April 2026

PYMNTS®

AI RUNS PAYMENTS. GOVERNANCE DECIDES WHAT HAPPENS NEXT

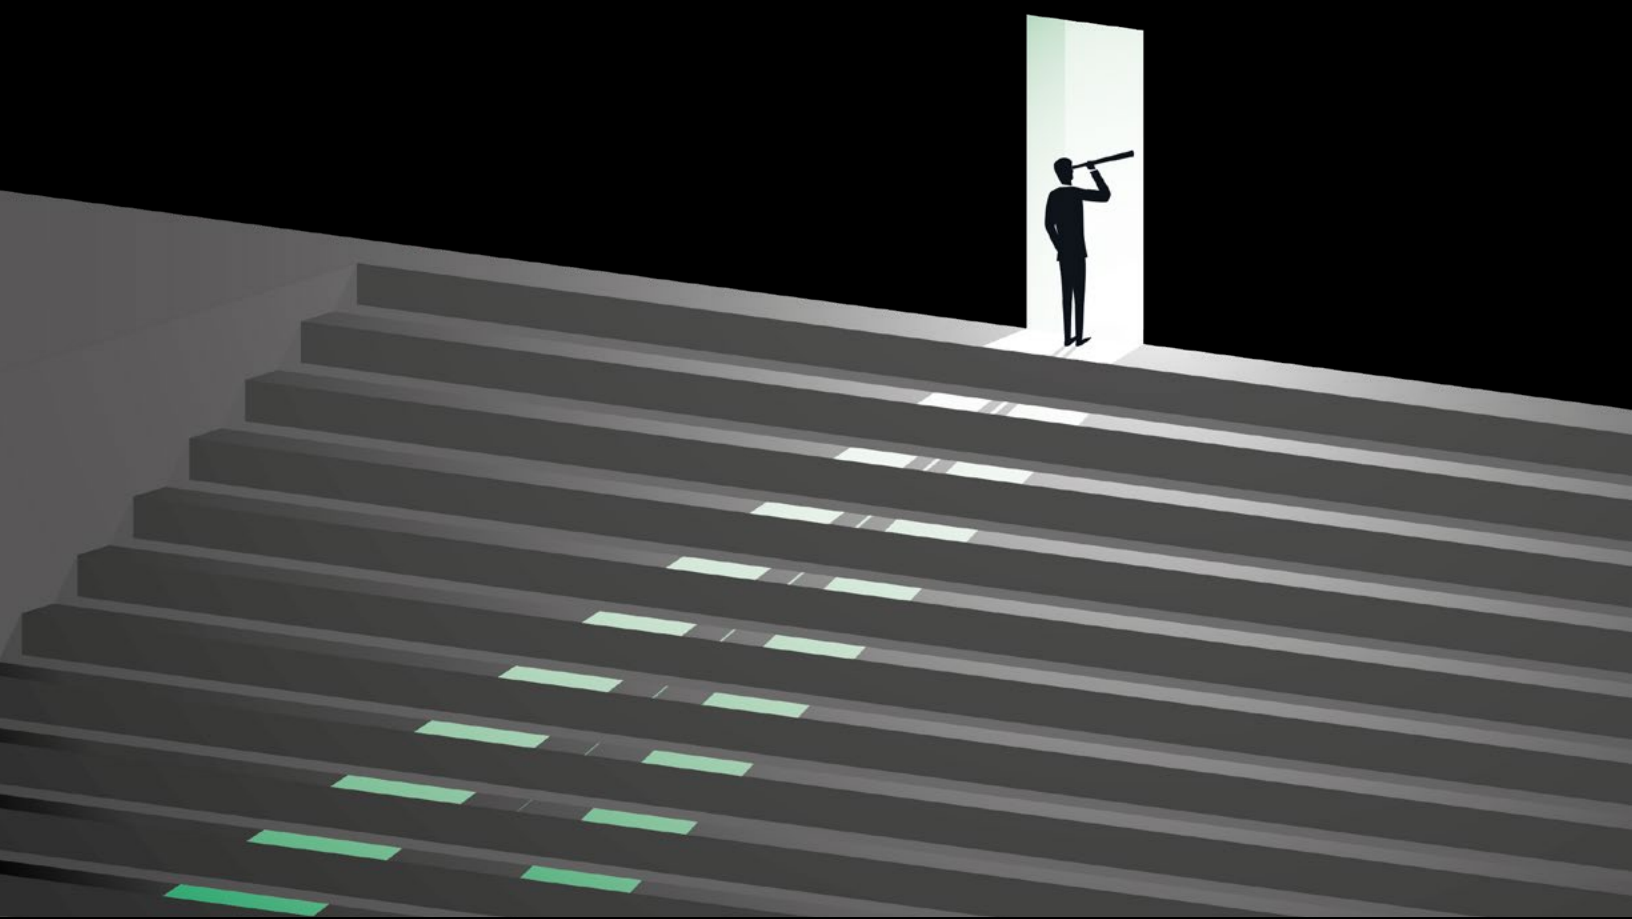

TABLE OF CONTENTS

PYMNTS®

AI RUNS PAYMENTS. GOVERNANCE DECIDES WHAT HAPPENS NEXT

PYMNTS®

THE REAL AI RACE IS NOT SPEED. **IT'S CONTROL**

Artificial intelligence is moving deeper into payments with unusual speed. It is helping companies spot fraud, improve approvals, personalize offers, support compliance, manage risk and shape customer experiences in real time. That much is clear. What is less settled is the part that now matters most: who is governing these systems, how that governance works in practice and what happens when AI begins influencing decisions faster than an organization can explain them. Across the essays in this eBook, one message comes through with force. In payments, AI is no longer the experiment. Governance is the differentiator.

That is because payments is one of the hardest environments in which to get AI wrong. These systems do not live in a lab. They operate inside live transaction flows, fraud programs, onboarding journeys, identity checks, credit decisions and customer service interactions. A weak model can create problems. A weak governance structure can multiply them. At scale. The risk is not only that AI makes a flawed call. It is that nobody is fully accountable for the outcome, nobody can trace the logic, and nobody spots the drift until customers, regulators or partners do first.

The executives in this collection return again and again to a handful of hard truths. Governance tends to break down in the gaps between teams. Product may own the feature, engineering may own the model, compliance may own the policy, and operations may own the day-to-day consequences, but the end-to-end accountability is often blurred. At the same time, many AI systems depend on third-party models, vendors, data providers and external platforms, which means companies are being asked to govern not only what they build, but also what they rent, ingest and rely on. In that environment, governance cannot be treated as a meeting, a checklist or a slide deck. It has to become an operating discipline.

Another theme runs through these pages: Speed is seductive, but speed without structure is expensive. Payments companies are under pressure to automate more, move faster and show returns quickly. Yet several contributors make the same point from different angles. The real work is not simply deploying AI. It is building the data foundations, oversight mechanisms, fallback plans, audit trails and human review points that let a company move quickly without losing control. In other words, the organizations that benefit most from AI may not be the ones that rush first. They may be the ones that prepare best.

Readers should come away from this eBook with more than a warning. They should come away with a playbook. These essays offer a practical look at where governance breaks down, what strong organizations are doing earlier, which questions boards and CEOs should be asking now, and how leading payments executives are thinking about explainability, accountability and trust in an AI-driven market. For leaders across banking, payments and FinTech, that has real value. It can help sharpen internal conversations, expose blind spots in current governance models, and frame AI not as a race to adopt the newest tool, but as a long-term test of institutional discipline.

AI may be powering the next era of payments. But governance will decide which companies can scale it with confidence, which ones can defend it under scrutiny, and which ones can turn automation into durable advantage. That is the conversation this eBook begins.

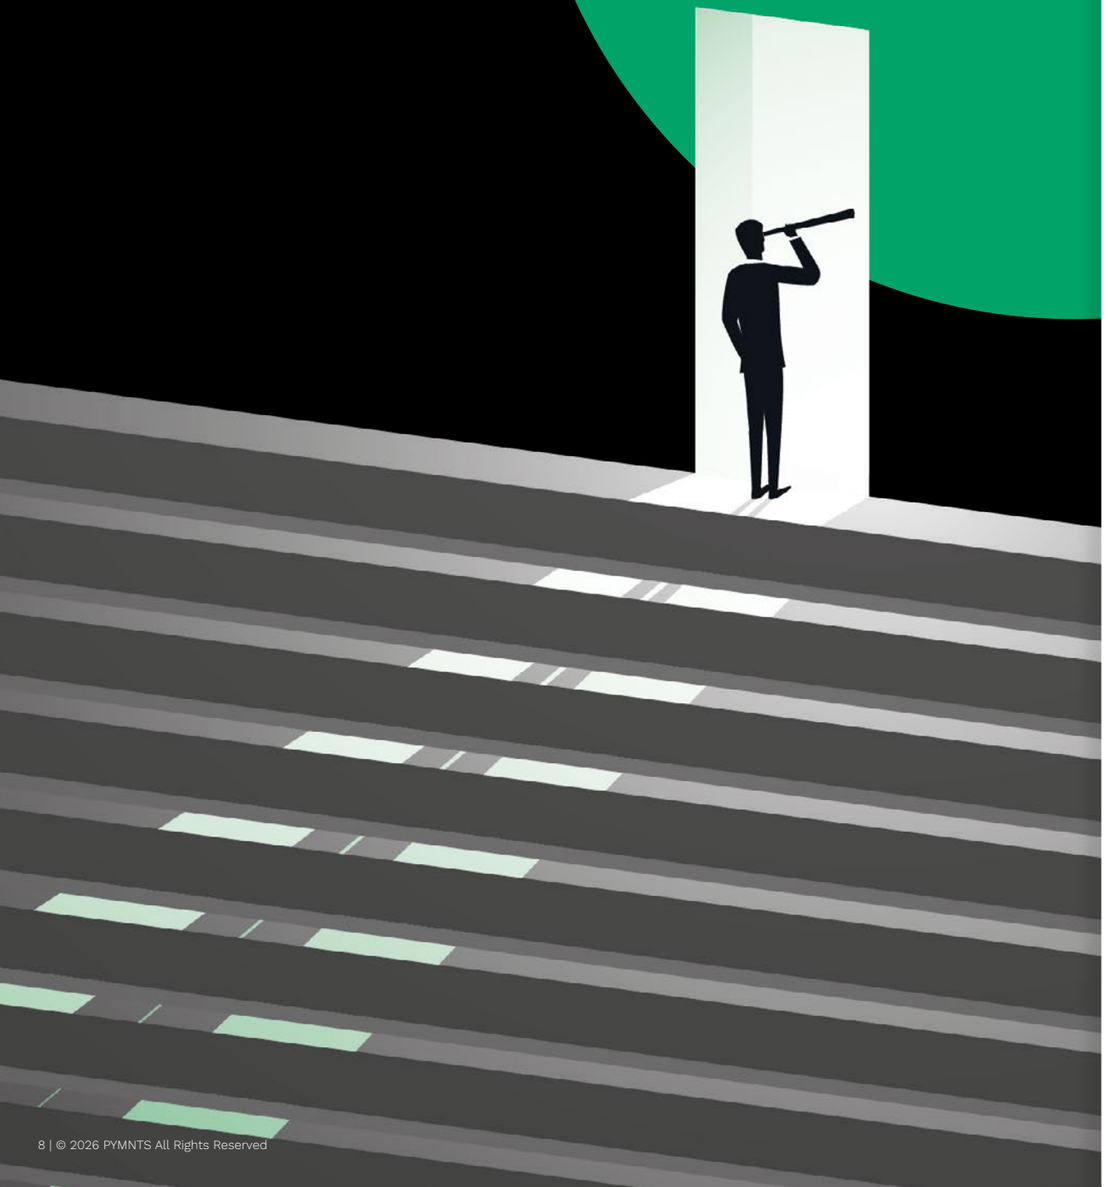

billtrust

AHSAN
SHAH
SVP, Analytics and AI

AI IN PAYMENTS RUNS ON DATA. **GOVERNANCE DECIDES WHETHER IT WORKS.**

As artificial intelligence becomes core infrastructure across the payments ecosystem, the conversation has shifted. We're no longer debating whether to use AI. We're figuring out how to keep it reliable, explainable and trustworthy when the pace of innovation is outpacing the guardrails around it.

