HARVARD BUSINESS REVIEW

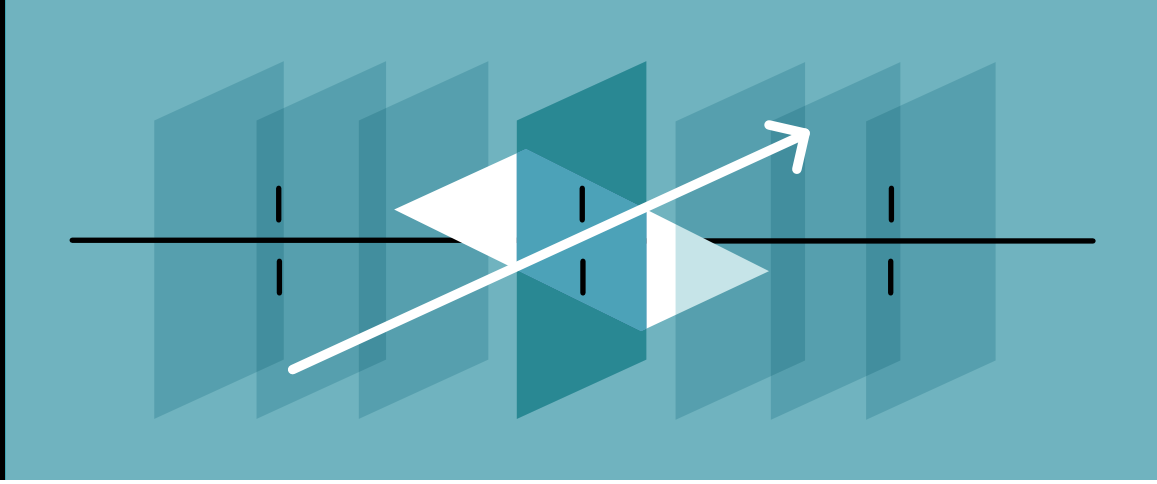

The Resilient Enterprise:

Using AI to Connect Governance, Risk, and Compliance

HBR Analytic Services **BRIEFING PAPER**

Sponsored by

Optro

Headline

Today, risk surfaces have grown far beyond the capacity of traditional governance, risk, and compliance (GRC) teams. Risks, from supply chain disruptions to cyber threats and beyond, evolve continuously, while the systems designed to oversee them remain static, manual, and episodic. The result is a widening gap between organizational accountability and what teams can realistically manage. Let's be clear — this gap is not caused by a lack of expertise. In fact, GRC teams are often among the most experienced, highly trained professionals in the business. Yet outdated processes trap them in repetitive work that does not scale and does not improve outcomes. AI gives teams a chance to rewrite this story.

AI is changing how work gets done inside organizations, not by replacing people, but by forcing a long-overdue reckoning with scale. The real opportunity of AI in GRC is not faster reports or more dashboards; it's removing the volume of rote work that forces risk teams into reactive, defensive postures. Tasks like reading and comparing frameworks, mapping policy changes, reconciling evidence, and tracing downstream impact can now be automated en masse. Lifting these burdens allows GRC professionals to return to where they drive the most value: judgment, prioritization, and accountability.

For decades, the lines of defense model has helped organizations clarify ownership of risk. But, in practice, it reinforced silos. Risk identification, control design, testing, and oversight became distributed across functions, systems, and tools that rarely share context. As risks accelerated and became more interconnected, these boundaries hardened into friction. Consequently, GRC functions became known less for enabling better decisions and more for stalling them.

However, a better way to think about this is through the lens of AI-enabled connected risk. This can fundamentally change how the lines of defense operate. Instead of acting as sequential checkpoints or "offices of no," risk and compliance functions can operate as continuous business partners within one platform. They can engage earlier, with richer context, and a clearer understanding of tradeoffs. This is how GRC teams level up and become strategic advisors to their organizations.

Risk management evolves from a tax paid after the fact into crucial intelligence that informs growth, expansion, and innovation. And more importantly, this evolution doesn't remove humans from the loop. While AI can surface discrepancies, propose changes, and highlight gaps, humans must still decide what is material, what is acceptable, and what actions to take. The most effective organizations will deliberately design this partnership, using automation to scale insight while preserving human control.

