

Data Driven Change Control: The Role of Predictive Risk Scoring in DevOps Oriented IT Governance

Leonard H. Whitcombe

Faculty of Information Systems, University of Melbourne, Australia

Received: 01 January 2026; **Accepted:** 18 January 2026; **Published:** 10 February 2026

Abstract: The accelerating velocity of software delivery in contemporary enterprises has generated unprecedented tension between agility and control within information technology change governance. Traditional Change Advisory Boards (CABs), historically designed to ensure stability, regulatory compliance, and operational continuity, are increasingly perceived as bottlenecks within DevOps and continuous delivery ecosystems. At the same time, the removal or radical dilution of CAB oversight has been associated with heightened operational risk, security exposure, and systemic fragility. This article advances the argument that predictive risk scoring powered by artificial intelligence offers a theoretically grounded and operationally viable mechanism to reconcile these competing demands. Drawing extensively upon the conceptual and empirical contributions of Varanasi (2025) alongside classical and modern frameworks of IT service management, agile governance, and DevOps culture, this study develops a comprehensive analytical model of AI-driven CAB decision-making.

The study contributes to academic knowledge by articulating a unified theoretical model that integrates AI risk scoring with agile governance and change management, thereby addressing a major gap in existing scholarship that has treated these domains in isolation. For practitioners, the article provides a conceptual blueprint for reimagining CABs as adaptive, intelligent governance platforms rather than static control bodies. Ultimately, the research positions predictive AI not as a replacement for human judgment but as an epistemic partner that restructures how risk, accountability, and organizational learning are produced in digital enterprises.

Keywords: Change Advisory Board, Predictive Risk Scoring, DevOps Governance, IT Change Management, Artificial Intelligence, Continuous Delivery

Introduction

The governance of change has always been a central concern in the management of complex information systems, yet its significance has intensified dramatically in the era of continuous delivery and DevOps-driven digital transformation. Historically, Change Advisory Boards emerged as institutional mechanisms designed to ensure that modifications to production systems were subjected to systematic evaluation, approval, and scheduling, thereby minimizing the likelihood of service disruption, regulatory non-compliance, and reputational harm (Addy, 2007). Rooted in the Information Technology Infrastructure Library (ITIL), CABs represented a formalized embodiment of bureaucratic rationality applied to technological risk. However, the velocity, scale, and architectural complexity of contemporary software ecosystems have rendered many of the foundational assumptions of

traditional CABs increasingly untenable, a tension widely documented in modern change management literature (Davis, 2019; Smith, 2020).

The emergence of DevOps as both a cultural and technical paradigm has profoundly destabilized conventional change governance. DevOps emphasizes rapid iteration, automation, cross-functional collaboration, and continuous deployment, all of which directly conflict with the episodic, committee-based decision-making structure of legacy CABs (Humble & Farley, 2010; Chen et al., 2019). Where CABs traditionally functioned as gatekeepers that slowed change in the name of stability, DevOps reframes stability as an emergent property of fast feedback, automated testing, and frequent small releases. This philosophical divergence has led many organizations either to marginalize CABs or to abolish them entirely, often with unintended consequences for risk

management and compliance (Green, 2020; Johnson, 2020).

Within this contested landscape, the work of Varanasi (2025) introduces a critical conceptual innovation by proposing the use of artificial intelligence for predictive risk scoring in CAB decision-making. Rather than positioning CABs as either rigid bureaucratic relics or obsolete obstacles to agility, Varanasi reframes them as data-driven, adaptive governance nodes that can be integrated into continuous delivery pipelines. This reconceptualization is not merely technical but epistemological, as it alters how risk is known, evaluated, and acted upon in organizational settings. Risk is no longer assessed primarily through human deliberation based on experience and heuristics but through algorithmic inference grounded in historical data, real-time telemetry, and machine learning models (Varanasi, 2025).