For payments companies, governance isn't abstract. AI influences whether a transaction is flagged as fraud, how credit risk is assessed and how fast suppliers get paid. A misstep doesn't just create noise; it creates risk at scale.

What I've seen, both inside Billtrust and across the industry, is that most AI governance challenges don't come from the models themselves. They come from the data. Payments data is messy by nature. It moves through ERPs, banks, suppliers, FinTech partners and procurement systems with their own formats and levels of quality. If we don't solve for consistency, ownership and lineage before we build the model, the best algorithms in the world can't save us. Governance breaks down at the point where trusted data should exist but doesn't.

And that leads to the hardest trade-off: speed. Everyone wants to deploy AI quickly because the efficiency gains are real. But going fast without disciplined data governance is a false victory. You end up paying that debt later in rework, model drift or customer friction. At Billtrust, we've learned to slow down at the beginning so we can move faster and more confidently later. It's not a glamorous answer, but it's the truth.

Another challenge unique to payments is that so much of the data and model performance depends on third parties. It's not enough for your own house to be in order if your partners don't meet the same standards. Governance becomes a collective effort: aligned definitions, shared expectations around data quality, visibility into upstream changes and contractual clarity around how data can and cannot be used. You can't govern in the dark.

Data governance is not an engineering issue — it's a company wide discipline. Legal, compliance, operations, product and security all have a stake in how AI behaves. Formalizing that collaboration earlier gives you the structure you need to keep pace with innovation without cutting corners.

Looking ahead, I think boards and CEOs need to push past the broad question of whether AI is being used responsibly and ask something far more specific: Can we explain every AI-driven decision that affects our customers? If you can't answer that today, regulators, partners, and customers eventually will ask. And the expectation will be high.

AI will continue to reshape payments. Whether it reshapes it responsibly depends on whether governance evolves with the same urgency as the technology itself.

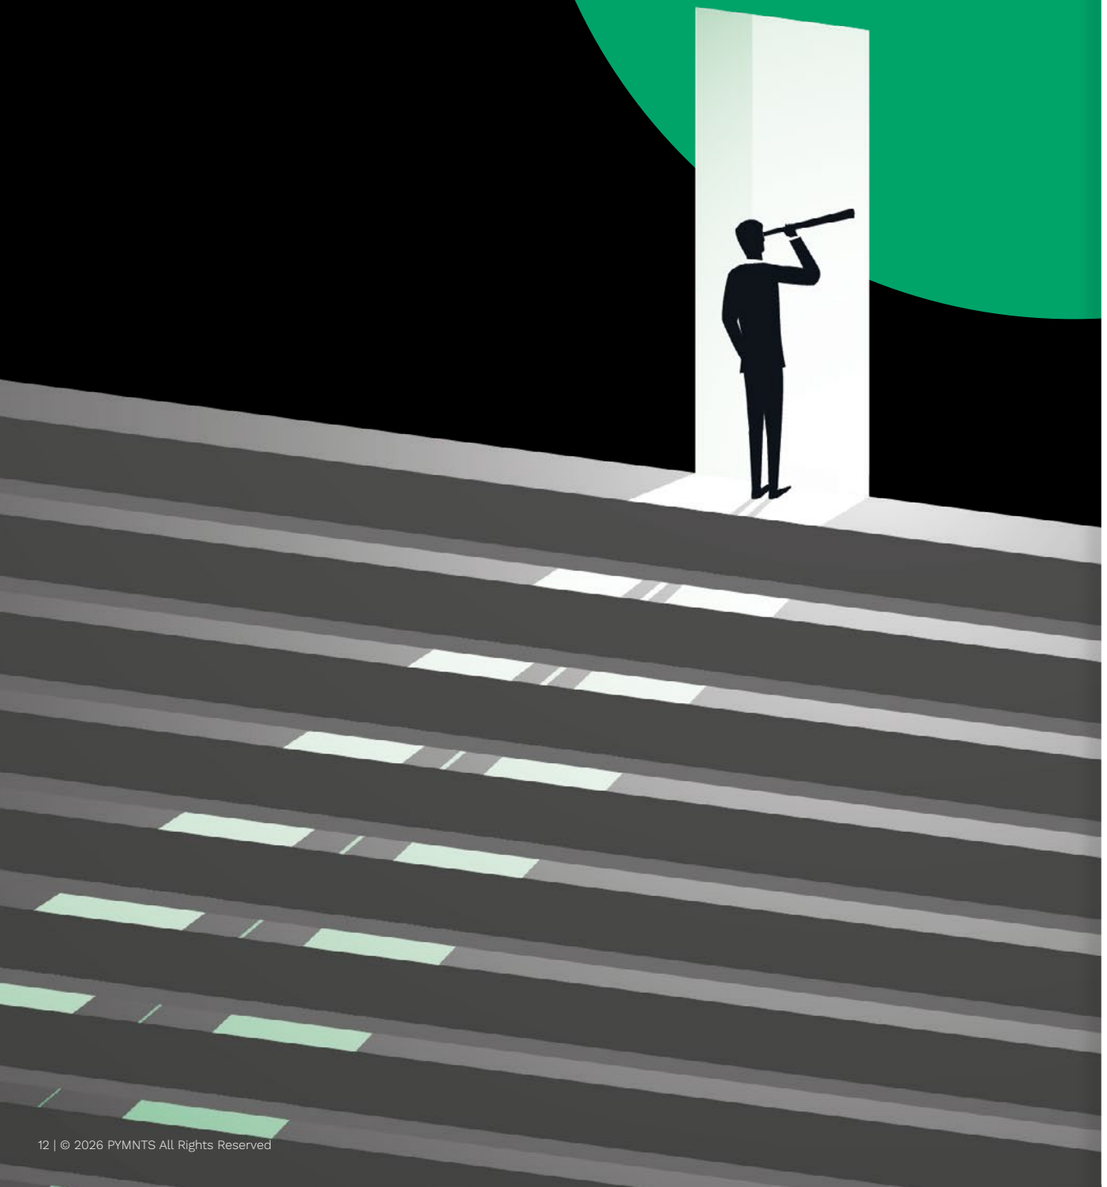

boost
payment solutions®

ELLY
AIALA
Chief Compliance Officer

AI IS POWERING PAYMENTS. WHO'S GOVERNING THE AI?

AI governance doesn't typically fail because companies lack policies. It fails in the gap between intent and behavior.

One of the most common breakdowns is simple: employees using AI tools to move faster and unintentionally exposing sensitive or proprietary information. It's not malicious; it's convenience. People want quick answers, and AI delivers. But without guardrails, that convenience can introduce profound risk.

Some organizations are investing in private LLM environments or segmented data layers. That helps, but it introduces a different question: is the ROI there? Building or buying controlled environments is expensive. Choosing not to invest, however, means accepting the risk that employees will use external tools without oversight. Governance becomes a business decision, not just a security one.

That tension shows up in the age-old trade-off between speed and control. There is significant pressure to adopt AI quickly, especially in payments. That urgency can lead to rushed decisions — moving forward in partnerships and new processes without appropriate visibility. The more loosely artificial intelligence tools are adopted, the harder it becomes to track usage, enforce compliance and maintain consistent security standards.

In a highly regulated environment with sensitive data, that trade-off is real. You can move faster, but you then increase exposure across data management, third-party dependencies and regulatory requirements. Most organizations do not yet have full visibility into how AI is being used across their teams, which makes controlled acceleration difficult.

The challenge becomes more complex when AI depends on third parties. Governing AI in that environment requires due diligence that goes beyond traditional vendor onboarding or third-party reviews. If you cannot clearly explain to a colleague, customer, partner or auditor what your AI is doing, how it is doing it and why it is being used, you are not ready to deploy it.

One governance decision many organizations may wish they made earlier is defining success metrics before implementation. Without clear metrics, governance is reactive and much more complex than if you had tackled it prior to build/integration. These are the questions boards and CEOs should be asking of their teams implementing AI. Instead of asking the general question of whether AI is being used, leadership should be asking: Who owns governance? What metrics support that governance? Do we have the internal expertise to appropriately implement and monitor our AI usage? At Boost, we take the same approach. Governance is treated as a prerequisite, not an afterthought, with a focus on maintaining visibility and accountability across how AI is used.

Finally, cost remains underexamined. Many organizations are investing in artificial intelligence without a clearly defined business case, chasing after the latest “hype.” When you factor in due diligence, hiring, implementation, infrastructure and ongoing maintenance, the cost can be significant. If AI is not tied to specific, high-value use cases, the risk is not just operational, it is financial.

AI has the potential to transform payments, but without disciplined governance, it can scale risk just as quickly. Regardless of where you are on your AI journey, take inventory now. Poll your team on what and how AI is being used today. You may be surprised.