Anton Dam
Senior Vice President, Product & AI
Optro

When risk is understood in context, aligned with business objectives and across functions, and addressed continuously, it stops acting as a brake on progress. Instead, it becomes a stabilizing force that allows organizations to move faster and with confidence. That's how we transform risk into opportunity. Optro sponsored this research in association with Harvard Business Review Analytic Services to explore how leading organizations are navigating this transition. The findings highlight a clear path forward: harness the power of connected risk to reposition GRC as a business enabler, powered by AI to amplify human intelligence.

The Resilient Enterprise:

Using AI to Connect Governance, Risk, and Compliance

Artificial intelligence (AI) is reshaping governance, risk, and compliance (GRC) from a largely reactive oversight function into a more proactive and increasingly predictive strategic capability. Once viewed primarily as a mechanism for enforcing policy and regulatory adherence, GRC is now increasingly being seen as a source of foresight, resilience, and decision support, helping organizations anticipate disruption, manage uncertainty, and sustain trust in a volatile global environment.

BUT AI'S INFLUENCE on GRC is still in its formative stages. "There are two sides to the AI coin," says Mark Beasley, the Alan T. Dickson Distinguished Professor of Accounting and director of the Enterprise Risk Management Initiative at North Carolina State University. "There's the 'Wow, I can't get there fast enough' from a strategic advantage perspective, and at the same time, there are a lot of risks I need to be thinking about as I go down that journey."

As organizations adopt AI, automation, and data-driven decision making at scale, the consequences of unmanaged risk are rising quickly. Cyber threats, supply chain disruption, regulatory change, geopolitical instability, workforce transformation, and AI governance itself are no longer isolated concerns. They intersect and amplify one another, creating a risk landscape that is broader, faster-moving, and more interconnected than ever before.

This tension between speed and control, innovation and oversight, helps explain why GRC has reached an inflection point. As businesses confront this widening spectrum of interconnected risks, GRC has moved decisively to the center of the enterprise agenda. Many executives now see it not as a back-office control function but as a stabilizing force that restores clarity, coherence, and confidence amid uncertainty.

The most significant structural shift underway in GRC is the move toward "connected risk," an operating model that integrates data, people, and processes across traditionally siloed

HIGHLIGHTS

The most significant structural shift underway in governance, risk, and compliance (GRC) is the move toward connected risk—an operating model that **integrates data, people, and processes across traditionally siloed functions.**

This form of artificial risk intelligence—continuous, contextual, and enterprise-wide—allows GRC teams to **move beyond hindsight toward foresight.**

Many executives now see it not as a back-office control function, but as a **stabilizing force that restores clarity, coherence, and confidence amid uncertainty.**

> **One of the biggest risks companies face right now is thinking about AI only as a technology initiative rather than as a governance and risk issue.**
>
> Mark Beasley, Alan T. Dickson Distinguished Professor of Accounting and director of the Enterprise Risk Management Initiative, North Carolina State University

functions. "We decided that it made sense to bring the front end of risk identification all the way through assurance into one area, because that breaks down so many silos," says Erin Banet, chief audit and risk officer at Humana, a Louisville, Ky.-based health insurance company. The results, she says, are a shared taxonomy, a common risk language, and end-to-end traceability from inherent risk through residual risk and assurance throughout the organization.

Humana isn't alone. Other organizations are no longer asking whether risk functions should be connected but rather how quickly they can make that connection real. However, there's a caveat. Instead of diminishing the role of human expertise, AI is magnifying it by elevating audit, risk, compliance, and information security professionals into strategic advisers who guide organizational resilience.

This paper examines how organizations are navigating this transformation and where AI is already reshaping the way governance, risk, and compliance functions operate. It explores how connected risk models are replacing siloed functions, how continuous and predictive assurance is changing the cadence of GRC, and how disciplined AI governance is becoming essential as automation accelerates. The paper also examines the data foundations, operating models, and human capabilities required to apply AI responsibly and at scale.

The Broader Risk Imperative

AI's growing influence on GRC reflects a broader tension facing organizations today, namely the promise of speed, scale, and insight on one hand and the need for discipline, judgment, and governance on the other. Across industries, many firms are experimenting with AI to identify, prioritize, and mitigate risk, drawing from sources such as audit reports, regulatory guidance, cybersecurity briefings, economic forecasts, and third-party risk data. NC State's Beasley points to early use cases in which risk and audit teams used AI to analyze global reports and external risk indices, distilling fragmented data into coherent themes for executive discussion. "AI can connect dots across data sets and elevate your game from information to insight," he says.