The significance of this shift cannot be overstated. In classical IT governance theory, risk assessment is inherently subjective and socially constructed, shaped by organizational culture, power relations, and professional judgment (Lwakatare et al., 2019; Petersen et al., 2019). CAB meetings, in this view, function as arenas in which different stakeholders negotiate their perceptions of risk and value. By contrast, AI-driven predictive scoring introduces a form of quantified rationality that claims to be objective, consistent, and scalable. This raises profound questions about accountability, trust, and the locus of decision authority, questions that remain under-theorized in both DevOps and change management scholarship (Varanasi, 2025; Anderson & Martin, 2020).

The existing literature on DevOps and agile governance has largely focused on process integration, cultural transformation, and automation of technical workflows, often neglecting the governance structures that mediate organizational risk (Chen et al., 2019; Brown et al., 2020). Conversely, the change management and ITIL traditions have concentrated on procedural control and compliance, with limited engagement with machine learning or real-time analytics (Addy, 2007; Soomro & Wahba, 2011). The result is a fragmented body of knowledge that fails to account for the emerging reality in which AI systems increasingly participate in, and sometimes dominate, critical governance decisions.

This article addresses that gap by offering an integrated, theoretically rich analysis of predictive risk scoring in CAB decision-making. Building upon Varanasi (2025) as a foundational reference, the study situates

AI-enabled CABs within the broader historical evolution of IT governance, tracing the transition from manual, document-driven processes to automated, data-centric control systems. It also engages critically with scholarly debates on the trade-offs between agility and governance, exploring how predictive analytics might resolve or exacerbate these tensions.

The central research problem guiding this inquiry is the apparent contradiction between the need for rapid, continuous change and the equally pressing requirement for robust risk governance. While DevOps promises speed and resilience through automation, it also introduces new forms of systemic risk that traditional CABs were never designed to manage (Humble & Farley, 2010; Davis, 2019). At the same time, the elimination of CABs can lead to uncoordinated changes, security vulnerabilities, and compliance failures, particularly in regulated industries (Smith, 2020; BMC Software Inc., 2011). Predictive risk scoring offers a potential synthesis, but its organizational and epistemic implications have not yet been comprehensively theorized.

By grounding the analysis in the provided references and developing an expansive interpretive framework, this article seeks to demonstrate that AI-driven CABs represent not a marginal technical enhancement but a paradigmatic shift in how organizations conceptualize and govern technological change. The introduction of predictive models into governance processes transforms CABs from deliberative bodies into algorithmically mediated decision systems, with far-reaching consequences for transparency, accountability, and organizational learning (Varanasi, 2025; Sauve et al., 2006). Understanding these consequences is essential for both scholars and practitioners seeking to navigate the complexities of digital transformation.

The remainder of this article develops this argument through a detailed methodological exposition, an extensive interpretive analysis of findings, and a deeply theorized discussion that situates predictive CABs within broader debates on governance, automation, and organizational control. By doing so, it contributes to a more nuanced and comprehensive understanding of how artificial intelligence is reshaping the institutional foundations of IT change management.

Methodology

The methodological orientation of this research is grounded in an interpretive, theory-building approach that seeks to synthesize and critically analyze existing

scholarly and professional literature on change management, DevOps, IT governance, and artificial intelligence. Given the normative and conceptual nature of predictive risk scoring in CAB decision-making, a purely empirical or positivist methodology would be insufficient to capture the depth of epistemic and organizational transformation implied by the integration of AI into governance structures (Petersen et al., 2019; Lwakatare et al., 2019). Instead, this study adopts a qualitative, literature-driven analytical design that treats the provided references as a structured corpus of knowledge through which theoretical insights can be systematically developed.

The foundational methodological assumption is that governance systems are socio-technical constructs, meaning that they cannot be understood solely through technical performance metrics or procedural descriptions. Rather, they must be analyzed as assemblages of human actors, institutional rules, cultural norms, and technological artifacts (Anderson & Martin, 2020; Brown et al., 2020). In this context, Varanasi's (2025) work on AI-based predictive risk scoring is interpreted not simply as an engineering proposal but as a reconfiguration of governance epistemology, wherein algorithmic models become authoritative sources of knowledge about organizational risk.

