



From sharper insights to structural edge

Why AI-native decision-making will
define winners in commodity markets

About the authors



Miguel Torreira

**Senior Managing Director -
CEO Advisory,
Commodity Markets
Global Lead**

Miguel Torreira is Accenture's global lead for Commodity Markets, advising leading trading organizations on strategy, growth and large-scale reinvention. Miguel brings deep expertise in trading, investment and optimization, helping clients unlock value through AI-enabled decision-making and innovative operating models. He works closely with C-suite leaders to transform how trading organizations generate alpha in increasingly volatile markets.



Ogan Kose

**Senior Managing Director -
CEO Advisory,
Reinvention Executive Advisory**


Ogan Kose serves as a trusted advisor to C-suite executives on enterprise-wide transformation, strategy and value creation. Ogan brings extensive experience in operating model design, large-scale reinvention and the application of advanced analytics and AI to complex business challenges. He partners with senior leaders to navigate disruption, accelerate growth and build future-ready organizations.



Nikiforos Atsikpasis

**Managing Director -
CEO Advisory,
Commodity Markets
EMEA Lead**


Nikiforos Atsikpasis is the EMEA lead for Accenture's Commodity Markets, specializing in the intersection of trading, strategy, and technology. He advises global trading organizations modernize decision-making and scale AI, analytics and technology across the trading lifecycle, driving measurable outcomes.



Lydia Karagianni

**Senior Manager -
CEO Advisory,
Commodity Markets**

Lydia Karagianni works closely with clients to translate complex data and market signals into actionable insights and deploy AI-enabled solutions across trading functions. She leads the delivery of these capabilities, helping organizations improve decision quality, execution discipline and operational performance in increasingly complex and volatile market environments.



Jia Liu

**Senior Manager -
CEO Advisory,
Commodity Markets
QuantAI lead**

Jia leads QuantAI, building the algorithmic and agentic AI capabilities behind the next generation of commodity trading decision support. Jia's work turns frontier AI research into backtested signals, autonomous agent systems and decision-grade tools that can be deployed into real trading workflows. He focuses on measurable edge, helping organizations improve forecasting, execution quality and decision speed in increasingly volatile markets.

A structural shift in how profits are made

Commodity trading has always rewarded speed, judgment and risk appetite. But the basis of advantage is changing. Markets now move faster than traditional organizational models can absorb. Volatility spikes more frequently.¹ Liquidity fragments. Geopolitical shocks transmit instantly across physical and financial flows. Renewables reshape intraday patterns. Signal density has exploded.

Advantage historically came from superior access: stronger relationships, better market color, deeper logistics insight. Today, the constraint is no longer access to information but the ability to interpret, prioritize and act on it faster than competitors.

AI is not just improving decisions. It is enabling a new system for making them. We call this the commodity decision engine. A system that ingests signals across markets, operations and external data; normalizes and enriches them into actionable intelligence; translates them into coordinated commercial decisions; executes across trading, procurement and operations; and learns continuously from outcomes.

Trading is evolving from periodic strategy deployment into a dynamic system of hypotheses, testing and adaptation: from optimizing individual decisions to orchestrating them across the entire value chain.

The shift is structural and it directly determines who captures profit. Over the next decade, commodity trading will separate into two groups: organizations constrained by static models, human-heavy workflows and episodic optimization, versus those operating continuously learning, AI-augmented trading systems.

The performance gap between these two groups will widen. In this environment, AI will not be an incremental enhancement but the execution engine of profitability itself.

The market that moves the world

Commodity trading is foundational to the global economy. The physical market alone—spanning oil, gas, metals, power and essential food staples such as rice and grains—represents approximately \$7 trillion annually.² These flows sustain energy systems, enable industrial production, secure supply chains and support food security at scale. Without efficient commodity markets, modern economies cannot function reliably or competitively.

Beyond the physical layer sits a significantly larger financial layer: derivatives markets several times the size of the underlying flows. This “paper” market is not speculative excess, but the pricing and risk infrastructure that enables capital allocation, investment confidence and operational resilience.