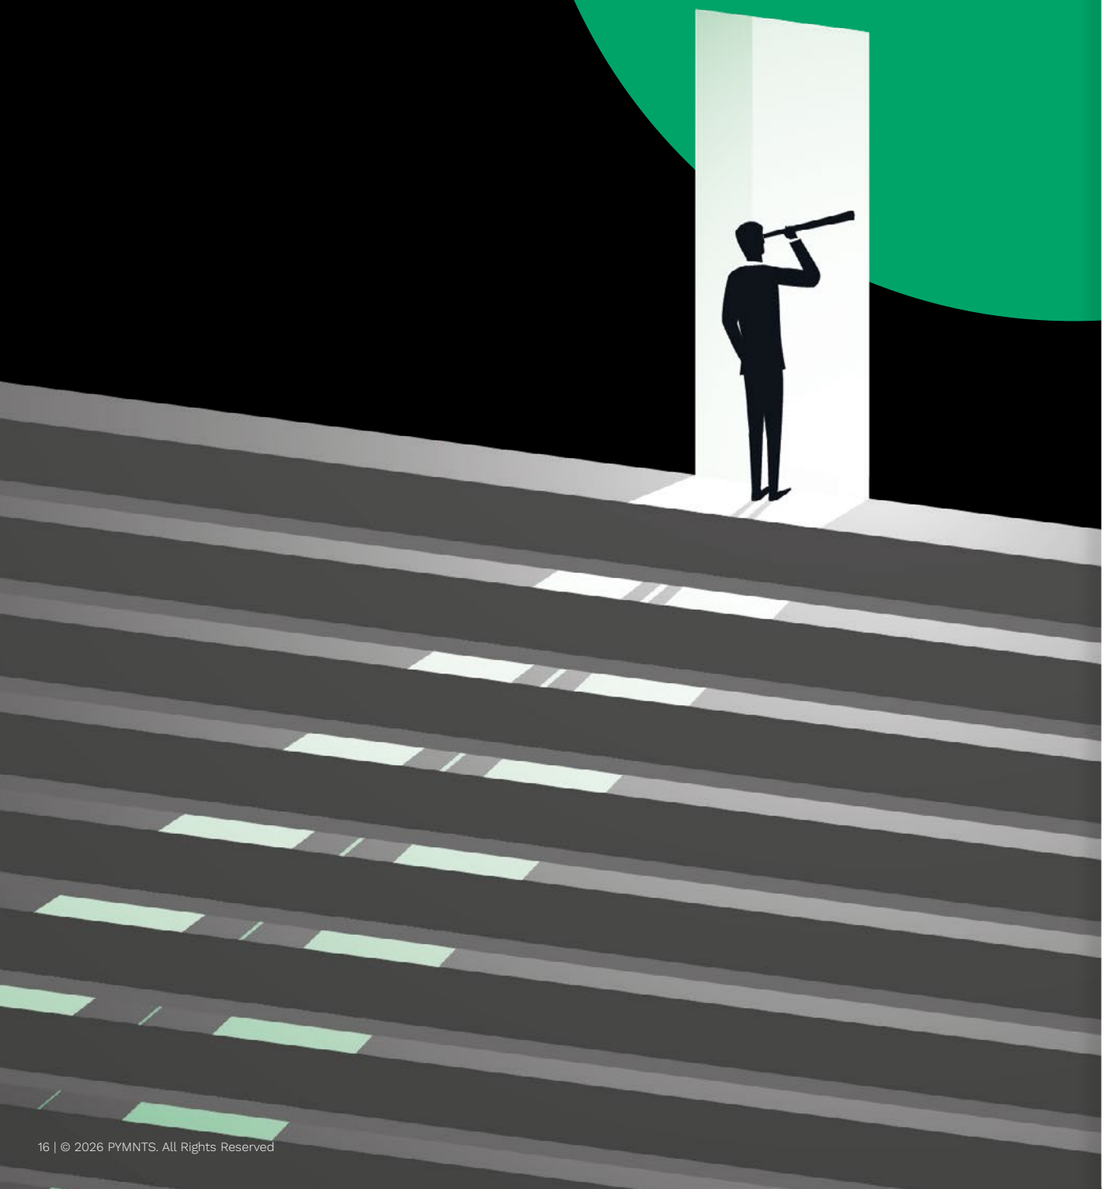

braze

KIPP
JOHNSON
Senior Director, AI Solutions Consulting

YOU FINALLY BUILT THE F1 CAR. **DID YOU HIRE THE RACE ENGINEER?**

Companies have long aspired to deliver highly personalized experiences — marketers dream of the right message, right channel, right offer, right time, to every customer. With AI agents, this is no longer hypothetical. It is happening on a true 1:1 basis, with companies delivering millions of unique experiences simultaneously. But as exciting as this sounds, it opens up unique risks requiring a rethink of historical governance practices.

Think of a cyclist in the Tour de France. They are in direct control of their bicycle — if a tire goes flat or a chain slips, they feel it immediately and react. Now imagine a Formula 1 (F1) driver competing in a Grand Prix. The car is incredibly powerful, highly complex, and moves faster than human reaction times can manage unassisted. If a component fails at 200 mph, the result is instantaneous and catastrophic. The old adage was if you want to drive fast, you need powerful brakes — but today it is more accurate to say: you need robust telemetry, race engineers on the pit wall and extensive evaluation protocols before the car touches the track.

The same is true of our increasingly powerful artificial intelligence agents. Last year, Braze AI Decisioning Studio made over 17 billion personalized decisions based on each customer's unique context. Like an F1 team, this was only possible because of deliberate investment in governance, observability and specialists to manage these programs.

WHERE GOVERNANCE MOST OFTEN BREAKS DOWN

No F1 team puts a car on the track before answering five questions: What is the goal? What are the upstream and downstream dependencies? What are the guardrails? What are the fallback options? How will performance be continuously measured? Too many organizations build before they can clearly answer those. One enterprise discovered mid-build that a critical data source carried a 72-hour delay — their models were learning from incomplete data, requiring a project redesign.

A second failure pattern: underestimating the expertise required. Even companies with strong data science teams often lack critical domains like reinforcement learning and underestimate the challenges of production at scale. The teams that succeed treat specialist support as an extension of their pit crew. One example: a company driving app usage discovered customers were opening the app directly after seeing marketing emails rather than clicking through — the AI was learning from flawed attribution signals. Catching it required domain expertise and tooling most internal teams don't have.

GOVERNING AI WHEN YOUR DATA DEPENDS ON THIRD PARTIES

An F1 car races in conditions nobody fully controls. Race engineers don't wait to see what happens — they red-team every scenario and design for graceful failovers before the car leaves the garage. The same discipline applies to third-party data dependencies. Assume something will be wrong. The most resilient programs maintain a non-AI-dependent control population as a fallback. When a data pipeline broke for three weeks on one program, deterministic rules kept operations running while the issue was resolved. Running underlying data assets alongside modeled scores lets you validate incremental model value and flags when performance starts to drift.

WHAT BOARDS AND CEOS SHOULD BE ASKING

Most conversations focus on AI strategy. The sharper questions are about observability: How are we evaluating systems that influence real-time financial decisions? What happens when an upstream dependency fails? Are our model inputs — not just outputs — transparent to compliance, risk and product teams, not just data scientists?

The old adage was right all along. The winning teams know that it doesn't only require a faster car — it requires a specialized pit crew and proactive governance from the start.

