EVALUESERVE

● Executive Exchange Summary

Risk Management Under Pressure: Adapting to a New Era of Volatility, Complexity, and AI Disruption

Takeaways from our Executive Exchange in New York City

Participants: BlackRock, BNP Paribas, FIS Global, HSBC, Mizuho, Morgan Stanley, MUFG, Nasdaq, RBC, Schroders, TD

June 2025

Introduction

We're living through a period of profound uncertainty. Some point to rising tariffs and the unraveling of the global order as the hallmarks of this time. Others warn that we're nearing the precipice of artificial general intelligence. And the tumult of ongoing climate change continues to shadow our days. It's not just that things feel strange—baseline assumptions are being dismantled in real-time.

No one understands this better than risk professionals.

Evalueserve's Executive Exchange series returned to New York with a roundtable at Hutong, where we gathered with senior risk leaders and Columbia Business School economics professor Brett House for a closed-door discussion. These are people whose day-to-day revolves around making projections and forecasts. But how does one stress test scenarios that would never have been considered in the first place?

Today's risk functions are being asked to do more, and do it faster, across increasingly complex, interdependent systems.

Technological advances offer new capabilities to meet this demand, but they also introduce new vulnerabilities, such as the possibility of systemic risk arising from generative AI.

"It's very scary" became the refrain of the evening. But not the crippling kind of fear. The sense in the room was that businesses can, and will, continue to manage risks and adapt, tempered by the healthy undercurrent of anxiety you'd expect in any serious conversation about risk.

What emerged was a thoughtful exchange around:

- Where are today's sources of volatility—the visible and less visible?
- How is the role of the risk practitioner evolving as complexity grows? and,
- What does safe, effective adoption of AI look like?

This paper offers a window into that conversation. At a time when many organizations are navigating similar challenges, we hope it contributes to the broader dialogue on how to move forward deliberately, with clear-eyed optimism.

About the Guest Speaker

Brett House is Professor of Professional Practice in the Economics Division at Columbia Business School. His research and writing are focused on macroeconomics and international finance, with interests in fiscal issues, monetary policy, international trade, financial crises, and debt markets.

Previously, Professor House was Deputy Chief Economist at Scotiabank, Chief Economist at a Toronto-based asset management start-up, and Global Strategist at a New York-based global macro hedge fund. In policy-making roles, Professor House was an Economist at the International Monetary Fund for nearly a decade, where he worked on emerging-markets financing. He was also Principal Advisor on Economic Issues in the Office of United Nations Secretary-General during the global financial crisis.

Brett House in
Columbia Business School

I. No Safe Anchors: When Policies Shift, Markets React, and AI Misfires

Volatility is coming from many angles, and much of it feels impossible to anticipate.

"It's like the New York subway system after 8 p.m.," according to Professor House. "You can assign a probability to a breakdancer swinging from the railings or a fantastic plant vendor appearing at 96th Street. But there's always the chance that something completely unexpected will happen out of the blue."

If you're not ready for an adventure, you can always skip the subway. Unfortunately, businesses don't have the luxury of opting out. You can't sit out geopolitics, shifting trade policy, or technological disruption. Institutions must find ways to operate through uncertainty. From planning private equity exits to hiring the next class of junior bankers, leaders are grappling with how to plan for the long term.

The discussion broke down four major sources of volatility that complicate these planning efforts: policy oscillations, reactive markets, unreliable macro data, and AI-induced errors.

1. Policy Oscillations

At the time of the event, the U.S. administration had just doubled tariffs on imported steel and aluminum. Risk models generally aren't designed to predict these kinds of moves.

"There's no grand plan in the White House's moves," noted Professor House. "On the one hand, it's easy for every country to show progress against the inflated tariff barriers ascribed to them on April 2nd. On the other hand, we could be in for a capricious back-and-forth on tariffs for the duration of this administration." The cycle of negotiation and rollback, he suggested, is largely performative. And there's no mathematical model or solid probability distribution for political theater.

But that doesn't mean that Washington's ups and downs can be ignored. Sweeping legislation and sudden policy shifts have very real implications for capital planning, operations, and exposure management.