Yet even as these capabilities mature, Beasley emphasizes that AI must remain a decision support tool, not a decision maker. "One of the biggest risks companies face right now is thinking about AI only as a technology initiative rather than as a governance and risk issue," he says. "The organizations that get into trouble are the ones that move fast without clearly understanding where accountability, oversight, and decision rights actually sit." That principle—AI as an accelerator of judgment rather than a substitute for it—runs throughout the emerging GRC landscape and defines the next phase of GRC transformation. The situation is urgent. Macroeconomic volatility, including inflationary pressure and currency risk, now intersects with fragile global supply chains, energy constraints, and escalating regulatory complexity. Meanwhile, the rapid adoption of cloud platforms and remote work habits has expanded the enterprise attack surface, increasing the likelihood and potential impact of cybersecurity breaches.

For senior executives, the reputational and economic consequences of these risks are no longer theoretical. A single disruption, whether a cyber incident, third-party failure, or regulatory lapse, can cascade across operations, customers, and markets. Yet many organizations continue to approach risk in narrow, function-specific ways, addressing threats in isolation rather than as part of a dynamic system. The result is a loss of enterprise-wide perspective that slows decision making precisely when speed and coordination matter most.

This gap between the complexity of modern risk and the structure of traditional oversight has created an opportunity, and an imperative, for GRC to evolve. Rather than focusing solely on historical compliance and control testing, organizations are increasingly seeking broader risk intelligence—insight that provides context, highlights interdependencies, and helps GRC professionals anticipate where disruption is most likely to emerge. AI is accelerating this shift, enabling organizations to move from static risk registers to having continuously updated views of enterprise exposure.

Modern GRC platforms centralize audit, risk, and compliance data into a common environment, enabling GRC professionals to standardize taxonomies, automate workflows, and maintain consistent documentation.

The Rise of Connected Risk Ecosystems

Rather than treating audit, compliance, risk management, information security, third-party risk, and business continuity as discrete activities, forward-looking GRC professionals are aligning them into a continuous, enterprise-wide view.

At Humana, GRC functions that were formerly managed by separate leadership teams are now explicitly combined. Banet oversees audit, risk, business continuity, third-party risk, and model governance as a single organizational function, with aligned processes and common reporting structures. However, she says connected risk goes far beyond reorganizing reporting lines or consolidating functions. It requires consistent data models, aligned methodologies, and shared visibility into how risks and controls interact across the enterprise. AI will become a key enabler of that visibility and be part of her team's future strategy, helping them analyze risk and control inventories, benchmark against external data, and surface interdependencies that might otherwise go unnoticed.

A similar philosophy is evident at Navan, a Palo Alto, Calif.-based travel, payments, and expense management platform, though applied in a very different context. Erin Dempsey Heuwetter, head of audit, risk, and compliance at Navan, is building a connected risk ecosystem from the ground up. "We're trying to develop a connected risk program from the start instead of bolting it on later," she confirms. "By establishing a single source that links risks, controls, and policies across the enterprise, we're making decision making more efficient and risk-driven because we can easily see the full picture."

For example, Heuwetter and her team have deployed AI tools to centralize and analyze data, automate workflows, and standardize reporting. GRC professionals rely on these intelligent agents to perform risk assessments, recommend controls, and generate audit work papers—compressing what once took days of manual synthesis into minutes. This architecture allows risk owners across business units to participate directly in risk assessments while maintaining consistent oversight, giving teams direct access to risk tools and embedding risk ownership closer to day-to-day decision making. "The goal isn't to centralize control; it's to distribute risk intelligence," Heuwetter notes. "When business owners can access the same risk data we do, they make better decisions without waiting for us. That's the unlock—GRC as an enabler, not a bottleneck."

Underpinning this shift is a new generation of GRC technology designed to support integration, automation, and decisions at scale. Modern GRC platforms centralize audit, risk, and compliance data into a common environment, enabling GRC professionals to standardize taxonomies, automate workflows, and maintain consistent documentation.

These systems support capabilities such as automated control testing, continuous issue management, centralized policy and regulatory content, and integrated reporting across GRC. When augmented with AI, they also enable advanced analytics, pattern recognition, and contextual insight, helping GRC teams like the one at Humana prioritize risk, focus resources, and surface emerging issues quickly. "With the help of our outsourced partners, we use AI for risk and control builds," Banet explains. "This helps us analyze what risks we should be considering, bringing in industry perspectives and benchmarking from other industries." In this sense, technology becomes an enabler of connected intelligence, not an end in itself, she adds.

