Counterparty Credit Risk · Opinion

Counterparty Risk as a Conversation

CVA numbers exist inside systems that only quants can query. Risk managers and traders work around them rather than with them. AI changes that — but only if the underlying calculations are sufficiently rigorous, transparent and reliable: open-source software and LSEG's Models-as-a-Service architecture make that possible.

By Scott Sobolewski
LSEG Post Trade Solutions
October, 2026



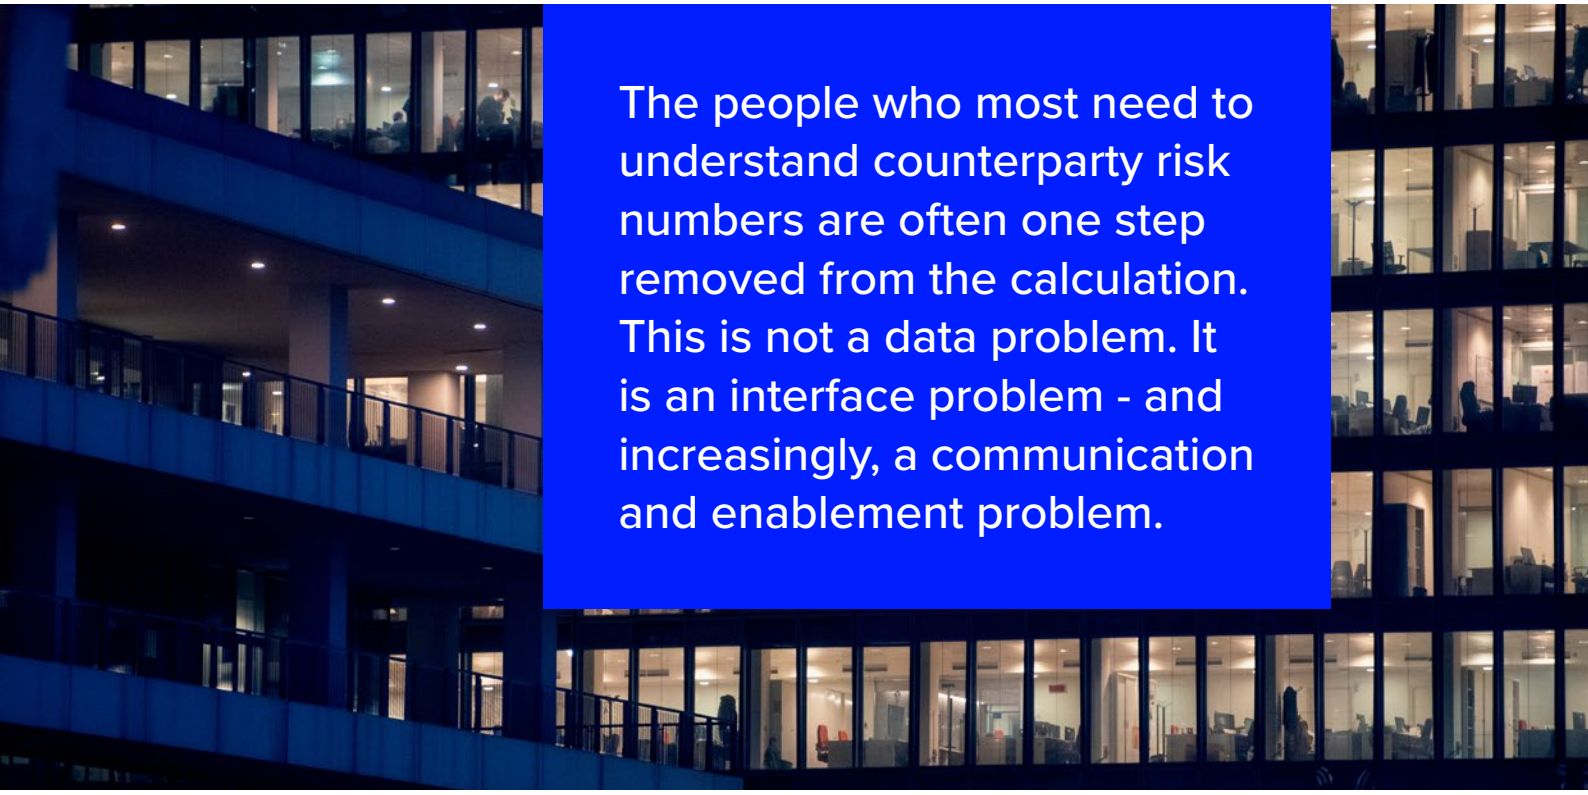
LSEG DATA & ANALYTICS

The gap nobody talks about

Every major derivatives dealer has a CVA desk. Every sophisticated buy-side firm tracks counterparty credit exposure. And yet, in most institutions, the people who most need to understand these numbers —portfolio managers, traders, senior risk officers— interact with them at one step removed from the calculation: through a weekly report, a morning briefing, or a dashboard that shows a static, often stale aggregated figure.