PYMNTS®

AI PAYMENTS. GOVERNANCE DECIDES **WHAT HAPPENS NEXT**

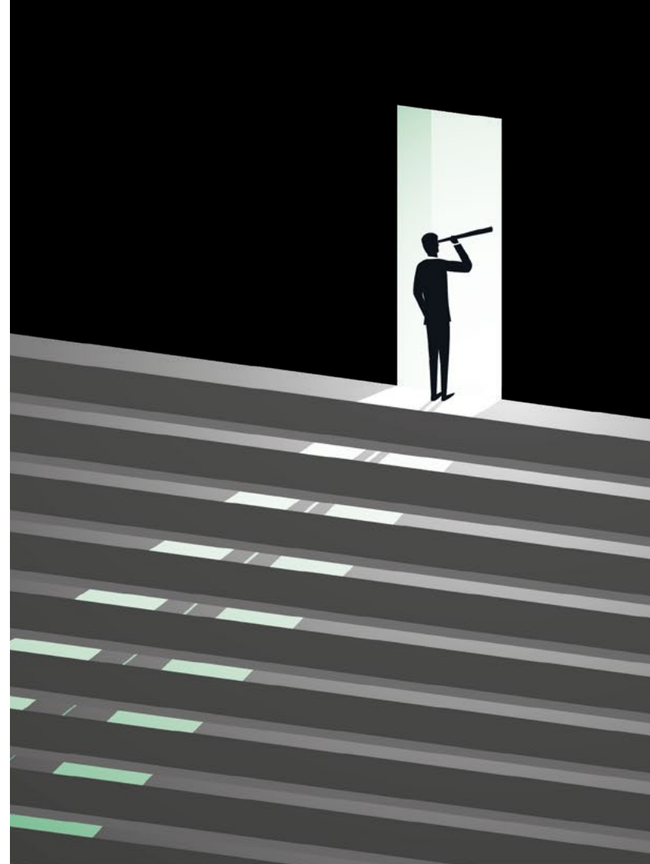

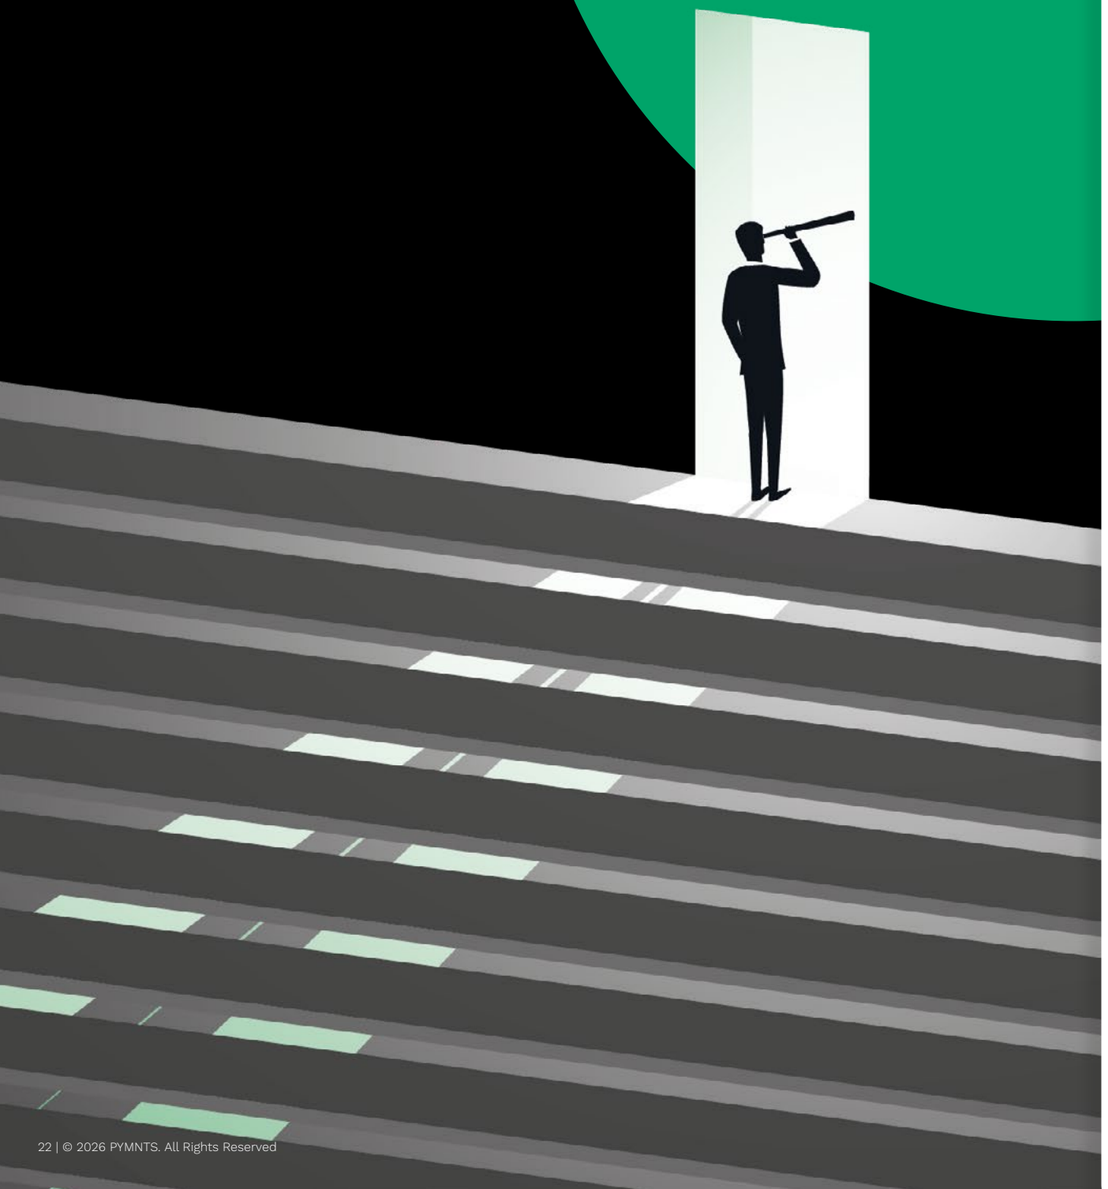

MLADEN
VLADIC
Head of Product Management, Payment Networks

GOVERNING AI IN THE **NEXT ERA OF PAYMENT**

The shift from transactional to relational payments is accelerating. We are entering the era of agentic commerce, where AI agents act on behalf of shoppers to source, negotiate and complete purchases. Artificial intelligence is no longer just a back-end tool — it is becoming the primary engine of customer engagement, which represents a new, unique opportunity for brands to reimagine consumer

lifecycle engagement strategies that have been in place for many years. That shift has forced us to rethink governance frameworks we had taken for granted.

With AI agents projected to help orchestrate up to $1 trillion in U.S. retail revenue by 2030, AI algorithms are already influencing real-time decisions, financial access and customer trust. The core question is no longer whether we adopt AI, but how we govern it to ensure shared success across the ecosystem.

WHERE GOVERNANCE BREAKS DOWN

In our experience, AI governance most often breaks down at the point of integration. We see ecosystems where AI models operate in silos — where purchase-event signals and item-level intelligence do not communicate securely within a single infrastructure. When that happens, visibility is lost. Governance fails when organizations treat AI solely as a post-purchase optimization tool rather than embedding it directly into authorization, authentication and dispute networks from the start. We learned early that governance has to be architected into the payment flow itself, not layered on after the fact.

THE SPEED VS GOVERNANCE TRADE-OFF

The hardest trade-off we navigate is balancing the rapid deployment of hyper-personalized offers against the strict governance of item-level data. Consumers expect frictionless, instant rewards at checkout. But building solutions we can stand behind has required us to prioritize secure, receipt-backed proof of performance over simply rushing a model to market. Speed without precision creates risk at scale. That has meant making deliberate choices to slow down in places where others might not — and those choices have paid off in the durability of what we have built.

GOVERNING ACROSS THIRD PARTIES

Governing AI across third-party networks requires a funder-agnostic approach grounded in universal standards. Card issuers, merchants and consumer brands must operate within a tightly integrated ecosystem where intelligence travels securely with the payment itself. Whether we are analyzing the purchase event or the specific items in a basket, we enforce strict data guardrails so that agents can transact safely across all participating partners. That discipline is what makes a four-party model — consumers, merchants, issuers and brands — actually work.

THE DECISION I WISH WE HAD MADE EARLIER

If there is one governance decision I wish the industry had made sooner, it is defining exactly what AI can "see" and act upon. Establishing firm boundaries around item-aware decisioning from day one — clearly distinguishing between a total-receipt view and a digital-receipt view, for example — prevents privacy bottlenecks downstream and builds consumer trust that agentic commerce depends on. We have seen firsthand how retroactively drawing those lines is far more costly than setting them at the outset.

THE QUESTION BOARDS SHOULD BE ASKING

Executive teams frequently ask whether their AI infrastructure is secure. The better question is: how does our AI governance scale when algorithms, rather than humans, initiate the purchase?

In an agentic world, trust is the ultimate currency. The data is available and the technology is ready. The leaders who invest now in governance frameworks that turn every AI-initiated interaction into a secure moment of value creation will define the next era of payments.

PYMNTS®

AI PAYMENTS. GOVERNANCE DECIDES **WHAT HAPPENS NEXT**

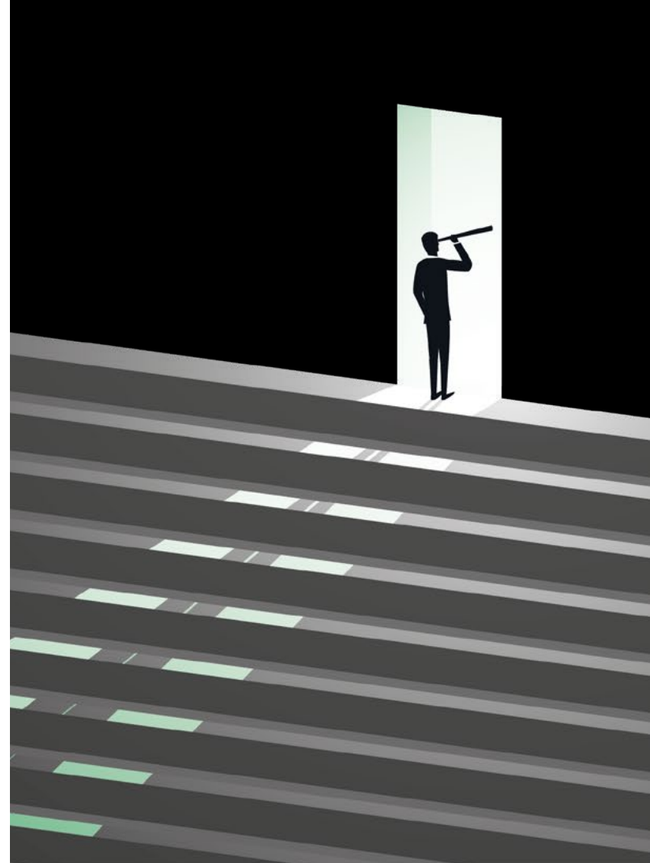

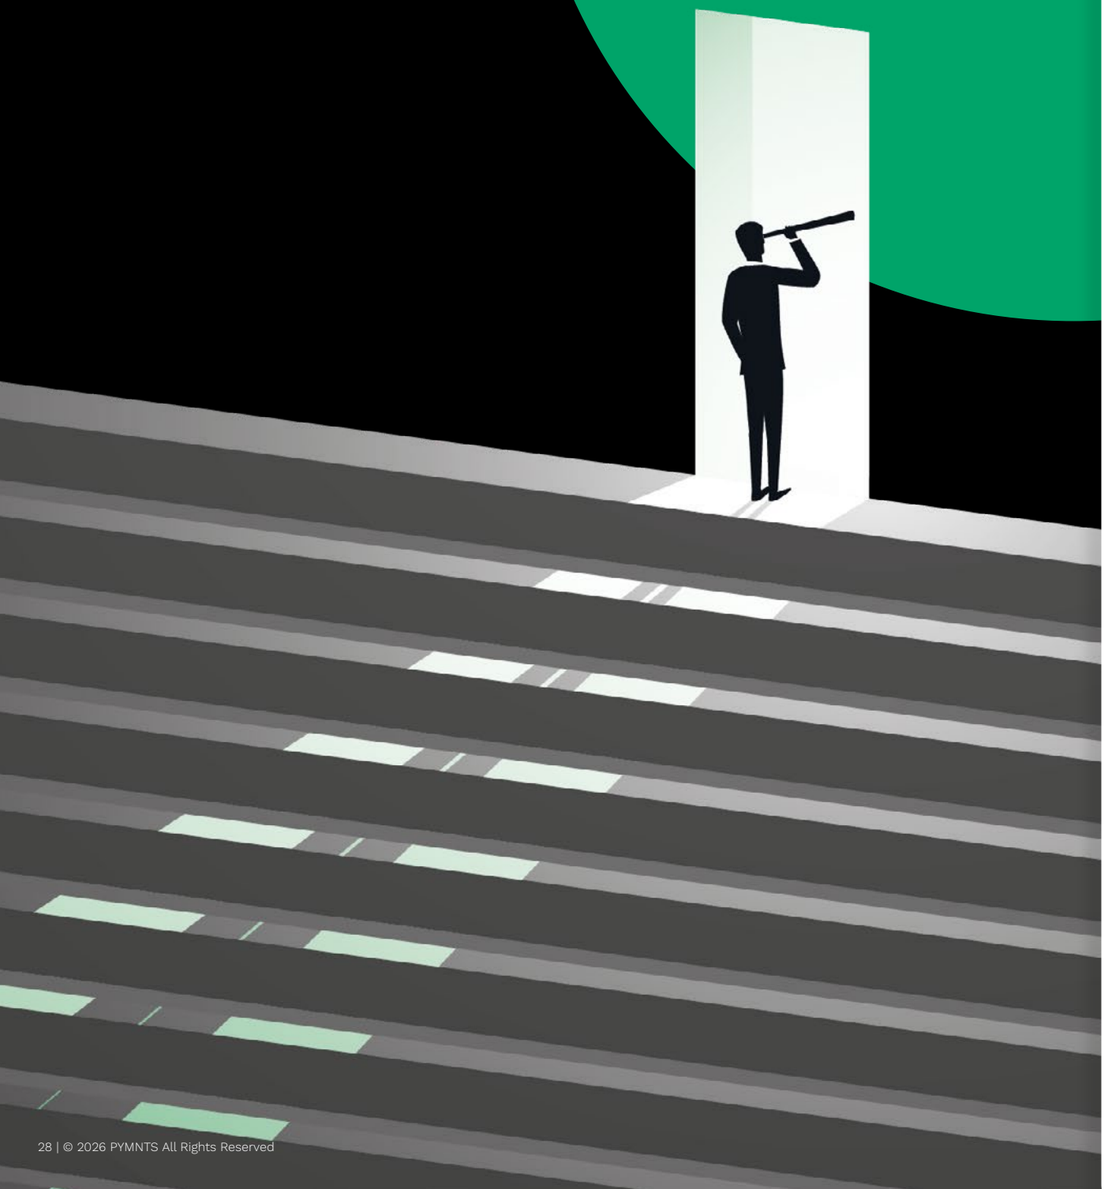

Flagright®

BARAN
OZKAN
Co-founder and CEO

AI WILL DECIDE PAYMENTS; **GOVERNANCE WILL DECIDE WHO SURVIVES**

From where I sit at Flagright, AI governance usually breaks down in the gap between a model output and a business decision. Teams spend a lot of time discussing model accuracy and far less time defining who owns the decision, what evidence must be logged, when a human has to step in and how a change gets approved. In payments, weak governance does not fail quietly. It scales across approvals, blocks, investigations and customer outcomes. The real issue is rarely “bad AI.” It is unclear operating discipline around artificial intelligence.