From Reactive to Proactive with AI

AI is also reshaping the cadence of GRC itself, shifting it from periodic, point-in-time reviews to continuous and increasingly predictive assurance. Advanced analytics, machine learning, and automation enable organizations to discover anomalies, monitor controls in near real time, and detect early risk signals before issues escalate into incidents.

For example, at Navan, Heuwetter's team is using AI to anticipate compliance issues and operational bottlenecks across rapidly evolving business processes. "We can start to see risk signals early, which helps us redirect resources before an issue becomes an incident," she says. "This [foresight]

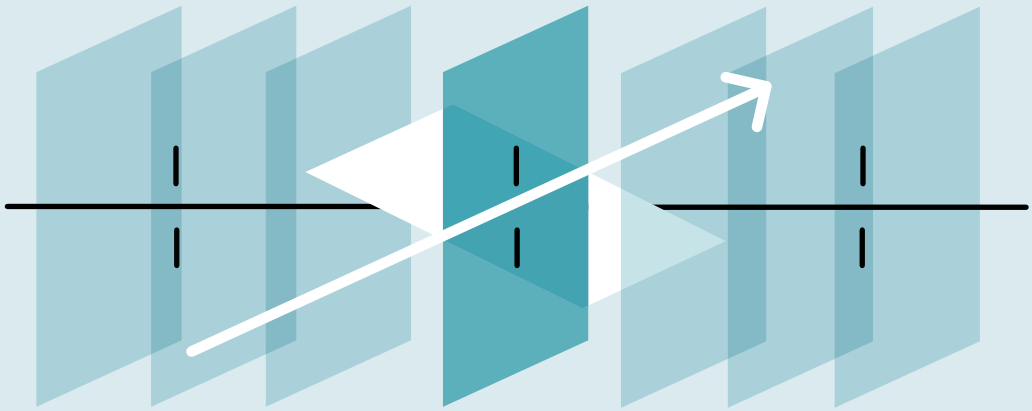

"We can start to see risk signals early, which helps us redirect resources before an issue becomes an incident. This [foresight] changes the entire value proposition of GRC."

Erin Dempsey Heuwetter
Head of audit, risk, and compliance
Navan

> **“With any AI or emerging technology, we can’t underscore enough the importance of clean data. If you put in garbage, it comes out garbage.”**
>
> Erin Banet, chief audit and risk officer, Humana

changes the entire value proposition of GRC—we become the early warning system, not the post-mortem.” These capabilities support proactive assurance, a core principle of connected risk. Tasks that once required days of manual synthesis can now be completed in hours, freeing audit and risk professionals to focus on higher-value activities like strategic guidance.

Beasley frames proactive assurance as a shift from episodic validation to continuous sensing, where AI helps audit and risk teams scan internal and external data for early signals and then intervenes before risks materialize as control failures or incidents.

Banet points to AI’s role in what she calls “storytelling,” linking audit findings to risk registers and control frameworks in ways that make residual risk more visible and actionable. To support that storytelling, her team plans to use AI to analyze years of audit reports, extract recurring themes, and benchmark risks against industry data.

AI also can help workers manage policies and retrieve knowledge by loading enterprise policies and standards into an AI-enabled GRC technology platform that allows teams to search, compare, and understand how requirements connect across the organization, pinpointing overlap and eliminating redundant controls that add complexity without reducing risk. “The system automatically loads policies so people can search and understand how things connect. Do we have duplication [and] overlapping? Are we instituting more risk?” Banet asks.

This form of artificial risk intelligence relies on AI and advanced analytics to continuously combine internal and external risk signals, add context across GRC activities, and emphasize data-driven decision making. In practice, it replaces static reports and periodic assessments with continuously refreshed insights to anticipate disruption, improve resilience, and become more proactive. However, both Heuwetter and Banet stress that predictive capability depends on a foundational requirement: data quality. Unified, trusted data—integrated across audit, risk, and compliance systems—is essential for AI to deliver reliable, contextual insight. Modern GRC platforms address this challenge by establishing a shared taxonomy and a single source of truth for risk, controls, and policies, replacing fragmented spreadsheets and siloed systems with consistent data models and traceable relationships. “With any AI or emerging technology, we can’t underscore enough the importance of clean data,” Banet cautions. “If you put in garbage, it comes out garbage.”