What risk managers must do is double down, run more differentiated models, and expand scenario planning. As one executive put it, "You really need to be able to look at what-if scenarios across complex happenings, under a primary actor in the middle, presuming things are totally unpredictable."

> What risk managers must do is **double down**, run more **differentiated models**, and **expand scenario planning.**

2. Reactive Markets

Since April 2nd, markets have dipped on tariff announcements, then bounced back when policies were walked back. This cycle of reaction and recovery has repeated with each update, creating a frustrating disconnect for economists.

"The equities market has rebounded," noted Professor House. "But it's clear the bond market doesn't share its hopeful outlook. The dollar doesn't buy it either." In other words, we're not in the clear. Equity prices assume a best-case scenario. But the fundamental data on the American economy, from unemployment to inflation, haven't yet reflected the real effects of recent policy shifts.

The challenge isn't just what's already happened, but what it sets in motion. Second- and third-order effects of recent policies may unfold slowly. Given this lag, Professor House believes midterm elections are unlikely to temper policy decisions. "We don't actually know how far the volatile policymaking will go."

3. Unreliable Macro Data

Adding to the challenge is the growing unreliability of key macroeconomic indicators. Professor House noted that data collection for the Consumer Price Index is being scaled back. This erosion in data quality is likely to extend to other economic indicators, undermining what were once considered relatively accurate anchor points for understanding the economy's ongoing evolution.

This complicates the use of basic benchmarks such as risk-free asset pricing. If publicly provided macro data becomes increasingly noisy or incomplete, some financial institutions may feel compelled to fill the gap themselves. But building this capability is costly, and likely only within reach of the largest players.

4. AI-Induced Errors

Setting aside the degradation of traditional macro data, the rise of AI introduces a new kind of uncertainty. As one risk executive remarked, "It's going to be easier to be wrong, let alone the right numbers are not really right."

As financial models grow more complex and AI becomes embedded across research and analytics workflows, the industry faces a growing risk of compounded errors. Flawed data feeds flawed models, which in turn drive flawed decisions upon flawed decisions.

"Portfolios today are more diverse than they used to be," one risk leader shared. "That helps reduce exposures in some sense, as errors tend to get diversified out. But it also means you have more exposure to more errors in the data."

The complexity of modern portfolios means that even small errors in the tails can have outsized impacts on stress testing. Portfolio managers and traders often lack the deep domain expertise required to catch issues buried in niche, highly specialized areas.

To make things more difficult, spotting something unusual is not the same as knowing it's wrong. The model might have misfired, or it might be surfacing something novel. As the old adage goes, past performance does not indicate future results. The problem is, we don't know what we don't know.

> It's going to be **easier to be wrong,** let alone the right numbers are not really right.

II. Rising Stakes for Risk Management: More Complex, Even More Critical

In stable markets, baseline assumptions often suffice. But in times of heightened uncertainty, portfolio optimization demands deeper insights and more agile modeling. As one leader framed it, the current environment is pushing risk managers into a more strategic role. Business stakeholders are hungry for any signal that might point them in the right direction.

But scaling insights comes with challenges. "There are not just a few scenarios," one practitioner noted. "There are lots of plausible, crazy scenarios that you need to be able to run."

In this increasingly complex environment, risk teams are operating more complex models. There's an opportunity to leverage AI advances to process more data, faster, but this introduces yet another layer of complexity and uncertainty to risk management systems.

Participants identified a number of recurring challenges:

- **Finding Value in the Noise:** In some areas, like equity markets, there's no shortage of data. The challenge is separating signal from noise. In other markets, such as private markets, information is limited, and analysts are expected to reach sound decisions without a full picture faster than the competition. In both cases, the goal is the same—turning imperfect inputs into value-driving decisions.
- **The Burden of Data Maintenance:** Financial institutions often manage thousands of components across sprawling data systems. For many, maintaining these platforms is a growing cost burden, especially for smaller players. Regardless of size, firms face increasing pressure to streamline data upkeep, improve integrations, and drive efficiency in risk-related middle- and back-office operations, and

- **Keeping Pace with Technology:** New capabilities, particularly in the field of AI, are accelerating rapidly. It seems like every day, there's a new, faster, more efficient way to mine insights. One practitioner described it as an arms race, expressing how it would be nice to work with providers who rent out these capabilities. In banking, building everything in-house isn't practical. But they must find ways to stay current and embed new capabilities into their workflows—or risk being left behind.