This is not a data problem. The calculations exist, often produced in nightly batches by legacy engines of considerable sophistication. It is a communication problem: the interface between the calculation and the decision-maker is broken. Counterparty risk has a last-mile problem, and it is more consequential than most institutions acknowledge.

Consider what happens when a trader wants to understand the CVA impact of a new trade before execution. In most organisations, the answer necessitates involvement of an overworked quant with direct system access, a proxy calculation that sacrifices accuracy for speed, or —most commonly— a decision based on stale data and/or intuition. None of these are satisfactory, and collectively they represent a significant source of both increased risk and missed optimisation opportunity.



The people who most need to understand counterparty risk numbers are often one step removed from the calculation. This is not a data problem. It is an interface problem - and increasingly, a communication and enablement problem.

Why the interface has not kept up

The sophistication of counterparty risk calculation has advanced enormously over the past fifteen years. Monte Carlo exposure simulation, XVA pricing, dynamic initial margin modelling, and Basel III SA-CCR capital calculations have all matured from academic concepts into production systems. The underlying models are sophisticated and can capture the risk and economic costs of trading with considerable granularity.

The interface to those calculations has not kept pace. Risk systems were built by quants, for quants. They surface their outputs in multi-dimensional report files, cube data structures, and configuration-heavy APIs. For the professionals for whom these systems were nominally built — risk managers, front-office practitioners, senior executives — the practical barrier to access remains high.

Many risk systems remain complex to access and use for non-technical users, even though AI Chat-based interfaces have made these traditional front-ends less relevant than ever.

There are structural reasons this gap has persisted. Building a genuinely usable interface on top of a complex risk engine requires investment that has traditionally been harder to justify than the underlying calculation capability. And the population of users sophisticated enough to interpret the raw outputs directly has been small enough that the gap has not registered as a crisis. It has simply been accepted as the cost of doing business with complex analytics.

The regulatory environment is now making that acceptance untenable. Supervisors expect firms to demonstrate not just that they have accurate XVA calculations, but that the outputs are understood and acted upon at an appropriate level of seniority. Model risk management frameworks require that model users — not just model developers — can explain and challenge the outputs in a meaningful way. The communication gap is becoming a model governance gap; and even though global regulatory pressures appear to have temporarily abated relative to the post-crisis Dodd-Frank era, expectations have shifted towards embedding on-going model monitoring frameworks into daily BAU processes. The US Federal Reserve's recent SR 26-2: Revised Guidance on Model Risk Management reflects the continued focus on model governance and oversight. Similar themes can be observed across a number of regulatory frameworks globally.

What a Risk Manager Should Be Able to Ask — and Currently Cannot

- "What are my top ten CVA concentrations by counterparty today, and how have they moved this week?"
- "If EUR rates rise 50bp, which netting sets see the largest PFE increase?"
- "Which trades are contributing most to my SA-CCR capital charge with counterparty X?"
- "What is the incremental CVA of this proposed trade before I execute?"
- "Which counterparty is most optimal for this proposed trade, relative to both to my funding cost of margin and regulatory capital requirements?"
- "Is my market risk optimally balanced across counterparties? Could I maintain the same net market risk by shifting trades from sub-optimal counterparties to others?"