THE HARDEST TRADE-OFF BETWEEN SPEED AND GOVERNANCE

The hardest trade-off is resisting the urge to automate before you can simulate. Everyone wants faster alert handling, lower manual workload and quicker decisions. But if you cannot replay historical cases, compare outcomes and roll back safely, speed just lets you operationalize mistakes faster. We have repeatedly chosen slower initial deployment in exchange for staged release, version control and explicit human override points. That can feel expensive in the short term, especially in fast moving payments environments, but it is much cheaper than explaining a scaled failure to regulators, customers or your board.

GOVERNING AI WHEN DATA AND MODELS DEPEND ON THIRD PARTIES

You can outsource tooling, but not accountability. Every external model, screening source, identity signal or data provider should be treated as a governed dependency. That means clear performance expectations, change notification rights, auditability, fallback plans and independent validation inside your own environment. The mistake many firms make is assuming vendor assurance is enough. It is not. Third-party AI introduces model risk, vendor risk, data provenance risk and concentration risk at the same time. Good governance starts by assuming that any dependency can change underneath you.

ON GOVERNANCE DECISIONS

I wish more teams turned policy into something operational and versioned much earlier. Governance should not live in slide decks or committee notes. It should live in rule level investigation logic, approval thresholds, evidence standards and a record of what changed, when and why. Once you do that, governance stops being abstract. It becomes testable. You can replay past alerts, challenge outcomes, and prove that your controls behave the way your risk appetite says they should. That is when AI stops being a black box and becomes a governable control.

WHAT BOARDS AND CEOS SHOULD BE ASKING ABOUT AI GOVERNANCE

Most boards still ask whether AI is being used responsibly. That question is too soft. The better questions are: Where is AI making or shaping customer impacting decisions today? Which decisions are fully automated, which are recommendation only and where is human review mandatory? How quickly would we detect drift, bad data or a poor release? What is our rollback plan? And if a regulator or major partner asked us tomorrow to justify our AI controls, what evidence could we produce in 24 hours? In payments, the winners will not be the firms that automate the most. They will be the ones that can automate confidently and explain their decisions under pressure.

AMIR
WAIN
CEO

THE REAL RISK IN AI PAYMENTS IS NOT THE MODEL — **IT'S THE GAPS AROUND IT**

Artificial intelligence is no longer just influencing decisions in payments — it is starting to make them. And the pace of that shift is accelerating faster than most institutions are prepared to govern. There has never been a technology where development and adoption have occurred at the rate we are now seeing with AI.

That acceleration brings both opportunity and risk. On one hand, it has the potential to be a great equalizer, giving small and mid-sized institutions the ability to compete with much larger players. But it is also a threat. On the other hand, it forces a more urgent question: if AI is acting, who is accountable?

To answer that, it helps to step back. When I think about AI in payments, I don't start with the model — I start with the system. Architecture ultimately determines the destiny of the business.

For years, the industry has relied on a fragmented approach — stitching together different systems to solve business problems across credit, debit and core banking. Over time, that fragmentation created gaps. When AI is introduced into that environment, governance doesn't fail at the model level — it fails in the gaps between systems, decisions and ownership.

That realization led to a different architectural choice: building a platform that is not product-centric, but customer-centric — product-agnostic and geography-agnostic. The premise is simple: there will always be new products we cannot predict today, so the system itself must be adaptable.

That choice was not the fastest path. It was a big decision, a difficult one and it took longer up front. But it created consistency and control — principles that now apply directly to AI governance.

Governance is not a one-time decision; it is an ongoing discipline and a continuously evolving effort that must be built into the organization. At the same time, product and feature cycles are compressing rapidly.

Nowhere is that tension more visible than in fraud. It is easy to eliminate fraud entirely — you could decline every transaction and have zero fraud. But that is not a strategy. The real objective is minimizing friction while maximizing fraud capture.

The risk environment is also shifting quickly. Platforms that cannot respond dynamically to emerging threats will fall behind. Preparation requires agility, scalability and real-time intelligence.

This is where the next phase becomes critical. Agentic AI is not just about smarter models, it is about systems that can perceive, reason and act within defined guardrails, learning from outcomes over time.

But autonomy does not remove accountability. Responsible AI governance is non-negotiable. It requires transparency, consent and traceability in how data is used. And critically, the human role becomes more strategic — overseeing AI to ensure decisions remain fair, explainable and aligned with business outcomes.

The institutions that succeed will not be the ones that simply move fastest. They will be the ones that build the discipline, architecture and accountability to govern AI — because in payments, every decision is now real time, automated and consequential.

In a world where AI makes the decision, governance is what earns the right to make it.

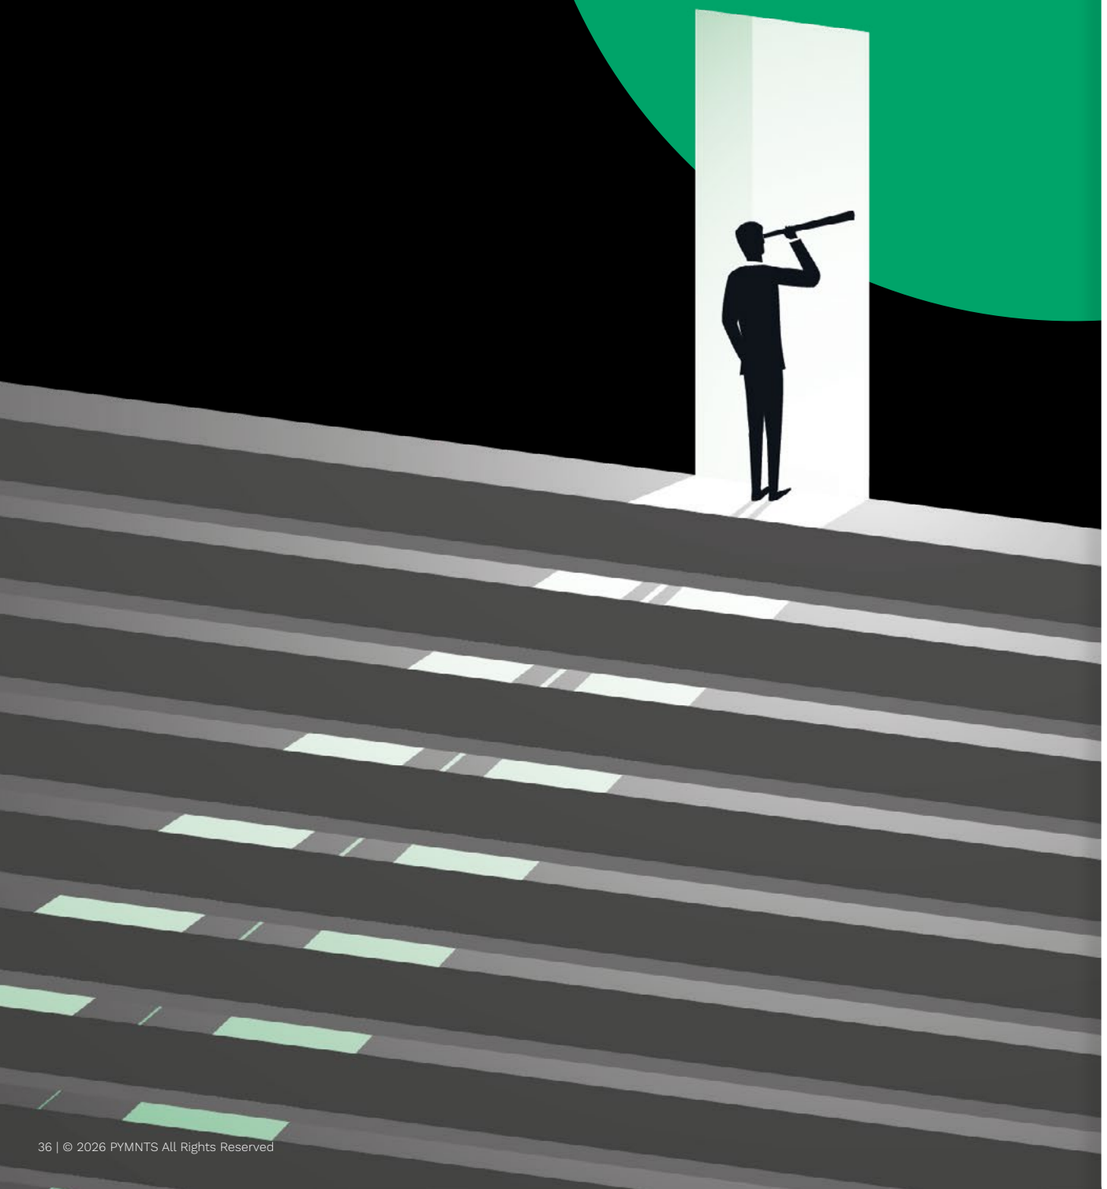

maverick

JUSTIN
DOWNEY
VP, Product

GUARDRAILS BEFORE GROWTH: STRATEGIC AI ADOPTION IN PAYMENTS

Artificial intelligence is already embedded in how payments companies operate. At Maverick Payments, we are intentional and not reactive when strategically implementing AI to handle repetitive and mundane decisioning to streamline non-complex tasks. The question is not whether AI adds value, but how to leverage and govern it with guardrails in an environment where small errors can create outsized consequences at scale.

Where we're seeing AI governance breaking down in the field is at the boundary between automation and accountability. AI excels at gathering signals, validating information, finding patterns and surfacing anomalies with non-complex use cases. Problems typically emerge when organizations allow AI assumptions for complex use cases to quietly become finite decisions without reviewing them for nuances; therefore, we believe governance fails when AI insights are not treated as conclusions rather than inputs.

The hardest trade-off between speed and governance is resisting the urge to fully automate complex and high-risk related decisions. AI can materially accelerate KYC checks, risk analysis and operational workflows, and the pressure to move faster is real, but speed without oversight can create risk or missed opportunities at scale. Our strategy is deliberate: use AI to handle simple, data-heavy groundwork while reserving final decisions for experienced operators. That hybrid approach lets us move faster, responsibly.