AI's rapid evolution also means it holds real potential to address many of these challenges: filtering out noise, streamlining backend processes, monitoring market signals, generating new scenarios, even coaching analysts on how to upskill in AI. Some risk officers have even asked whether generative AI could be used to run the scenario analysis process end-to-end. The vision is to surface a broader, more comprehensive range of scenarios, many times faster than traditional methods.

"It's a great idea, but not so easy to do," one of the executives noted. Risk professionals understand this better than most. While AI has promising capabilities, it also introduces its own set of risks. Few are ready to hand over the reins to autonomous AI systems. But in targeted areas, with clear oversight and guardrails, there was broad agreement that the technology can drive meaningful impact.

> There are lots of **plausible, crazy scenarios** that you need to be able to run.

III. Unlocking AI's Potential: Harnessing the Speed, Without Losing the Wheel

The Opportunity for AI

Risk managers work with models every day. At its core, a large language model (LLM) is simply another kind of model. Like any tool, the value of gen AI depends on where and how it's applied. Some use cases offer clear benefits. Others are less relevant or introduce new risks.

Participants highlighted how different capabilities of generative AI are likely to have varying levels of impact across risk management workflows:

- **Signal processing:** Impact here is limited, as market data remains largely numerical. However, generative AI can help scale sentiment analysis, adding a valuable qualitative layer to traditional signals.
- **Code generation:** This is one of the most universally beneficial applications, not just in risk, but across industries. Gen AI helps speed up model development, supports rapid testing, and reduces development time.

- **Research support:** Gen AI is already widely adopted to aid research. It can assist with gathering information for market scans and scenario planning, especially as conditions shift quickly. That said, participants cautioned that AI still struggles with basic errors, and human oversight remains essential; and
- **Agentic orchestration:** This was the area where participants saw the greatest potential. Beyond verbal or textual outputs, generative AI agents can be deployed to orchestrate and automate processes end-to-end.

"Where I'm pushing our team is on process automation—and it has to be a process we understand," one risk leader said. The business benefits are clear: faster hedging, quicker trade execution, and broader fund coverage per portfolio manager. Ultimately, the risk manager needs to ensure the process is well understood and the output can be explained. But if an AI agent can run that process autonomously and reliably, the potential to scale operational efficiencies is enormous.

Start With Use Cases Where Faster Is Better

Some outputs require absolute accuracy. For example, any public or client-facing figure must be precisely correct. But in other areas, being close is good enough. It's better to optimize for speed first and then refine for precision later.

"Particularly in private markets, the margins of error are pretty wide because there's a lot of uncertainty," one risk leader noted. "So your inaccuracy is pretty wide too. If you can filter through it rapidly, it can be quite valuable."

The discussion covered several areas where faster execution can drive real competitive advantage:

- High-frequency trading;
- Commercial real estate (CRE) bidding;
- Responding to RFPs and RFQs rapidly;
- Private credit deal review; and
- Mortgage analysis.

In these cases, we're talking about orders of magnitude faster. One private credit executive suggested, "It used to take 10 hours in a data room to create a summary for one deal. AI could draft it in minutes." This would be a clear example of process transformation.

Many in the room drew parallels to earlier waves of technology disruption. One executive recalled that when they first started working in mortgages, they covered one deal a week. Once everything moved to email, that jumped to 10 deals a week. Now, AI is set to supercharge these workflows even further.

Most agreed it's not a question of if, but when. The gains from AI will be clear. But with many firms still experimenting, the timeline may be longer than the market expects.