Together, the physical and financial markets form a vast, systemically critical ecosystem that underpins growth, stability and global economic security.

Contents

05

**How value gets
created and measured**

16

**Moving to
self-learning systems**

27

**Five actions to build
structural edge**

10

**Treating data as a
P&L foundation**

20

**Breaking through
to scale**

32


**The reinvention
window**

12

**Making AI the execution
engine for decisions**

25

**Second-order risks
CEOs must weigh**

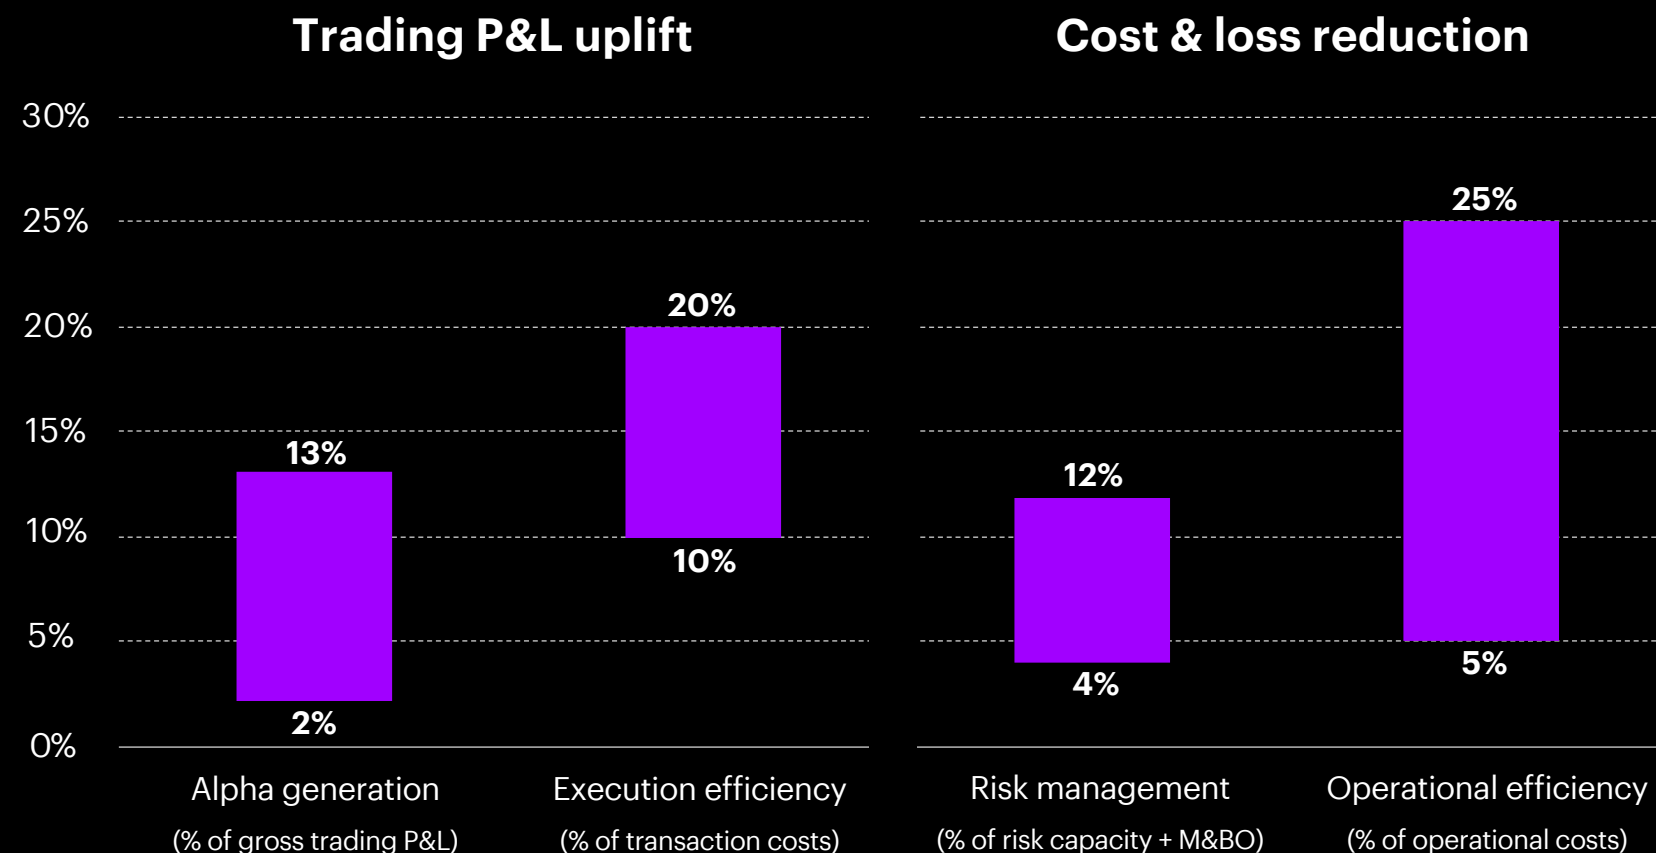


**How value gets
created and measured**

As trading executives spend more on AI, investments need to translate into visible, defensible business impact.

In commodity trading, the impact of AI consistently shows up in four areas: alpha generation, execution efficiency, improved risk management and operational efficiency. Each carries clear, quantifiable signals that separate genuine transformation from superficial experimentation (see Figure 1).

Figure 1: Impact of AI in commodity trading



Source: Accenture analysis of companies across sectors, 2026. Note: Averages have been used across commodity benches including Accenture proprietary benchmarks and publicly available figures. Each pool measured against its own denominator; bars are not additive. The typical values per trader will vary depending on the maturity of existing capabilities.

1. Alpha generation

AI increases alpha by improving how trading and commercial decisions are made, particularly under uncertainty and speed. By detecting patterns across weather, freight flows, regional markets, operational updates and sentiment, AI enables traders to identify opportunities traditional analysis may miss. The impact is twofold: improved directional signal accuracy on trades and better recognition of embedded optionality across the value chain.

Accenture analysis suggests **AI can deliver a 2–13% uplift in gross trading profit and loss (P&L)**, with the range reflecting the gap between mid-tier and leading implementations depending on data maturity, implementation quality, commodity mix and level of integration across the value chain. The direction is consistent: the ability of AI to extract structure from data that was previously too noisy or too vast for traditional analysis can translate into measurable P&L improvement.