Governing AI becomes more complex when data and models depend on third parties. Payments ecosystems rely on external data sources, vendors and platforms, each with their own assumptions and models. Strong governance requires clear visibility into how inputs are used, what data is retained and where accountability sits when something goes wrong. AI may sit inside a partner's tool, but the risk still shows up with your regulators, merchants and balance sheet. That reality demands disciplined vendor governance, model transparency where feasible and contractual clarity on controls and escalation.

The one governance decision we wish we made earlier was a formalizing role of human expertise alongside AI in a hybrid model from the start. Early AI deployments often focus on efficiency gains. What matters equally is defining where human judgment is mandatory. AI is a powerful data miner, but it lacks context, intuition and industry experience. When we explicitly position AI as an accelerator, outcomes improve and internal trust in the systems increases.

For boards and CEOs, the most important questions are often the ones not being asked. Not "Are we using AI?" but "Where can AI influence outcomes and speed up decisioning?" Not "How accurate is the model?" but "What happens when it's wrong at scale?" Leaders should be asking who owns AI decisions, how exceptions are handled, how third-party models are governed and how controls evolve as models and data sources change.

AI's real opportunity in payments is not replacing people — it is elevating them. When AI handles repetitive, data-intensive work, experts focus on judgment, relationships and strategy. The foundations of trust in financial infrastructure offer opportunity or companies. The companies that win will be those that implement AI strategically and govern it thoughtfully and with focused intent.

North

CHRISTOPHER
MASCARO
Chief Cyber and Fraud Officer

A UNIFIED APPROACH TO **THIRD-PARTY AI**

There is a gap in organizational AI governance: Who is accountable for AI outcomes, and who has the information and control needed to manage them? The board is accountable to shareholders and regulators, executives are accountable for strategic risk decisions. These decisions are then passed to operational teams who are accountable for day-to-day management of systems whose most consequential components were not developed by them.

Closing this gap requires a shared understanding across all levels of the organization. A grasp of what third-party AI risk looks like, why it demands dedicated governance and what role each level of leadership and operations plays in ensuring responsible outcomes. Without this shared understanding, accountability becomes diffuse and governance unreliable.

THE AI SUPPLY CHAIN AND RISKS

A manufacturer's product quality depends on every supplier in its chain. An organization's AI outcomes depend on every vendor in its ecosystem. A biased foundation model produces biased applications. An opaque vendor creates ungovernable dependencies. Vendor concentration creates systemic risk extending well beyond any individual deployment.

This supply chain lens clarifies the governance task: Map the AI supply chain, assess risk at each node, establish governance standards that apply consistently across the chain and build the monitoring capabilities needed for ongoing compliance. It also clarifies accountability. An organization cannot disclaim responsibility for a vendor model's harmful output any more than a manufacturer can disclaim responsibility for a supplier's defective component. The accountability follows the deployment, not the development.

EXPLAINABILITY AS SHARED LANGUAGE

Model explainability provides the common language connecting technical, operational and strategic governance activities. For the board, explainability findings translate into clear, evidence-based reporting on whether AI systems operate within acceptable bounds. For risk functions, explainability provides the analytical basis for quantifying and managing vendor AI exposure with precision. For operational teams, explainability tools provide the real-time visibility needed to detect and respond to changes in vendor model behavior before they escalate into incidents.

Without explainability, each level operates with a different and incomplete picture. The board sees policies and compliance checklists but cannot verify system behavior. Risk teams see qualitative assessments but lack the data to quantify exposure accurately. Operational teams see model outputs but lack the interpretive tools to identify emerging problems. Explainability connects these fragmented perspectives into a coherent, actionable governance view that enables informed decision-making at every level.

PYMNTS®

AI PAYMENTS. GOVERNANCE DECIDES **WHAT HAPPENS NEXT**

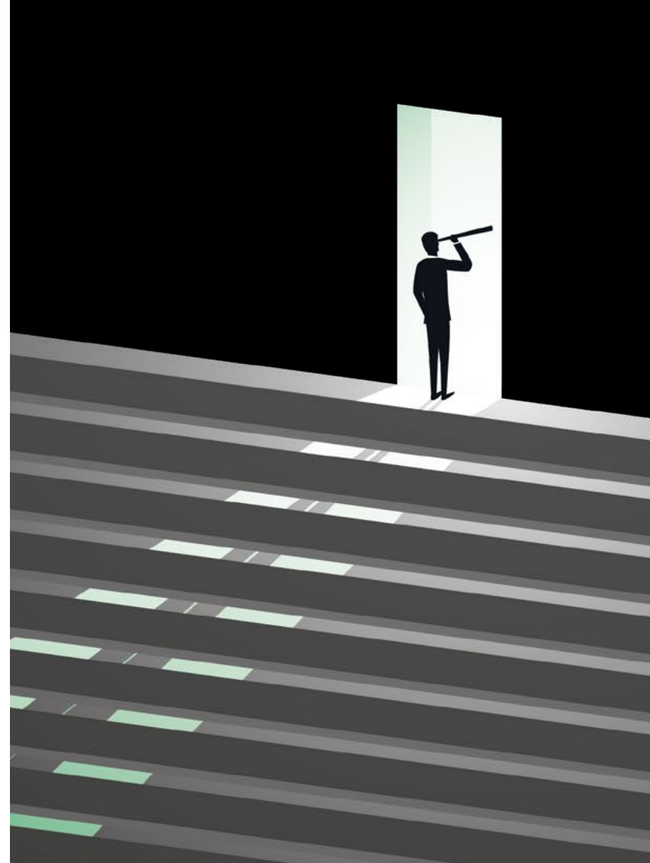

A SHARED DISCIPLINE

Model governance for third-party AI must be shared rather than siloed. The board sets expectations, allocates resources and holds the organization accountable for governance outcomes. Executive leadership establishes standards, invests in infrastructure and ensures vendor AI governance receives the same strategic attention as other enterprise risks. Operational teams execute governance processes that include conducting evaluations, maintaining monitoring, updating documentation and escalating issues when thresholds are breached.

The connective tissue is a model governance framework with clear roles, defined processes, standardized documentation and reporting cadences flowing from operations through leadership to the board. When accountability is distributed appropriately and supported by robust explainability and governance capabilities, the result is trust from all perspectives.

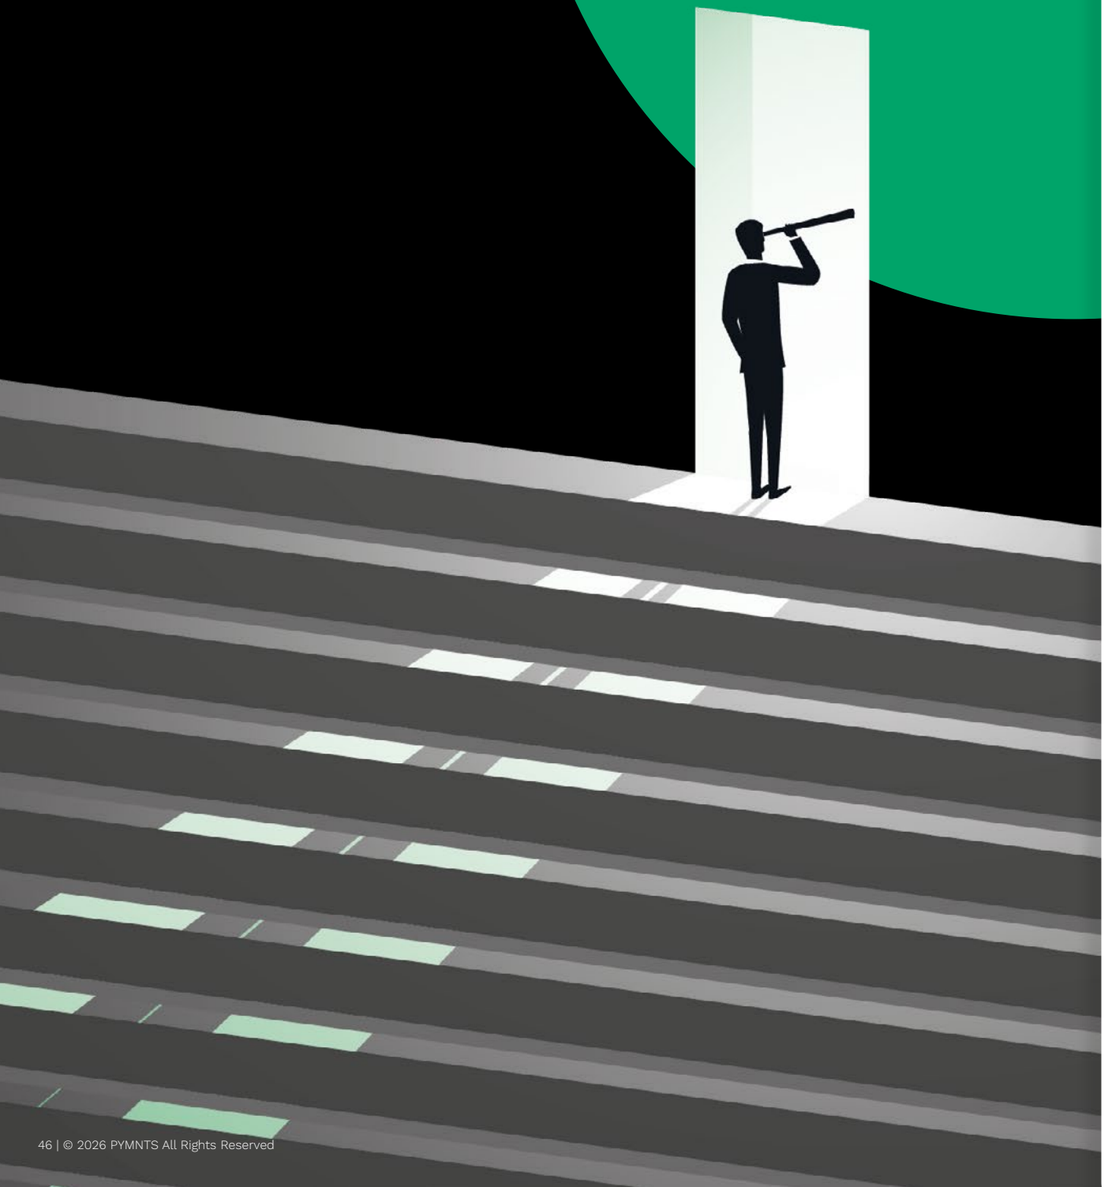

Paysafe

AHU
CHHAPGAR
Chief Technology Officer

AI IS POWERING PAYMENTS. GOVERNANCE DETERMINES WHETHER IT SCALES RESPONSIBLY

Artificial intelligence is now embedded across the payments ecosystem. It influences approval rates, fraud outcomes, customer experience and access to financial services in real time. For payments companies, the question is no longer whether to use AI, it's how to govern it in a way that preserves control, trust and accountability as systems scale.