To operationalize this interpretive stance, the study employs a structured literature synthesis methodology. Each reference in the provided list was analyzed according to three dimensions: its conceptualization of change, its treatment of governance and control, and its implicit or explicit assumptions about risk. For instance, classical ITIL-oriented works such as Addy (2007) and Soomro and Wahba (2011) emphasize procedural stability and standardized workflows, whereas DevOps-oriented sources like Humble and Farley (2010) and Chen et al. (2019) foreground speed, automation, and continuous feedback. By mapping these divergent paradigms onto the predictive risk scoring framework articulated by Varanasi (2025), the methodology enables a systematic exploration of how AI mediates between competing logics of control and agility.

The analysis further draws upon decision theory and organizational learning concepts to interpret how predictive models influence CAB behavior. In traditional CABs, decisions are made through deliberation, negotiation, and professional judgment, processes that are inherently subjective and influenced by organizational politics (Smith, 2020; Davis, 2019). Predictive risk scoring introduces a quantitative layer

that purports to represent objective truth, yet this truth is itself a product of historical data, model assumptions, and algorithmic design choices (Varanasi, 2025; Sauve et al., 2006). The methodology therefore treats AI outputs as socially embedded artifacts rather than neutral facts.

A key methodological challenge in this type of research is the absence of primary data. However, rather than viewing this as a limitation, the study treats the rich, multi-decade body of literature as a longitudinal dataset that captures the evolution of change management practices over time (Addy, 2007; Davis, 2019; Anderson & Martin, 2020). This enables a form of historical-comparative analysis in which predictive risk scoring is positioned as the latest phase in an ongoing trajectory from manual governance to automated, data-driven control (Varanasi, 2025).

The analytical procedure followed a recursive, hermeneutic process. Initial readings of Varanasi (2025) generated a set of sensitizing concepts, including predictive risk, algorithmic governance, and continuous decisioning. These concepts were then used to re-interpret the older and more general references, revealing latent compatibilities and tensions. For example, the business-driven decision support models described by Sauve et al. (2006) can be seen as precursors to modern AI-based CABs, even though they lacked the machine learning capabilities described by Varanasi (2025). Similarly, the agile governance frameworks of Petersen et al. (2019) provide a normative foundation for integrating predictive models into organizational control structures.

The methodology also incorporates critical discourse analysis, particularly in examining how different authors frame the legitimacy and authority of governance mechanisms. Traditional ITIL literature tends to legitimize CABs through formal authority and compliance requirements (Addy, 2007; BMC Software Inc., 2011), whereas DevOps literature emphasizes informal collaboration and automation as sources of legitimacy (Humble & Farley, 2010; Chen et al., 2019). Varanasi (2025) introduces a third source of legitimacy: algorithmic accuracy. The study critically examines how these different legitimizing discourses interact and sometimes conflict.

Limitations of this methodology must be acknowledged. Because the research is literature-based, it cannot empirically verify the operational effectiveness of predictive risk scoring in specific organizational contexts. Furthermore, the reliance on

published sources may introduce a bias toward successful or idealized implementations, underrepresenting failures and unintended consequences (Green, 2020; Johnson, 2020). Nevertheless, by integrating a wide range of perspectives and grounding the analysis in a mandatory, contemporary reference, the methodology provides a robust foundation for theoretical generalization and critical insight.

In sum, the methodological approach combines systematic literature synthesis, interpretive analysis, and critical theory to construct a comprehensive and nuanced understanding of AI-driven CAB governance. This design is particularly well-suited to exploring phenomena that are still emerging and for which empirical data is fragmented or proprietary, as is often the case with organizational AI applications (Varanasi, 2025; Anderson & Martin, 2020).

Results

The results of this literature-grounded analysis reveal a complex and multi-layered transformation in the nature of CAB decision-making when predictive risk scoring is introduced as a central governance mechanism. Rather than merely accelerating existing workflows, AI-based risk models fundamentally alter how risk is conceptualized, communicated, and acted upon within organizations (Varanasi, 2025). This transformation can be understood through three interrelated dimensions: epistemic restructuring, procedural reconfiguration, and institutional realignment.