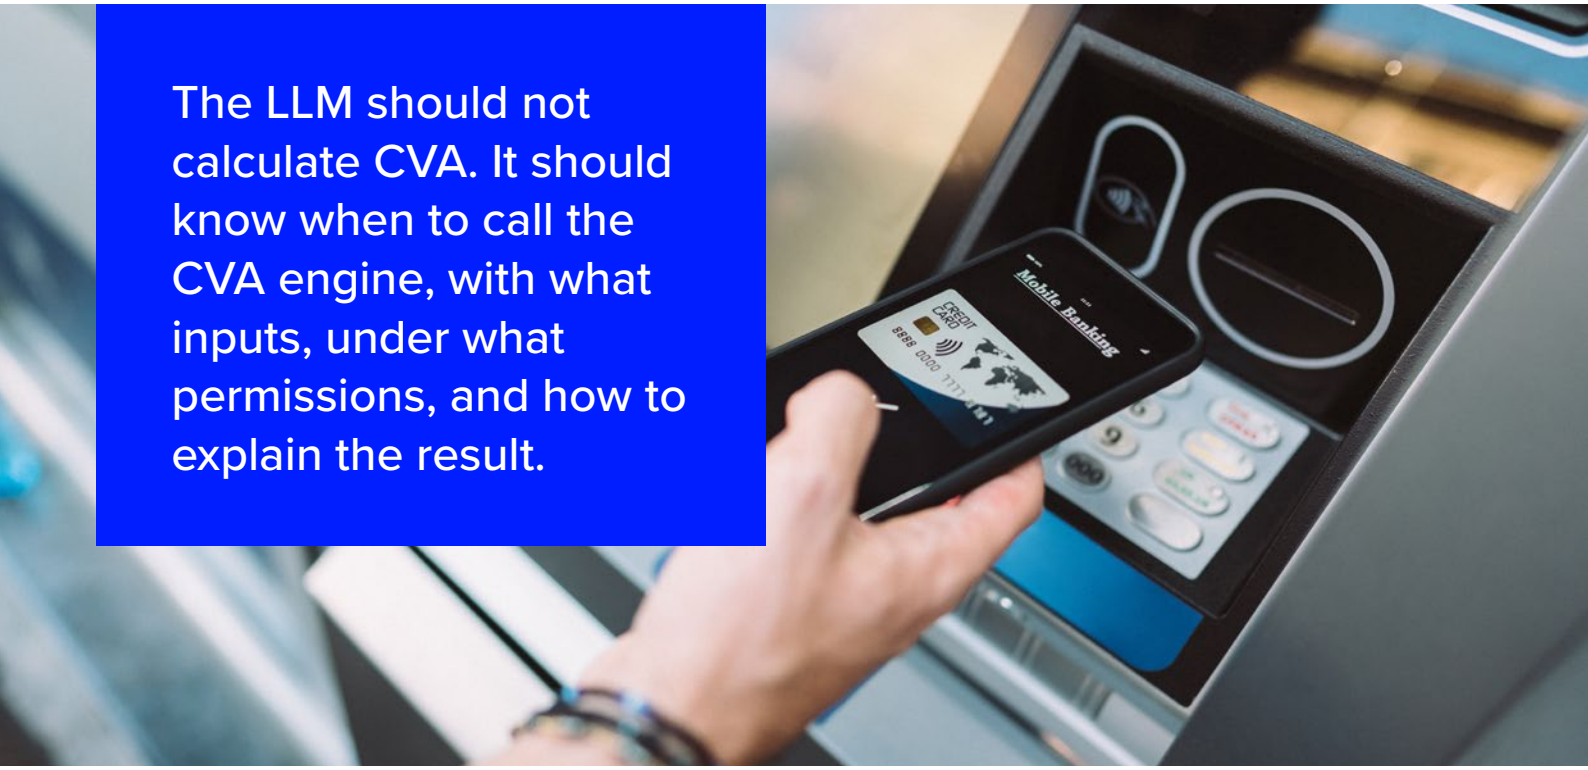
These are not exotic questions. They are the questions that drive risk management decisions daily. In most institutions, answering any of them requires a quant with direct system access.

What AI actually changes

Large language models have attracted a great deal of attention for their ability to generate text, summarise documents, and answer general knowledge questions —to the average user, they are search engine replacements. Those capabilities are useful, but they are not what makes AI transformative for counterparty credit risk. The more important capability is analytics integration and accessibility, which requires significant data plumbing “under the hood”: translating a natural-language question into a sequence of approved analytical operations, assembling the required inputs, invoking the relevant model, and explaining the returned result.

This is qualitatively different from a chatbot that retrieves static information. When a risk manager asks, 'why did our CVA increase by $8 million yesterday?', a useful answer requires decomposition: which counterparties, netting sets, and trades moved; how much came from market-factor changes versus portfolio activity; whether collateral or initial margin changed the exposure profile; and which part of the movement is genuinely material. Can the AI agent initiate a timely re-run of the full risk stack? An LLM connected to live analytical tools can orchestrate that investigation. A static dashboard cannot.

> The LLM should not calculate CVA. It should know when to call the CVA engine, with what inputs, under what permissions, and how to explain the result.



That distinction is central to trusted and reliable AI adoption in Trading Book risk management. The value of AI as a risk interface depends entirely on the quality and timeliness of the underlying calculation. An AI that can articulate a CVA figure confidently but incorrectly is worse than no AI at all —it creates the appearance of understanding without the substance. This is why the combination of a rigorous, open-source calculation engine with an AI interface is not an incidental pairing. It is the essential architecture. The AI provides accessibility; the open-source calculation engine can support model governance and ongoing monitoring by improving transparency and explainability relative to less transparent alternatives.