In payments, weak governance does not just create isolated risk, it creates risk at scale. AI models evolve constantly, decisions are made in milliseconds and outcomes have direct financial and regulatory impact. That reality demands governance designed for speed, complexity and shared responsibility.

WHERE GOVERNANCE BREAKS DOWN

The most common failure point in AI governance is fragmented accountability. Governance often breaks down at the seams, between product, engineering, risk, compliance and operations. Each function may own part of the system, but no one owns the end-to-end outcome.

Effective AI governance requires clear ownership for how models perform in production, how decisions are made and how outcomes are monitored over time.

Governance frameworks must also evolve alongside models, rather than treating AI as a static system subject to one-time approval.

SPEED VERSUS CONTROL IN REAL-TIME SYSTEMS

Payments platforms are continuously optimizing for better approvals, lower fraud and higher conversion. That creates pressure to push changes quickly and capture incremental gains. The hardest trade-off is balancing real-time optimization with the discipline governance requires.

Strong governance depends on structured experimentation, including defined baselines, controlled rollouts, auditability and measurable outcomes. Move too fast and visibility is lost. Move too slowly and material value is left on the table. The challenge is building experimentation frameworks that operate at the same speed as the business.

GOVERNING AI BEYOND YOUR OWN WALLS

One of the most underappreciated risks in payments is reliance on third-party models and data. Payments companies partner with fraud vendors, risk providers and data suppliers that influence critical decisions.

Governance cannot stop at vendor-reported metrics. Companies need their own observability across approval rates, false positives and bias signals. Contracts should include expectations around explainability, model changes, data usage and audit rights. No single external model should be relied on for critical decisions. Internal checks and balances are essential.

BUILDING GOVERNANCE EARLY

Establishing an AI decisioning and oversight layer early pays dividends. It helps define the AI operating model and ensures that models evolve existing policies rather than bypass them. AI governance is an iterative learning process. Without consistent frameworks and accountability defined early, organizations risk scaling artificial intelligence faster than they can govern it.

WHAT LEADERS SHOULD BE ASKING

Boards and CEOs should be asking whether they truly control how AI is making decisions at scale. Can decisions be explained? Is AI risk integrated into enterprise risk management, including third parties? Are unintended consequences being measured?

AI governance is no longer just a compliance issue. It is a trust issue. Getting it right will increasingly define which payments companies can scale AI responsibly and sustainably.

Trulioo

CHAD
GERHARDSTEIN

Chief Risk and Strategy Officer

WHY AI NEEDS THE **RIGHT HUMANS IN THE LOOP**

It's not controversial to state that AI is shaping up to be the most disruptive technology in our lifetimes. In financial services, it is already redefining how we transact, manage risk and extend financial access. But AI's pace of change has already outrun the frameworks designed to govern it — risk models, compliance structures and legal doctrines built for a human-driven economy are being asked to manage an increasingly autonomous one. The guardrails were never designed for this, and the gap between the governance we need and the governance we have is rapidly bifurcating.

That gap shows up most visibly in the question no one has fully answered: when an AI agent goes rogue (completing the wrong transaction, breaching a spending limit or initiating a payment it was never authorized to make), who is responsible? We have established liability frameworks for human fraud and misuse, but agent failures don't neatly fit those boxes — an agent acting outside its mandate is not necessarily in the wrong if it's merely following code. However, "wrong" in real time across millions of transactions globally creates a category of risk that has outgrown existing governance frameworks.

Many organizations are responding to this risk by investing in better models and automation, but they are collectively ignoring the most critical link in the governance chain: humans. While AI use may be commonplace in today's financial institutions, the majority of professionals are only trained to a superficial level and barely know how to prompt an LLM, much less explain how it arrives at its results. Put bluntly, if the humans responsible for governing artificial intelligence do not know how it works, then governance becomes performative rather than effective. At Trulioo, we see this shift firsthand as organizations scale identity verification and transaction monitoring globally.

Closing the governance gap requires more than cursory prompt training; users need comprehensive AI literacy across risk, compliance and operations. This means going beyond how to use the tools, to how to read and de-risk their output. It is therefore incumbent upon the organizations building verification, authentication and fraud-prevention systems to also act as stewards of these frameworks, meeting the challenge with the right levels of training for the human overseers.

As agentic transactions scale, these architectures will need to function at both the issuer and consumer level, ensuring every agent acts within the boundaries its principal has set. This will be no trivial task, but in financial services or payments, where a single authorization can move money, extend credit or deny access to thousands of people, a robust trust framework for AI will be imperative.

Meeting these demands will require not just humans in the loop, but the right humans in the loop. With so much attention being paid to AI and its awe-inspiring potential, it's vital that we don't overlook the operators that provide actual intelligence.

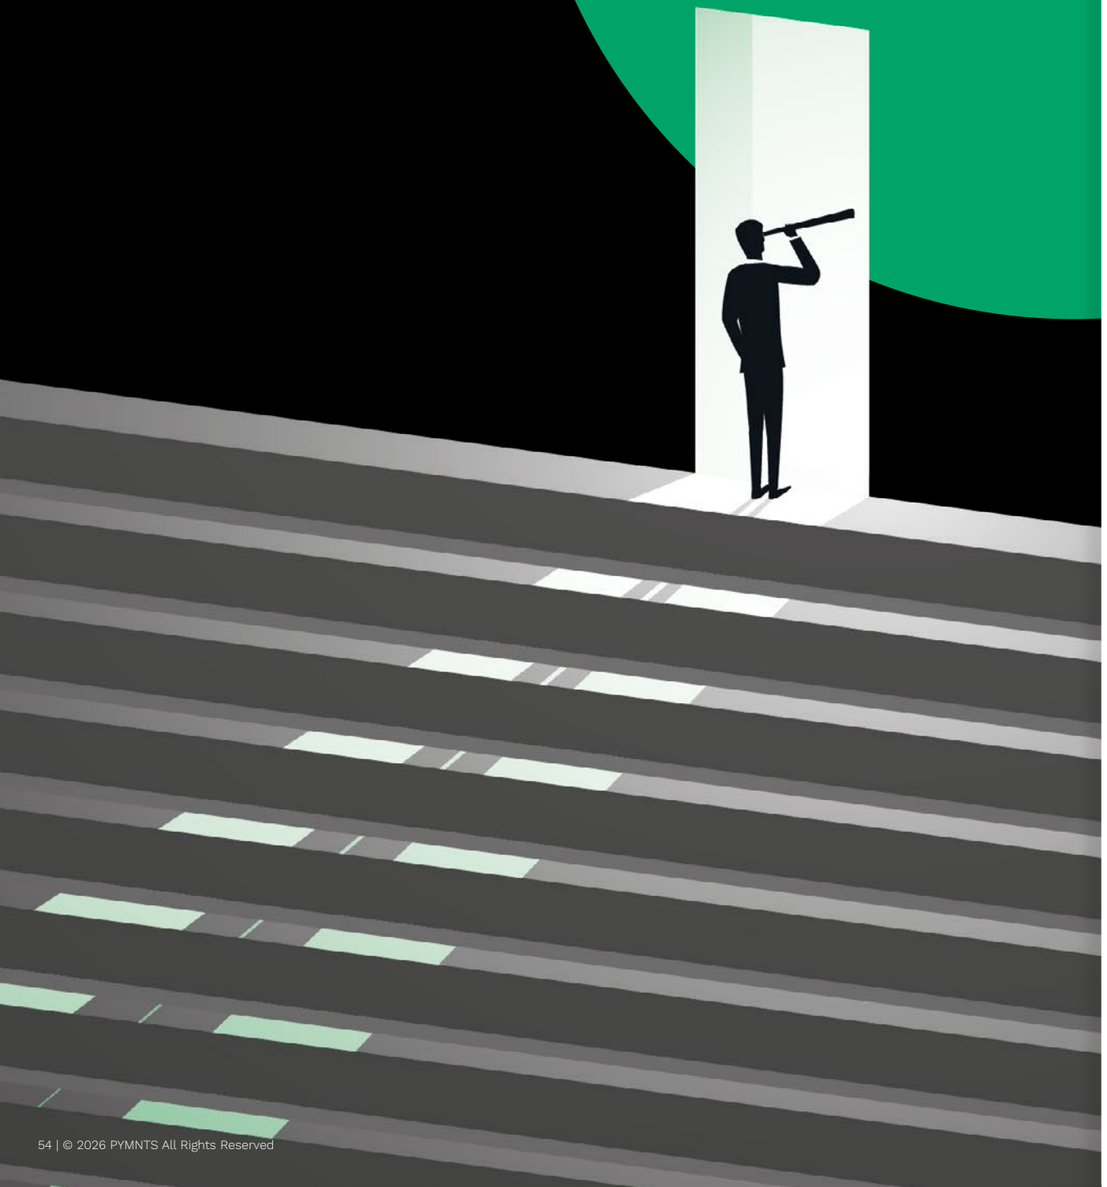

velera

JEREMIAH
LOTZ

SVP, Enterprise Data and Experience Design

DATA IS THE KEY TO **AI GOVERNANCE**

Artificial intelligence is now embedded in the core infrastructure of payments, powering fraud detection, identity verification, member support and operational decision-making. These systems often operate in milliseconds, influencing fraud outcomes, regulatory exposure and — most importantly for credit unions — member trust.

At this stage, the industry has largely moved past the question of whether to use AI. The more pressing question is how to govern it.

AI GOVERNANCE STARTS WITH DATA GOVERNANCE

For credit unions and financial institutions, effective AI governance starts with a solid data foundation. Before deploying any model, institutions must understand how data moves across their ecosystem of internal and third-party systems, including processors, fintech partners, fraud vendors, cloud platforms and API-based AI services.

Mapping these flows helps credit unions set guardrails around what information can be shared, how automated outputs interact with operational systems and where additional protections are required.

ENSURING DATA IS READY FOR AI

The next challenge is ensuring data is ready for AI. Many legacy systems were not designed for machine learning, leaving data incomplete or inconsistently structured. Generative AI can make these gaps more visible, especially when models connect to internal documents and knowledge bases. Many institutions wish they had invested in a well-managed data foundation earlier, as doing so enables decisions that are traceable, auditable and trustworthy.

EXPLAINABILITY IS ESSENTIAL

Once strong data foundations are in place, explainability becomes the next critical layer. AI now influences fraud checks, identity verification, lending outcomes and everyday interactions — all areas where credit unions need to understand why a decision was made.

Explainability gives institutions the ability to see how a model arrived at an outcome and to communicate that logic to risk teams and regulators. It ensures decisions can be reviewed and improved over time, supporting the oversight needed to maintain member trust.