2. Execution efficiency

AI optimizes how a trader gets into and out of positions that have already been decided—reducing market impact, slippage, bid-ask spread paid and timing cost. This is distinct from alpha generation: it does not affect which trades are taken, only how efficiently they are executed. In physical commodity markets, this extends beyond tighter execution on a futures screen. AI can select the optimal vessel and route based on live freight rates, port congestion and weather;

predict demurrage risk and reschedules cargoes to avoid waiting time; and model the cheapest blend of grades that meets contractual quality specifications, preventing reblending claims and quality rejections at destination.

For a \$10B physical book, **AI can reduce all-in transaction costs by 10–20%**, driven primarily by execution timing, freight and logistics optimization, demurrage reduction and grade blending. The gains compound across trades: small improvements in slippage, routing and quality become a persistent edge in margin per cargo.

3. Risk management

Risk management is often where the impact of AI becomes visible first, but its real value lies in how it reshapes risk decisions. Traditionally, risk has operated as a control layer: monitoring exposures, enforcing limits and intervening after positions are taken. In fast-moving commodity markets, this model is no longer sufficient.

AI shifts risk management from reactive oversight to real-time decision support. With continuous exposure visibility, dynamically optimized hedges and automated limit monitoring, companies no longer assess risk at discrete checkpoints but actively shape it as decisions are made. Operational errors that create risk exposure—limit breaches, position reconciliation breaks—can be caught earlier, reducing the likelihood that small errors escalate into material losses. Accenture estimates that **AI can deliver a 4–12% reduction** against the combined risk capacity

and middle- and back-office cost base, concentrated in hedging effectiveness, limit monitoring and dispute resolution. The result is a more predictable risk profile, fewer end-of-day surprises and greater confidence in how risk and return interact.