Instituting AI Governance and Guardrails

As AI becomes more deeply embedded in GRC processes, having disciplined governance procedures becomes increasingly important. The traditional Three Lines of Defense model is evolving as new operating models call for tighter connections across lines. First formalized by the Institute of Internal Auditors in 2013, the model was created to clarify roles and accountability in risk management after a series of high-profile corporate failures and financial crises exposed gaps in oversight. Since then, the model has provided a simple but powerful organizing principle for GRC teams: Frontline management owns and manages risk in day-to-day operations; the second line—risk management, compliance, information security, and related oversight functions—establishes policies, frameworks, risk appetite, and ongoing monitoring; and the third line—internal auditing—provides independent assurance to senior management and the board that governance, risk management, and internal controls are operating as intended.

Today, according to Beasley, the three lines are becoming blurred. And while the basic principles remain intact, execution is changing. With the first line of defense, he says, AI-enabled monitoring now allows operational teams to detect anomalies and control failures earlier, shifting risk management from reactive response to preventive action. With the second line, AI supports greater synthesis and correlation across risk domains—connecting compliance, cybersecurity, and operational data to provide context-rich oversight, Beasley explains. With the third line, he adds, internal auditing is

Some organizations are establishing cross-functional AI governance models that span the full life cycle of AI use—from design and deployment to monitoring, assurance, and accountability.

moving beyond episodic reviews toward continuous assurance, using AI-driven insights to focus on areas of highest risk while preserving independence and objectivity.

Both Banet and Heuwetter emphasize the importance of strong guardrails and human oversight as AI adoption accelerates. "Governance over AI—having the right tollgates in process, the right guardrails, training, and monitoring—all those things are extremely important," Banet says. "There always has to be a human in the loop."

The need for strong guardrails and human oversight increasingly extends to AI systems themselves, including tracking AI use cases across the organization, approving and documenting models before deployment, validating training data and assumptions, and assessing third-party AI providers. AI has become a distinct category of enterprise risk, not simply because it is widely deployed but also because the stakeholders often lack clear visibility into how AI systems operate, how decisions are generated, and how outputs can be explained, traced, and governed. This lack of clarity introduces new challenges related to model opacity, data lineage, hallucinations, bias, and accountability.

In response, regulatory bodies are beginning to mandate greater transparency into how AI systems are trained, governed, and monitored, particularly in high-impact use cases that can directly affect people, financial outcomes, or safety. In the European Union, the Artificial Intelligence Act establishes formal requirements for transparency, documentation, risk management, and human oversight of high-risk AI systems. In parallel, U.S.-based organizations are increasingly aligning with guidance such as the National Institute of Standards and Technology's AI Risk Management Framework, which provides a voluntary but influential structure for identifying, assessing, and managing AI-related risks across the system life cycle.

Leading organizations are already adapting their GRC practices to operate across this evolving regulatory landscape. At Navan, Heuwetter describes how legal, privacy, compliance, and information security teams work together to evaluate AI-enabled tools against regional regulatory expectations, ensuring that risk assessments, documentation, and human oversight requirements are consistent across jurisdictions. AI governance includes rigorous review of AI-enabled vendors, documentation of model behavior, and clear accountability for outputs. "AI can accelerate our work, but it doesn't absolve us from judgment," Heuwetter says. "If anything, it raises the bar for human accountability." Banet notes that regulatory scrutiny, particularly in Europe and other highly regulated markets, has reinforced the importance of traceability, documentation, and clear accountability, pushing audit and risk teams to demonstrate not only that controls exist but also how AI-driven decisions can be explained, validated, and defended to regulators.

For audit, risk, and compliance professionals, this convergence of AI adoption and regulatory scrutiny is further elevating GRC's strategic importance and necessitating even greater human oversight. Some organizations are establishing cross-functional AI governance models that span the full life cycle of AI use—from design and deployment to monitoring, assurance, and accountability. Internal audit teams are often drawn into this process, Banet adds, not only to provide assurance but also to help their organizations understand how AI risk is being managed across the enterprise ecosystem, including third-party vendors whose systems influence risk, compliance, and decision making.