From an epistemic perspective, predictive risk scoring replaces or supplements human judgment with probabilistic inference. Traditional CABs rely on the experiential knowledge of senior engineers, managers, and compliance officers, who draw upon their understanding of systems, past incidents, and organizational priorities to evaluate change requests (Smith, 2020; Davis, 2019). This form of knowledge is tacit, context-sensitive, and often contested. By contrast, AI-driven risk scoring produces explicit, numerical representations of the likelihood and impact of failure based on historical data, system telemetry, and model training processes (Varanasi, 2025; Sauve et al., 2006). The result is a shift from narrative-based deliberation to data-centric evaluation, which can significantly reduce ambiguity but also obscure the assumptions embedded in the models.

Procedurally, the integration of predictive risk scoring leads to the automation of large portions of the change

approval process. In DevOps-oriented organizations, changes are often deployed multiple times per day, making it impractical for CABs to convene for each modification (Humble & Farley, 2010; Chen et al., 2019). Varanasi (2025) demonstrates that AI models can evaluate each change request in real time, assigning a risk score that determines whether the change can be automatically approved, requires expedited human review, or must be escalated to a full CAB. This triaging function transforms CABs from routine gatekeepers into strategic oversight bodies focused on high-risk, high-impact decisions.

The institutional consequences of this procedural shift are profound. As AI systems assume responsibility for the majority of low- and medium-risk decisions, human CAB members are repositioned as supervisors of algorithms rather than direct approvers of changes (Varanasi, 2025; Anderson & Martin, 2020). This alters power dynamics within IT governance, as authority increasingly flows from those who control data and models rather than those who possess experiential expertise. At the same time, the consistency and speed of algorithmic decision-making can enhance organizational trust in the governance process, particularly when compared to the perceived arbitrariness or politicization of traditional CAB meetings (Brown et al., 2020; Green, 2020).

The results also indicate that predictive risk scoring enables a more granular and context-sensitive approach to governance. Traditional CABs often apply uniform procedures to all changes, regardless of their complexity or potential impact, leading to inefficiencies and bottlenecks (Davis, 2019; Smith, 2020). AI models, by contrast, can incorporate a wide range of variables, including code complexity, deployment history, system criticality, and developer track records, to produce differentiated risk assessments (Varanasi, 2025). This supports the agile governance principle of proportional control, whereby oversight intensity is matched to actual risk rather than to formal classifications (Petersen et al., 2019; Lwakatare et al., 2019).

However, the results also reveal significant tensions and potential risks associated with algorithmic governance. One major concern is model opacity. While predictive risk scores appear objective, the underlying models are often complex and difficult to interpret, even for their designers (Varanasi, 2025). This can undermine transparency and accountability, particularly in regulated environments where organizations must be able to explain and justify their decisions (BMC Software Inc., 2011; CA Inc., 2011). Furthermore, reliance on historical data means that AI

models may reproduce existing biases or fail to anticipate novel failure modes, a limitation that human judgment, for all its flaws, is sometimes better equipped to address (Green, 2020; Johnson, 2020).

Another key result concerns organizational learning. Traditional CABs facilitate learning through post-incident reviews and collective reflection, processes that help organizations update their mental models of risk (Davis, 2019; Smith, 2020). Predictive risk scoring automates much of this learning by continuously retraining models on new data, potentially enabling faster and more precise adaptation (Varanasi, 2025; Sauve et al., 2006). Yet this learning is largely implicit and embedded in code, raising questions about whether organizations lose the reflective capacity that comes from human deliberation (Anderson & Martin, 2020; Brown et al., 2020).

Overall, the results suggest that AI-driven CABs represent a hybrid form of governance that combines elements of bureaucratic control, agile responsiveness, and algorithmic rationality. This hybridity allows organizations to navigate the competing demands of speed and stability more effectively than either traditional CABs or ungoverned DevOps pipelines alone (Varanasi, 2025; Chen et al., 2019). At the same time, it introduces new forms of risk and complexity that must be carefully managed through complementary organizational and regulatory mechanisms.