4. Operational efficiency

Operational efficiency has long been a focus in commodity trading, with firms investing heavily in straight-through processing (STP) to reduce manual intervention. Traditional automation and RPA improved standardized workflows but struggled with the variability, exceptions and unstructured data that characterize real trading environments.

AI, particularly agent-based systems, extends automation into these more complex scenarios. Rather than following predefined rules, AI can interpret contracts, reconcile inconsistencies and manage exceptions dynamically, enabling higher and more resilient STP rates even in fragmented or bespoke workflows. By reducing manual intervention in confirmations, settlements and reconciliation, firms can lower error rates, shorten cycle times and improve data quality at source.

AI can reduce the operational cost base by 5–25%, with reconciliation and trade lifecycle automation accounting for roughly 75% of the potential impact. The objective is not headcount reduction but capacity reallocation: freeing experienced professionals to focus on judgment-intensive activities such as counterparty risk assessment, deal structuring and market positioning.

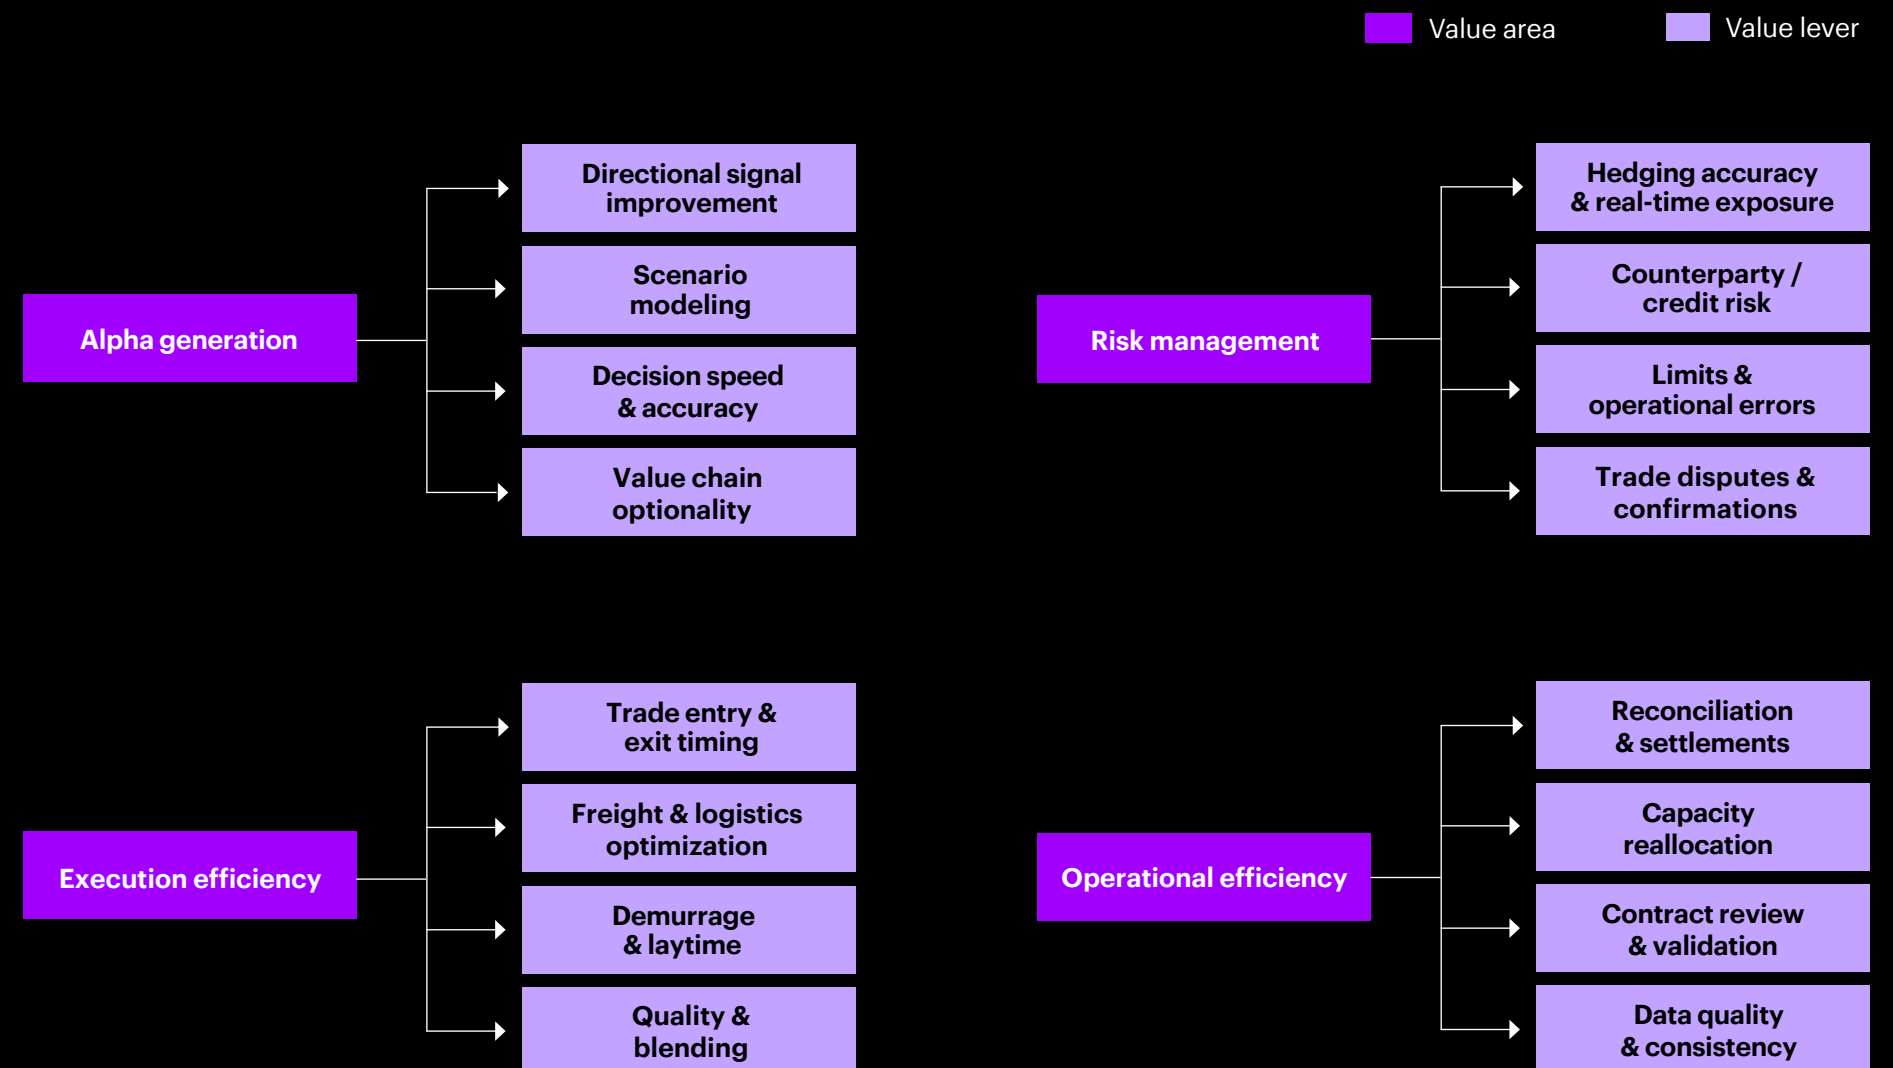
The objective is not headcount reduction but capacity reallocation: freeing experienced professionals to focus on judgment-intensive activities such as counterparty risk assessment, deal structuring and market positioning. Firms measure impact through end-to-end STP rates, trade-processing cycle times and the proportion of exceptions resolved autonomously versus requiring manual intervention.

The unifying theme? Measurability.

Improvements in alpha generation, execution efficiency, risk management and operational efficiency each map to specific metrics, from P&L attribution and latency reduction to hedge accuracy, exception frequency and cycle times (see Figure 2).