Talent, Culture, and Workforce Alignment

The shift toward connected, AI-enabled GRC is not only changing how risk is identified and monitored but also fundamentally reshaping who does the work and how they collaborate. As routine work is automated, the traditional pyramid management structure is flattening into a diamond-shaped workforce model, with fewer entry-level roles and greater demand for mid-career professionals who can interpret data, apply judgment, and manage complexity. "Some firms are cutting back on hiring entry-level talent, but they're not thinking

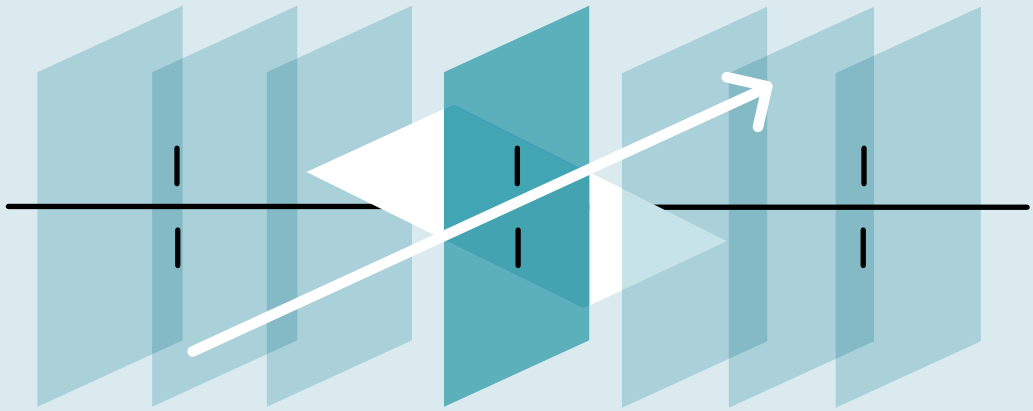

"Governance over AI—having the right tollgates in process, the right guardrails, training, and monitoring—all those things are extremely important. There always has to be a human in the loop."

Humana's Banet

> **“AI can be a powerful first draft, but it’s human judgment that ultimately determines whether it makes sense.”**
>
> NC State’s Beasley

about two to three years out, when they’ll need seasoned people,” Beasley warns.

Beasley’s warning underscores a broader challenge now facing audit and risk professionals: As entry-level roles shrink, organizations must find new ways to develop analytical, judgment-driven skills within their existing workforces, including greater emphasis on data literacy, AI fluency, and internal capability development. Banet has witnessed this shift firsthand. “A few years ago, it wasn’t required for audit professionals to have knowledge of AI or data analytics. Now, it’s a given,” she says, acknowledging the unresolved challenge of developing new hires without traditional training pipelines.

Heuwetter reinforces the point, emphasizing that future GRC professionals must be fluent not only in controls and compliance but also in data literacy, model oversight, and ethical reasoning. “There’s a counterintuitive dynamic at play,” she notes. “As we enhance these technologies, we’re actually likely to see more demand for auditors and risk professionals, not less. Better tools reveal just how much ground we weren’t covering before. The appetite for risk coverage grows with our capability to deliver it.”

As GRC professionals take on more analytical and technically sophisticated roles, those new capabilities must be matched by changes in how organizations collaborate and share information. Corporate culture is the connective tissue that holds these changes together, solidified by shared data. “It’s not just about systems talking to each other,” Heuwetter observes. “It’s [also] about people sharing information across departments.” When risk data is visible, consistent, and shared, alignment follows, she adds, and connected risk becomes not just an operating model but also an organizational mindset.

Conclusion

At its core, the evolution of GRC reflects a single imperative—to combine risk insight, assurance, governance, and human judgment into a connected intelligence fabric that helps organizations navigate uncertainty with confidence. AI is the thread that stitches that fabric together, transforming GRC from a retroactive control function into a forward-looking source of insight—one that helps organizations anticipate risk, adapt faster, and make better-informed decisions. This transformation hinges on several overlapping elements: connected risk operating models, continuous and predictive assurance, disciplined AI governance, unified data foundations, and a workforce equipped to apply human judgment alongside automation technology.

Together, these shifts signal a structural redefinition of GRC’s role in the enterprise. What was once episodic oversight is now becoming continuous insight, what was once siloed compliance is now becoming integrated risk intelligence, and what was once a defensive function is now emerging as a strategic stabilizer amid uncertainty.

For organizations navigating today’s complex risk landscape, the opportunity is clear: The convergence of global volatility, digital transformation, and AI-driven risk has elevated GRC from a necessary control function to a strategic stabilizer, one capable of helping forward-thinking executives make sense of disruption and reestablish trust.

By combining AI-enabled insights with human accountability, enterprises can move beyond reacting to disruption and toward anticipating it. “AI can be a powerful first draft,” says Beasley, “but it’s human judgment that ultimately determines whether it makes sense.”

VISIT US ONLINE

hbr.org/hbr-analytic-services

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

CRE74930326