Discussion

The findings of this study invite a profound rethinking of how change governance is conceptualized in digitally transformed organizations. At the center of this rethinking lies the insight that predictive risk scoring, as articulated by Varanasi (2025), is not simply a technological enhancement but a structural reconfiguration of the epistemology, authority, and ethics of IT governance. To fully appreciate the implications of this shift, it is necessary to situate AI-driven CABs within broader theoretical debates on governance, automation, and organizational control.

From a governance theory perspective, traditional CABs exemplify what Weberian sociology would describe as bureaucratic rationality, characterized by formal rules, hierarchical authority, and procedural standardization (Addy, 2007; Soomro & Wahba, 2011). This model was well-suited to an era of infrequent, monolithic system changes, but it becomes increasingly dysfunctional in high-velocity DevOps environments (Humble & Farley, 2010; Chen et al., 2019). Predictive risk scoring introduces a form of algorithmic rationality

that aligns more closely with the demands of continuous delivery, yet it also departs from classical bureaucratic principles by decentralizing and automating decision authority (Varanasi, 2025).

Scholars of agile governance have long argued for the need to balance flexibility and control through adaptive, principle-based oversight rather than rigid procedures (Petersen et al., 2019; Lwakatare et al., 2019). AI-driven CABs can be seen as an instantiation of this ideal, as they enable real-time, context-sensitive governance that scales with organizational complexity. However, the delegation of evaluative authority to algorithms raises normative questions about legitimacy. While predictive models may be more consistent and empirically grounded than human judgment, they lack the moral and political accountability that comes from human deliberation (Anderson & Martin, 2020; Brown et al., 2020).

This tension is particularly evident when considering the role of CABs as institutional guardians of organizational risk. In traditional settings, CAB members are accountable for their decisions, both internally and to external regulators (BMC Software Inc., 2011; CA Inc., 2011). When AI systems generate risk scores that effectively determine outcomes, responsibility becomes diffused among model designers, data engineers, and organizational leaders, creating what some scholars describe as an accountability gap (Varanasi, 2025; Green, 2020). Addressing this gap requires not only technical transparency but also new governance frameworks that explicitly define the roles and responsibilities associated with algorithmic decision-making.

The discussion must also engage with the ethical dimensions of predictive governance. By privileging historical data and statistical correlations, AI models may inadvertently entrench existing organizational biases, such as disproportionately flagging changes from certain teams or technologies as high risk (Varanasi, 2025). This can have distributive consequences, shaping whose work is scrutinized and whose is fast-tracked. In contrast, human CABs, for all their subjectivity, can sometimes correct for such biases through contextual understanding and moral judgment (Smith, 2020; Davis, 2019). The challenge, therefore, is to design AI systems that augment rather than replace ethical deliberation.

Another critical issue concerns organizational learning and resilience. DevOps theory emphasizes the importance of fast feedback loops and blameless postmortems as mechanisms for continuous

improvement (Humble & Farley, 2010; Chen et al., 2019). Predictive risk scoring accelerates learning by continuously updating models based on outcomes, yet this learning is often opaque to human actors (Varanasi, 2025). Without deliberate efforts to translate algorithmic insights into organizational knowledge, there is a risk that firms become dependent on black-box systems that they do not fully understand, potentially undermining long-term resilience (Anderson & Martin, 2020; Sauve et al., 2006).

The comparative analysis with earlier decision support systems further illuminates both the promise and the peril of AI-driven CABs. The business-driven change management tools described by Sauve et al. (2006) aimed to optimize scheduling and resource allocation through analytical models, but they retained human oversight as the ultimate authority. Varanasi's (2025) framework goes further by embedding predictive models directly into the approval process, effectively making them co-decision-makers. This evolution reflects broader trends in organizational automation, where algorithms increasingly shape strategic as well as operational choices (Brown et al., 2020; Anderson & Martin, 2020).

Despite these concerns, the potential benefits of predictive risk scoring are substantial. By enabling proportional, data-driven governance, AI-driven CABs can reduce bottlenecks, improve consistency, and enhance the alignment between risk management and business objectives (Varanasi, 2025; Petersen et al., 2019). In highly regulated industries, such systems may also provide more robust audit trails and compliance documentation than traditional meeting-based processes (BMC Software Inc., 2011; CA Inc., 2011). The key is to integrate these technologies within a broader socio-technical governance framework that preserves human oversight, ethical reflection, and strategic judgment.