This measurability gives CEOs a practical way to track impact, double down on what works and scale AI initiatives with confidence. But reaching repeatable, enterprise-level results requires a structural shift in the trading operating model built on three core capabilities: **data as the P&L foundation, AI as the execution engine** and **continuous learning** as the mechanism that sustains advantage.

Figure 2: Making AI impact measurable in commodity trading



From electrons to insights

In utilities, resilience is becoming the defining challenge. As systems grow more complex, the industry faces a delicate balance between affordability, security and stability. This shift is nowhere more visible than in the rise of renewables.

There is strong pressure to deliver cheaper energy, often through renewable generation. Yet, as Konstantinos Nazos, Chief Energy Management & Trading Officer and Deputy CEO for Public Power Corporation (PPC), puts it, this creates a tension: “It introduces volatility. Supply becomes less predictable, price signals more erratic and system stability harder to maintain.”

The result is a mismatch between perception and reality. “It creates a false impression of abundant energy,” Mr. Nazos warns. Beneath the surface, critical infrastructure, particularly networks, faces mounting strain. “At the end of the day, we are transferring electrons, and the network is the most critical part.”

Energy may appear plentiful, but resilience is increasingly fragile.

To manage this complexity, utility companies are turning to new technologies. AI is emerging as a central component. “My ambition is to use this technology to build a new business model,” Mr. Nazos says. “We no longer need the same investment or time to build capabilities,” nor must firms “repeat the journeys others have already taken.” Trading capabilities are becoming more accessible, challenging incumbents.

As in commodity markets more broadly, the source of advantage for utility providers is shifting. AI compresses the time between signal and decision, enabling faster, more precise responses.

"I don't see this just as a technology shift, but also a leadership challenge."

— Konstantinos Nazos, Chief Energy Management & Trading Officer and Deputy CEO, PPC

Yet technology alone is not enough. “There’s also a personal challenge involved. I don’t see this just as a technology shift, but also as a leadership challenge. One that requires us to learn how to ask better questions. For example, how do we expand this knowledge across our team? We need to switch from finding solutions to asking the right questions,” he says. Systems may analyze data, but professionals must interpret outputs and decide how to act.

That requires a different mindset. “We need to develop systems thinking,” Mr. Nazos explains, “moving from detail to a broader, first-principles view.” This also blurs boundaries. “We need to see ourselves as having one foot in the industry and one in technology.”

For utilities, the implications are vast. They operate at the intersection of infrastructure and markets, where decisions must account for both price signals and system security.

As data and AI reshape decision-making, advantage will belong to those who can manage both energy and information while keeping the system stable.



**Treating data as
a P&L foundation**

Modern commodity markets reward organizations that detect, interpret and act on change faster than competitors. This ability to make better decisions depends on data quality, latency, breadth and governance.

Historically, trading firms tolerated fragmented data landscapes because human teams could still synthesize key insights. That era has ended.

Today's trading environment produces vast volumes of interconnected signals, ranging from weather anomalies and satellite imagery to vessel behavior, refining outages, policy announcements, supply and demand movements, load curves and emissions updates.

AI, however, changes what is possible with data. Large language models (LLM) can read and interpret unstructured documents, such as contracts, regulatory filings, operational bulletins and earnings transcripts, at a scale that was previously unthinkable.

Beyond text, AI frameworks extend to visual and sensor data—critical for commodity trading, where physical flows and logistics drive price formation. Computer vision models convert satellite imagery and port-activity footage into structured insights. Multimodal agents evaluate these inputs in context, cross-checking them against liquidity, margin impact, logistics constraints or hedging commitments. In this model, AI becomes the orchestrator of “source-to-signal” workflows that ensure data hitting the desk is clean, timely and relevant.

Treating data as P&L infrastructure requires discipline. Leading organizations now enforce latency targets, lineage tracking, schema validation and anomaly detection before any signal reaches downstream applications. Rather than relying on the data to be perfect, they build systems that can work with messier data, leveraging the use of synthetic data for completeness, while still enforcing quality thresholds by establishing clear ownership, measurable service level agreements,


accountability for data quality and a governance model that ensures AI systems never operate on unverified or stale inputs. The reward is earlier visibility into dislocations such as a pending outage, a shift in cross-border flows or a weather-driven imbalance that may only last minutes.