Now, AI is set to **supercharge** these workflows even further.

                    evalueserve.com

The Challenge of Trust

The time horizon for AI's ROI may be disappointing, but it's far from the only concern. Risk managers are also looking at the broader risks tied to widespread adoption of LLMs. The most alarming possibility raised was the potential for generative AI to introduce systemic failure, where the entire market is exposed to a single underlying system. In short, the creation of a kernel problem.

As one leader raised, "Is gen AI the new asbestos?"; that is, it's everywhere, you don't know what it does, and it's almost impossible to remove.

It's an unsettling but plausible analogy. Generative AI is powerful and increasingly easy to deploy—even when users don't fully understand how it works. If it saves time and money, businesses are likely going to run with it.

But the risks are real: hallucinations, feedback loops, and hidden dependencies can quietly undermine decision-making at scale. That's why putting guardrails in place matters.

Effective AI governance isn't straightforward and requires careful consideration across several dimensions:

- **Use-Case Specificity:** First and foremost, AI guardrails need to be driven by business needs. Gen AI should be purpose-built with clearly defined boundaries. Like any model, expect things to break down at the edges when it's used outside its intended scope;
- **Fit-for-Purpose Training Data:** The major LLMs in the market are trained on the world's data. So, for any given use case, it's worth asking: Do I trust the world's data to make this decision? In highly specialized areas, the answer is likely no and will require creating a bespoke dataset. This can be quite costly;

- **Multi-Engine Diversification:** By running multiple gen AI models for each process, users can reduce their risks. This diversifies exposure and builds redundancy into the system. But again, it's an expensive undertaking;
- **Human-in-the-Loop:** AI-generated insights should assist, not replace, human decisions. It's important to design opt-in workflows that allow users to decide when and how to bring generative AI into the process, as a calculator, a sounding board, or a tool for creating a first draft. Otherwise, as many leaders in the room warned, you're going to get burned, and
- **Transparent and testable:** AI outputs need to be auditable. If a model suggests a trade, the reasoning behind it should be traceable—something you can compare against how a human analyst would approach the same decision. That means building an ecosystem of standardized workflows, benchmarking protocols, and stress testing as part of the governance process. And if firms don't take the lead, regulators will. As risk professionals know quite well, they're quite good at forcing transparency.

At the end of the day, a model is only as useful as it is trusted. If users can't rely on gen AI to perform consistently, mistakes may end up costing far more than anticipated, undermining the very efficiency gains the technology is meant to deliver. Governance isn't just about control. It's about managing the risks of AI, so businesses can depend on it to drive value.

Is **generative AI** the new **asbestos**?

Conclusion

In an environment shaped by macroeconomic volatility, increasingly complex financial systems, and rapid technological disruption, risk and resilience are not just back-office concerns; they're central to business strategy. Risk professionals may feel the pressure rising, but this moment also presents an opportunity to step into a more strategic role and partner with business leaders in navigating uncertainty.

This discussion showed that in many cases, moving forward will require businesses to focus on the basics. Volatility demands more robust scenario planning and a willingness to challenge assumptions. The rise of AI calls for the same rigor applied to any widely used model: know its limits, question the data, define clear use cases, and build systems for transparency and governance. Ultimately, businesses that build trusted systems, generate sound insights, and adapt with discipline will be best positioned to deliver value in any environment.
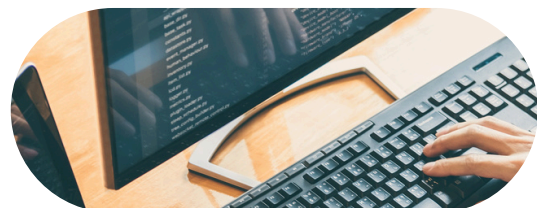
About Our Executive Exchange Series

Evalueserve's Executive Exchange events provide an unparalleled opportunity for senior financial services executives to connect, share insights, and explore innovative ideas in a relaxed setting over dinner and drinks. Our exchanges feature engaging expert-led discussions on trends such as digital transformation and AI.

These intimate, invitation-only events are carefully curated to bring together leaders in banking, lending, and private capital in global financial hubs including New York, Toronto, London, and more.

Learn More