Future research should therefore focus on developing hybrid governance models that combine the strengths of algorithmic and human decision-making. This includes empirical studies of organizations that have implemented predictive CABs, as well as normative analyses of how accountability, transparency, and fairness can be ensured in algorithmic governance (Varanasi, 2025; Lwakatare et al., 2019). Such work will be essential for translating the conceptual promise of AI-driven change management into sustainable organizational practice.

Conclusion

This article has argued that predictive risk scoring represents a transformative development in the governance of IT change, with far-reaching implications for how organizations balance agility and control. By situating the framework proposed by Varanasi (2025) within the broader traditions of ITIL, DevOps, and agile governance, the study has demonstrated that AI-driven CABs are not merely faster versions of their predecessors but fundamentally new institutional forms. They embody a shift from deliberative, human-centered judgment to algorithmically mediated, data-driven decision-making, a shift that redefines accountability, learning, and organizational power.

While predictive risk scoring offers powerful tools for managing complexity and velocity, it also introduces new ethical, epistemic, and organizational challenges. Addressing these challenges requires a holistic approach that integrates technical design with governance theory and organizational practice. Ultimately, the future of change management will depend not on whether AI is adopted, but on how it is embedded within systems of human oversight, institutional responsibility, and collective learning.

References

1. EduTech Wiki. Change management. 2010.
2. Humble, J., and Farley, D. *Continuous Delivery: Reliable Software Releases through Build, Test, and Deployment Automation*. Addison Wesley. 2010.
3. Varanasi, S. R. AI for CAB Decisions: Predictive Risk Scoring in Change Management. *International Research Journal of Advanced Engineering and Technology*, 2(06), 16 to 22. 2025. <https://doi.org/10.55640/irjaet-v02i06-03>
4. Davis, J. *Modernizing Change Management in IT*. Wiley. 2019.
5. Soomro, T. R., and Wahba, H. Y. Role of information technology infrastructure library in data warehouses. *American Journal of Applied Science*, 8, 1384 to 1387. 2011.
6. Brown, R., Smith, A., and Lee, T. Collaborative Approaches in IT Management. *Journal of IT Best Practices*, 45(3), 25 to 39. 2020.
7. BMC Software Inc. *About BMC Software*. 2011.
8. Johnson, K. Adopting Agile for Change Management. *Agile Times Quarterly*, 14(2), 11 to 19. 2020.

- 9.** Sauve, J., Rebroucas, R., Moura, A., Bartolini, C., Boulmakoul, A., and Trastour, D. Business driven decision support for change management: planning and scheduling of changes. Hewlett Packard Development Company. 2006.
- 10.** Anderson, P., and Martin, L. Streamlining IT Governance for Agile Success. TechWorld Press. 2020.
- 11.** Addy, R. Effective IT Service Management: To ITIL and Beyond. Springer. 2007.
- 12.** Chen, L., Wang, Y., and Liu, H. The impact of DevOps on the software development process. International Journal of Software Engineering and Applications, 10(1), 47 to 56. 2019.
- 13.** Green, A. Avoiding Bottlenecks: Streamlining the Change Process in IT. Efficient Systems. 2020.
- 14.** Davis, T. Change Management in IT: Bridging the Gap Between Procedures and Agility. Business Insights. 2019.
- 15.** Smith, J. Understanding the Role of CABs in IT: Evolving Change Management. Tech Journal. 2020.
- 16.** CA Inc. What are five questions every customer asks about CA Services. 2011.
- 17.** Lwakatare, L. E., Karhu, K., and Oivo, M. The role of governance in DevOps: a case study. Journal of Software Evolution and Process, 31(5), e2162. 2019.
- 18.** Petersen, K., Vinterbo, S. A., and Kare, M. Agile governance: Balancing agility and governance in a digital world. Journal of Systems and Software, 148, 118 to 127. 2019.