The commercial benefits are direct. AI-enhanced data improves trading decisions, reveals optionality across the value chain, identifies new arbitrage patterns and clarifies how physical and financial conditions are evolving from intraday to macro timelines. Most importantly, it gives front-office teams confidence that the intelligence behind critical decisions is both timely and governed, allowing them to move faster without increasing risk.

For leadership, the shift is philosophical as much as technical.

It entails treating data as a strategic asset where quality directly determines competitive advantage.

Making AI the execution engine for decisions



If data defines what a firm can see, execution defines what it can capture. In commodity trading, advantages of execution quality compound. Small improvements in slippage, timing, routing or hedge performance add up significantly across trades, creating a persistent edge.

AI elevates how teams execute decisions across the entire lifecycle—front, middle and back office—improving speed, accuracy and resilience.

Front office

In the front office, AI-supported pre-trade analysis is reshaping how trading decisions are prepared and executed. Rather than relying on static heuristics, AI systems continuously scan alternative data, market signals and operational inputs to identify opportunities, assess conditions and propose actions aligned with cost and risk appetite.

AI models can detect patterns across structured and unstructured data that are difficult to capture through manual analysis, providing traders with a forward-looking view of price distributions and market dynamics.

Examples:

In the UK electricity market, Accenture developed an execution-support engine using time-series foundation models with lightweight Low-Rank Adaptation (LoRA), generating 30-minute-ahead probabilistic forecasts 48 times per day.

In out-of-sample testing, the model achieved ~62% directional accuracy vs. 56–57.5% in comparable benchmarks—translating improved foresight into tighter bid/offer placement, hedge timing and dispatch decisions.

Using our fast adaptive scenario testing methodology, we developed a conflict impact simulation capability modeling hundreds of geopolitical scenarios and their effects on oil and gas prices, logistics constraints and portfolio valuation.

By combining multi-agent simulation, forecasting and Monte Carlo techniques, the system provided a probabilistic view of outcomes, allowing traders to evaluate strategies across a wide range of conditions.

We combined signal processing, gen AI and agentic simulation techniques to analyze earnings transcript data. The prototype was designed to anticipate quarterly equity performance by extracting predictive sentiment signals from management tone, linguistic patterns and narrative consistency.

Back-tested over 11 years against Accenture's own stock, the model achieved a Sharpe ratio of 2.07, comparable to global top-20 quant hedge funds. Notably, it outperformed a buy-and-hold S&P 500 strategy by 3.69x over the same period.

This illustrates how AI can surface alpha hidden in unstructured text, a capability directly relevant to commodity trading, where market-moving information is often buried in non-standard formats.

AI systems can monitor large numbers of signals simultaneously, identify relationships and refine strategies dynamically—something that quickly exceeds human capacity. Reinforcement learning-based optimizers can enhance execution and portfolio allocation by continuously selecting and adapting strategies within defined risk guardrails. In practice, these systems improve execution discipline by monitoring positions, market conditions and contract exposures in real time, raising alerts or adjusting actions when execution quality deteriorates. Rather than replacing traders, they augment decision-making where speed and complexity are critical.

AI enables scenario-based decisioning at scale. Instead of relying on a limited set of predefined scenarios, firms can simulate hundreds of potential market developments and assess their impact on positions, logistics and portfolio value in near real time.

Across these capabilities, execution becomes more disciplined and adaptive. AI agents continuously monitor markets against internal positions and constraints, responding to anomalies, enforcing limits and refining actions as conditions evolve. The result is not just faster execution, but better-informed and more resilient decision-making at the point of action.

Middle office

In the middle office, AI replaces manual, reactive oversight with real-time intelligence. Margin requirements, collateral availability, credit line utilization and limit adherence are updated continuously rather than at predetermined checkpoints. P&L explanations become more granular: firms can separate model-driven gains from execution-driven ones and understand how both controllable and uncontrollable factors contributed to P&L movements. This helps chief risk officers, chief financial officers and heads of trading better anticipate stress, improves forecasting and reduces end-of-day surprises that undermine confidence.