LSEG Financial Analytics, MCP, and the AI tool layer

LSEG's Models-as-a-Service (MaaS) architecture brings these components together. Open Risk Analytics makes the Open Source Risk Engine's (ORE) transparent, cross-asset models available as a managed service with professionally tuned configurations, integrated market data and scalable compute, accessible through LSEG's Analytics API. The service is designed for familiar developer environments including Python, Jupyter and Visual Studio Code, so firms can use ORE analytics without building and operating the full calculation infrastructure themselves.

Model Context Protocol (MCP) adds the AI layer. MCP gives a compatible AI application the structured layer to discover approved tools, understand their inputs and outputs, pass authorised context, and invoke them. In this setting, an LLM can treat an Open Risk Analytics calculation as a governed tool: identify the appropriate calculation, construct the request from permitted portfolio and CSA context, call the ORE-based calculation service, receive the structured output, and then explain what the result means in both simplified text and navigable charts/tables. The number comes from ORE; the LLM provides orchestration and interpretation.

LSEG's integration with Anthropic's Claude brings the AI conversation layer to a risk calculation stack that is specifically designed to support it. ORE —stewarded by LSEG Post Trade Solutions— provides the underlying counterparty risk calculations: Monte Carlo exposure simulation, full XVA suite (CVA, DVA, FVA, KVA, MVA), PFE profiles, SA-CCR capital, and ISDA SIMM. These calculations are based on models that are used by a range of financial institutions in production environments.

This division of labour is what makes the combination powerful. ORE provides transparent and auditable quantitative logic. Open Risk Analytics provides managed deployment, model configuration, integrated data and scalable execution. The Analytics API makes those capabilities callable. MCP makes them available to AI ecosystems. None of those layers has to pretend to be the others.

Open source raises the ceiling further. Because ORE's methodology, source code and documentation are 100% transparent down to the code-level, an approved AI assistant or coding agent can do more than read an API schema: it can trace model dependencies, help generate integration and test code, and relate outputs back to implementation details. The same transparency appreciated by a human model validator makes ORE more legible to machines. AI makes a rigorous model more usable —and more developable— without making it less rigorous.

Critically, ORE's outputs are structured, decomposable, and attributable in ways that support AI explanation. Because the engine can identify the contribution of individual trades to netting-set exposure, isolate the effect of specific risk factors on CVA, and produce granular sensitivity measures alongside aggregate figures, the AI interface has something substantive to explain —not just a number, but the anatomy of that number.

LSEG's market data infrastructure completes the picture. Real-time interest rate curves, credit spreads, FX rates, and volatility surfaces —delivered natively through LSEG's trusted data platform— mean that the calculations underlying the AI's answers are grounded in the current state of the market, not end-of-day snapshots. When a risk manager asks about CVA exposure at 11am, the answer reflects 11am market conditions.

The conversation still has a last-mile engineering problem

A Claude chat interface looks nice in demos and can be appreciated by non-technical users; making it safe and meaningful in production environments is much harder. Portfolios and CSA data must be mapped, entitlements enforced, model versions controlled, analytics calls constrained and evaluated, and outputs returned into the workflows where decisions are made. This is why enterprise AI is increasingly adopting a Forward-Deployed Engineering (FDE) model: expert engineers work alongside clients at the boundary between a platform and the client's data, controls and operating processes rather than handing over a generic chatbot and leaving the last mile unresolved.

This model is becoming increasingly mainstream and one LSEG has solved for within Open Risk Analytics. AWS launched a $1 billion FDE organisation in 2026 to embed engineers with customers building agentic systems. OpenAI's Deployment Company similarly embeds FDEs to redesign workflows and move frontier models into production. Anthropic is working with DXC to train Claude-certified FDEs for regulated industries, while ServiceNow and Accenture have launched a joint FDE programme. The common lesson is that access to a powerful model is not enough; value depends on engineering it into a governed production workflow.

LSEG Post Trade Solutions provides this quant risk analogue through Expert Services. In combination with an ORA subscription, our engineers can support assessment, reconciliation and benchmarking to other client systems, architecture design, implementation and internal API calls, model validation, and on-going production support —all while approved AI coding agents can work against ORE's open source, tests and documentation to accelerate integration and maintenance.

