption, and are therefore easy to systematically undervalue in ordinary budgetary and performance-evaluation cycles that reward measurable short-run efficiency gains and rarely reward the avoidance of low-probability, high-impact disruptions that have not yet occurred. This temporal asymmetry helps explain why organizations that have suffered a clear resilience failure often subsequently overcorrect, investing heavily in redundancy immediately following a crisis before gradually eroding that redundancy again as the memory of the disruption fades and competitive pressure toward efficiency reasserts itself, a cyclical pattern that resilience scholars have observed in other high-reliability domains and that this paper suggests is likely to recur in the management of algorithmic decision support absent structural mechanisms that insulate resilience investment from short-term performance pressure.

One plausible structural mechanism is to formally separate the evaluation of algorithmic decision-support systems from the teams responsible for their deployment and ongoing operation, analogous to the separation between risk management and trading functions that has become standard practice in financial institutions following earlier episodes of correlated model failure. An independent review function, reporting through a distinct organizational channel and evaluated on criteria that explicitly include the maintenance of override capacity and diversity of oversight rather than solely on operational efficiency metrics, would be structurally better positioned to resist the gradual erosion of resilience-preserving redundancy that this paper identifies as a predictable organizational tendency. This is consistent with broader institutional theory, which suggests that organizational practices persist most reliably when they are embedded in formal structures and evaluation criteria rather than left to the discretionary judgment of actors whose incentives favor short-run efficiency.

Methodologically, this paper is conceptual rather than empirical, developing its argument by integrating resilience theory with existing organizational and information systems literature on algorithmic management. Future empirical work should test the specific mechanisms proposed here, particularly the hypothesized erosion of judgment capacity under sustained algorithmic mediation, through longitudinal study of organizations at varying stages of algorithmic adoption, and should examine whether the structural interventions proposed, manual override rehearsal, distributional drift monitoring, and diversity of oversight, measurably improve organizational performance under disruption relative to organizations that have adopted algorithmic decision support without corresponding resilience safeguards.

In conclusion, the widespread organizational adoption of algorithmic decision support represents a genuine operational advance under stable conditions, but resilience theory suggests that the same properties that generate efficiency gains in stable environments, reliance on historical pattern-matching, reduced deployment of distributed human judgment, and opaque internal reasoning, systematically undermine the adaptive capacity that organizations most need under disruption. Preserving organizational resilience in this environment is not automatic and does not follow from the technical sophistication of the algorithmic systems themselves. It requires deliberate managerial investment in maintaining human judgment capacity, monitoring the boundaries of model validity, and preserving genuine diversity of oversight, treating it as a necessary cost of long-run adaptive capacity rather than as an inefficiency to be eliminated. The organizations most likely to withstand the next major disruption to their operating environment are unlikely to be those with the most sophisticated algorithmic decision-support architecture as such, but rather those that have deliberately preserved the human capacity to recognize when that architecture has quietly stopped fitting the world it was built to describe.

References

1. Walker, B., Holling, C. S., Carpenter, S. R., & Kinzig, A. (2004). Resilience, adaptability and transformability in social–ecological systems. *Ecology and Society, 9*(2), 5.

2. Woods, D. D. (2015). Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering & System Safety, 141*, 5–9. https://doi.org/10.1016/j.ress.2015.03.018

3. Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., & Lawrence, N. D. (2009). *Dataset shift in machine learning*. MIT Press.

4. Woods, D. D. (2018). The theory of graceful extensibility: Basic rules that govern adaptive systems. *Environment Systems and Decisions, 38*(5), 433–457. https://doi.org/10.1007/s10669-018-9708-3

5. Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the unexpected: Sustained performance in a complex world* (3rd ed.). Wiley.

6. Gatrell, C., Muzio, D., Post, C. and Wickert, C. (2024), Here, There and Everywhere: On the Responsible Use of Artificial Intelligence (AI) in Management Research and the Peer-Review Process. J. Manage. Stud., 61: 739-751. https://doi.org/10.1111/joms.13045

7. Anthropic. Claude (Sonnet 5 version) [Large language model], 2026. https://claude.ai

8. Kim, S. (2025). Strategic human resource management in the era of algorithmic technologies: Key insights and future research agenda. *Human Resource Management*. https://doi.org/10.1002/hrm.22268

9. Uhl-Bien, M., & Arena, M. (2017). Complexity Leadership: Enabling People and Organizations for Adaptability. Organizational Dynamics, 46, 9-20. https://doi.org/10.1016/j.orgdyn.2016.12.001

10. Teece, D.J. (2007), Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. Strat. Mgmt. J., 28: 1319-1350. https://doi.org/10.1002/smj.640

Scientific publications

MATERIALS

The XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

Prague, Czech Republic
(August 24–26, 2026)