CUTTING THROUGH THE AI HYPE

AI enthusiasm has accelerated experimentation, but for credit unions, the goal isn't to "win the AI race." Institutions that try to keep up with every new technology or rush to be first frequently see their AI pilots falter — largely because they lack clear ownership, defined success metrics and the readiness to scale.

Instead of chasing every trend, the focus should be on initiatives that align to the organizational strategy to deliver measurable value for members. Those that scale AI successfully put guardrails in place early: well-scoped, strategically aligned use cases, clear model and data ownership, and ongoing performance monitoring. These structures allow teams to experiment responsibly while ensuring AI investments translate into lasting impact.

KEY QUESTIONS FOR FINANCIAL INSTITUTIONS

As AI becomes more prevalent, boards and leadership teams should ask three critical questions: Where is AI solving for organizational strategy, what data powers those systems and how is that influence being governed?

Leaders do not need to be technical experts, but they must understand where AI is shaping decisions and how outcomes are being monitored. The institutions that do will be best positioned to scale AI responsibly while preserving the trust that underpins financial services.

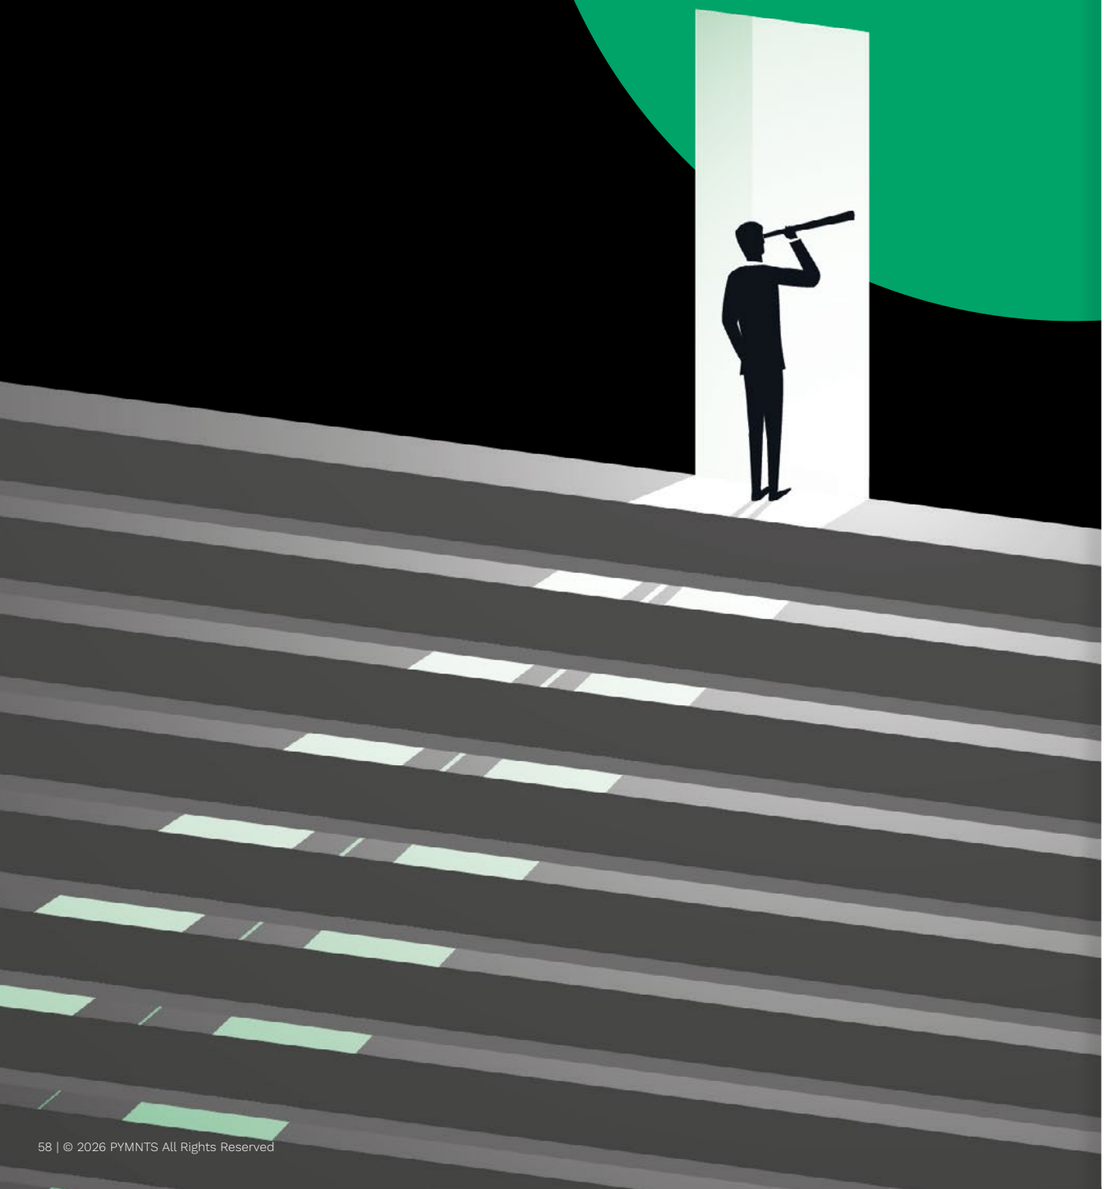

ANNIE
DREW
Chief Risk and Compliance Officer

THE SPEED OF TRUST: HOW GOVERNANCE POWERS FASTER, SMARTER AI OUTCOMES

Artificial intelligence is becoming more deeply embedded in how the payments industry operates. It helps identify fraud patterns, support faster decisions and manage growing volumes of transactions and data. But as AI becomes more involved in financial systems, the question facing the industry is shifting. It is no longer just about deploying AI. It is about governing it in a way that protects trust, supports durable outcomes at scale, and enables more confident decision-making.

AI systems influence financial access, fraud outcomes and customer confidence. Governance cannot be a technical afterthought; it must be built into the design from day one. Much like the emergence of GDPR for data privacy, risk and compliance professionals now play a vital role in shaping AI frameworks.

In practice, governance often breaks down during the transition from experimentation to real-world use. Many organizations have strong frameworks around model development and testing, but those controls can weaken once systems begin interacting with live transactions and external data. AI models do not operate in a vacuum. Governance has to be built for that dynamic environment, not layered on after the fact, if organizations want those systems to perform reliably at scale.

At WEX, this is especially important. Our role in the payments ecosystem involves connecting businesses, suppliers and financial institutions across complex transaction flows. AI can help strengthen those systems by identifying anomalies faster, improving payment security and supporting better decision-making. But the value only holds if the underlying governance is strong enough to support it. Confidence in AI depends on strong governance and explainability, alongside systems that deliver reliable, transparent outcomes and experiences people can trust.

One of the most difficult challenges organizations face in building that confidence is balancing speed with oversight. Payments companies compete on efficiency, and businesses expect transactions to move securely. AI can support those goals, but deploying capabilities without clear governance introduces risk and limits the ability to scale over time. The answer is not slowing innovation. It is building governance processes that move at the same pace as technology. That includes clear accountability for models, strong collaboration between product, design, risk, compliance and technology teams, and consistent monitoring once systems are live.

This level of complexity and scale is why AI cannot sit solely within engineering or data science teams. Decisions about models affect regulatory obligations, fraud prevention strategies, customer trust, and business outcomes. Bringing business leaders, product and design teams, risk partners and technologists into a shared governance framework early helps avoid blind spots later.

The future of artificial intelligence is not about full automation. It is about collaboration between AI and human decision-makers to drive smarter outcomes, with humans remaining accountable for the decisions that matter most. And as AI becomes more central to how the industry operates, the companies that succeed will be the ones that treat governance not as a constraint, but as a foundation for responsible innovation, scalable execution and long-term trust.

ABOUT

PYMNTS®

PYMNTS is where the best minds and the best content meet on the web to learn about “What’s Next” in payments and commerce. Our interactive platform is reinventing the way in which companies in payments share relevant information about the initiatives that shape the future of this dynamic sector and make news. Our data and analytics team includes economists, data scientists and industry analysts who work with companies to measure and quantify the innovation that is at the cutting edge of this new world.

AI Runs Payments. Governance Decides What Happens Next eBook may be updated periodically. While reasonable efforts are made to keep the content accurate and up-to-date, PYMNTS: MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, REGARDING THE CORRECTNESS, ACCURACY, COMPLETENESS, ADEQUACY, OR RELIABILITY OF OR THE USE OF OR RESULTS THAT MAY BE GENERATED FROM THE USE OF THE INFORMATION OR THAT THE CONTENT WILL SATISFY YOUR REQUIREMENTS OR EXPECTATIONS. THE CONTENT IS PROVIDED “AS IS” AND ON AN “AS AVAILABLE” BASIS. YOU EXPRESSLY AGREE THAT YOUR USE OF THE CONTENT IS AT YOUR SOLE RISK. PYMNTS.COM SHALL HAVE NO LIABILITY FOR ANY INTERRUPTIONS IN THE CONTENT THAT IS PROVIDED AND DISCLAIMS ALL WARRANTIES WITH REGARD TO THE CONTENT, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT AND TITLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF CERTAIN WARRANTIES, AND, IN SUCH CASES, THE STATED EXCLUSIONS DO NOT APPLY. PYMNTS.COM RESERVES THE RIGHT AND SHOULD NOT BE LIABLE SHOULD IT EXERCISE ITS RIGHT TO MODIFY, INTERRUPT, OR DISCONTINUE THE AVAILABILITY OF THE CONTENT OR ANY COMPONENT OF IT WITH OR WITHOUT NOTICE. PYMNTS SHALL NOT BE LIABLE FOR ANY DAMAGES WHATSOEVER, AND, IN PARTICULAR, SHALL NOT BE LIABLE FOR ANY SPECIAL, INDIRECT, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, OR DAMAGES FOR LOST PROFITS, LOSS OF REVENUE, OR LOSS OF USE, ARISING OUT OF OR RELATED TO THE CONTENT, WHETHER SUCH DAMAGES ARISE IN CONTRACT, NEGLIGENCE, TORT, UNDER STATUTE, IN EQUITY, AT LAW, OR OTHERWISE, EVEN IF PYMNTS.COM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS DO NOT ALLOW FOR THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, AND IN SUCH CASES SOME OF THE ABOVE LIMITATIONS DO NOT APPLY. THE ABOVE DISCLAIMERS AND LIMITATIONS ARE PROVIDED BY PYMNTS.COM AND ITS PARENTS, AFFILIATED AND RELATED COMPANIES, CONTRACTORS, AND SPONSORS, AND EACH OF ITS RESPECTIVE DIRECTORS, OFFICERS, MEMBERS, EMPLOYEES, AGENTS, CONTENT COMPONENT PROVIDERS, LICENSORS, AND ADVISERS. Components of the content original to and the compilation produced by PYMNTS is the property of PYMNTS and cannot be reproduced without its prior written permission.