The FDE emphasis can also shift even further toward orchestration. MaaS and the Analytics API remove much of the engine-hosting burden, so Expert Services can focus on mapping portfolio context, defining entitlements and approved data schemas, exposing the right ORE calculations through MCP, building evaluation and audit controls, and embedding LLM-driven workflows inside a client's existing AI or risk environment. ORE calculates; the LLM orchestrates and explains; LSEG Expert Services expedite making the workflow production-ready.



Practical starting points

Institutions considering AI-augmented risk analytics do not need to solve the entire problem at once. The most effective implementations tend to start with a specific, high-frequency question that currently requires deep quant involvement to answer, and move towards eliminating that dependency.

Pre-trade XVA estimation is a compelling example. The question 'what is the incremental CVA of this proposed trade?' may arise repeatedly on an active derivatives desk. Today, obtaining the answer can require a specialist workflow or a simplified proxy. With Open Risk Analytics exposed as an approved tool through an MCP-enabled AI interface, a user could ask the question conversationally; the LLM could assemble the permitted trade and counterparty context, invoke the calculation, and return the result with the assumptions and key drivers made explicit.

Counterparty concentration and scenario analysis is another natural starting point, particularly for asset managers and regional banks. A risk manager could ask which netting sets drive today's largest PFE, how those concentrations change under a rates or FX shock, or which trades explain a material day-on-day movement. The value is not that AI creates a new risk model, it removes the manual translation layer between the person asking the question and the model already capable of answering it.

The differentiator is controlled accessibility. Questions that previously entered a queue for an overburdened quant team can become direct calls to governed analytics, while model owners retain control over approved calculations, versions, data, parameters and entitlements. Every tool call can be logged, every result can remain reproducible, with all models subject to regular regression testing and monitoring. And because ORE is open source, the institution can inspect the underlying methodology —and use AI to help navigate it— rather than accepting an AI explanation of a “black-box” on faith, which should be particularly compelling to LSEG’s highly-regulated financial services clients.



Conclusion:

The conversation is the interface

Counterparty risk management has both a computation problem and a communication problem. The computation problem is increasingly well-served by mature pricing, exposure and capital models —as long as the compute performance remains scalable. What remains inefficient is the last mile between those models and the people making trading, limit, capital allocation, and hedging decisions.

AI changes the economics of that last mile, but only when it is connected to tools that deserve to be trusted. The combination of open-source ORE, Open Risk Analytics delivered through LSEG's MaaS framework, the LSEG Analytics API, governed market data and MCP-enabled language models creates a clear separation between calculation and conversation. The model remains deterministic and auditable; the interface becomes adaptive and natural-language driven.

That architecture exists right now and can broaden model use without weakening model governance. LSEG helps reduce dependency on scarce specialists for routine questions, and creates a more direct link between calculation and action. Counterparty risk as a conversation is not shorthand for letting an LLM make risk decisions. It is an AI-enabled interface so the language model knows how to ask the right model the right question —and the institution can still inspect and understand how the answer was produced. Institutions that adopt these capabilities may benefit from improved accessibility to risk analytics, more efficient workflows and enhanced governance processes.

About the Author

Scott Sobolewski is Head of Quantitative Services at LSEG Post Trade Solutions. LSEG Post Trade Solutions is the steward of the Open Source Risk Engine (ORE) and provides risk analytics infrastructure to financial institutions globally through the LSEG Risk Analytics platform.

To learn more about counterparty risk analytics via Open Risk Analytics and LSEG's MCP-enabled AI capabilities, contact analyticssales@lseg.com.

Contact us

To learn more, contact us at analyticssales@lseg.com or visit our website: https://www.lseg.com/en/data-analytics/lseg-analytics-solutions

Americas +1 646 989 2200

Tokyo +81 3 6441 1015

EMEA +44 20 7334 8963

APAC +886 2 8729 5130

© LSEG 2026. All rights reserved.

London Stock Exchange Group plc and its Group companies (“LSEG”) do not guarantee that any information contained in thisdocument is and will remain accurate or that use of the information will ensure correct and faultless operation of the relevant service or equipment. LSEG, its agents and employees, accepts no liability for any loss or damage resulting from reliance on the information contained in this document. This document contains information proprietary to LSEG and may not be reproduced, disclosed, or used in whole or part without the express written permission of LSEG. Any software, including but not limited to, the code, screen, structure, sequence, and organization thereof, and documentation are protected by national copyright laws and international treaty provisions. This document is subject to U.S. and other national export regulations. Nothing in this document is intended, nor does it, alter the legal obligations, responsibilities or relationship between yourself and LSEG as set out in the contract existing between us.

LSEG Data & Analytics is part of London Stock Exchange Group plc Group companies.

LSEG DATA & ANALYTICS