Example:

Accenture developed an exhaustive document consistency engine that detects contradictions across large trading and regulatory documents such as master agreements, trade confirmations and compliance filings, where manual review is labor intensive and costly. The system compares every sentence pair in a document (a 500-sentence filing generates 124,750 comparisons), scores each pair with a deterministic natural-language inference model, applies algorithmic change point detection to surface only genuinely anomalous candidates and validates the highest-scoring pairs through an LLM-as-judge layer.

The result is a fully auditable pipeline that processes in roughly 90 minutes what would take a human reviewer over 1,000 hours, with reproducible, deterministic findings that stand up to regulatory scrutiny.

For commodity trading middle offices, this means clause mismatches between master agreements and confirmations, inconsistent risk assessments across filing sections or contradictory limit definitions can be surfaced automatically and systematically rather than discovered during audits or disputes.

Back office

Post-trade operations also benefit. AI systems automate contract parsing, confirmation matching, settlements, reconciliation and invoice validation. These are tasks that traditionally consumed large amounts of manual effort and introduced operational risk.

By automating these steps, firms reduce lifecycle bottlenecks and create cleaner, faster workflows that support higher trading velocity.

Example:

One trader deployed a multi-agent AI solution that automates repetitive trading tasks with greater speed and precision than human teams. The system reconciles positions, monitors settlements and invoicing, validates compliance and extracts insights from large volumes of trade documentation.

By standardizing workflows and using AI to handle communications, flag workflows and trigger automation widgets incrementally, the client has reduced errors and costs while releasing 27% capacity and enabling skilled talent to focus on higher-value risk and commercial decisions.

The pattern is consistent across commodity groups: human operators maintain oversight, ensuring automated actions align with compliance rules, exposure limits or market-design requirements. This hybrid model illustrates where the industry is heading.

Execution uplift is one of the most reliable and measurable AI value pools.

When aligned with proper controls, AI-driven execution improves both P&L and risk management, and often does so faster than any other type of digital investment.



**Moving to
self-learning systems**

Many trading organizations still rely on static, manual or infrequently updated models for decision-making. This is a structural disadvantage in volatile commodity markets. Conditions evolve constantly: weather patterns shift, trade flows reroute, geopolitical triggers reshape supply chains and renewables introduce new forms of generation variability.

Competitive advantage now belongs to firms that can learn continuously; where AI agents generate, test, deploy and retire signals in rapid cycles.

Example:

In a month-long, walk-forward pilot, two classic S&P 500 signals, RSI divergence and a 20/50-day moving average crossover, were traded daily, yielding cumulative losses of 20–30%.

When managed by a reinforcement learning (RL) selector that treated each signal as an RL agent and chose daily which to activate based on next-day profitability, cumulative returns shifted to +10–15% without changing position sizing.

The RL layer introduced no new signals but optimized real-time capital allocation, turning a losing baseline into a winning outcome through effective trading strategy orchestration.

A self-learning architecture enables this gain. It connects agents through shared memory and unified policy frameworks. Signals move through an automated pipeline in which they are designed, stress-tested, validated and deployed with speed. Poor performers are quickly culled while strong performers scale. Rather than waiting months to operationalize new insights, firms can deploy improvements in days.

Example:

In one simulation, a continuous-learning anomaly detection system aimed to reduce false alerts in hourly energy-load monitoring.

An inner loop scored each building's load with a windowed k-Nearest Neighbor detector (7-day context, stride=1) and applied an unsupervised 99.5th-percentile threshold that recalibrated as new data arrived. An outer loop ("Evolutionary AI") used LLM-assisted evolutionary search to propose code-level improvements, re-run walk-forward evaluation and retain only variants that improve precision/recall.

In one evolution session (100 generations), the mean F1 score, a metric balancing detection accuracy and false-alarm control, rose 4x, from ~0.15 to ~0.64.

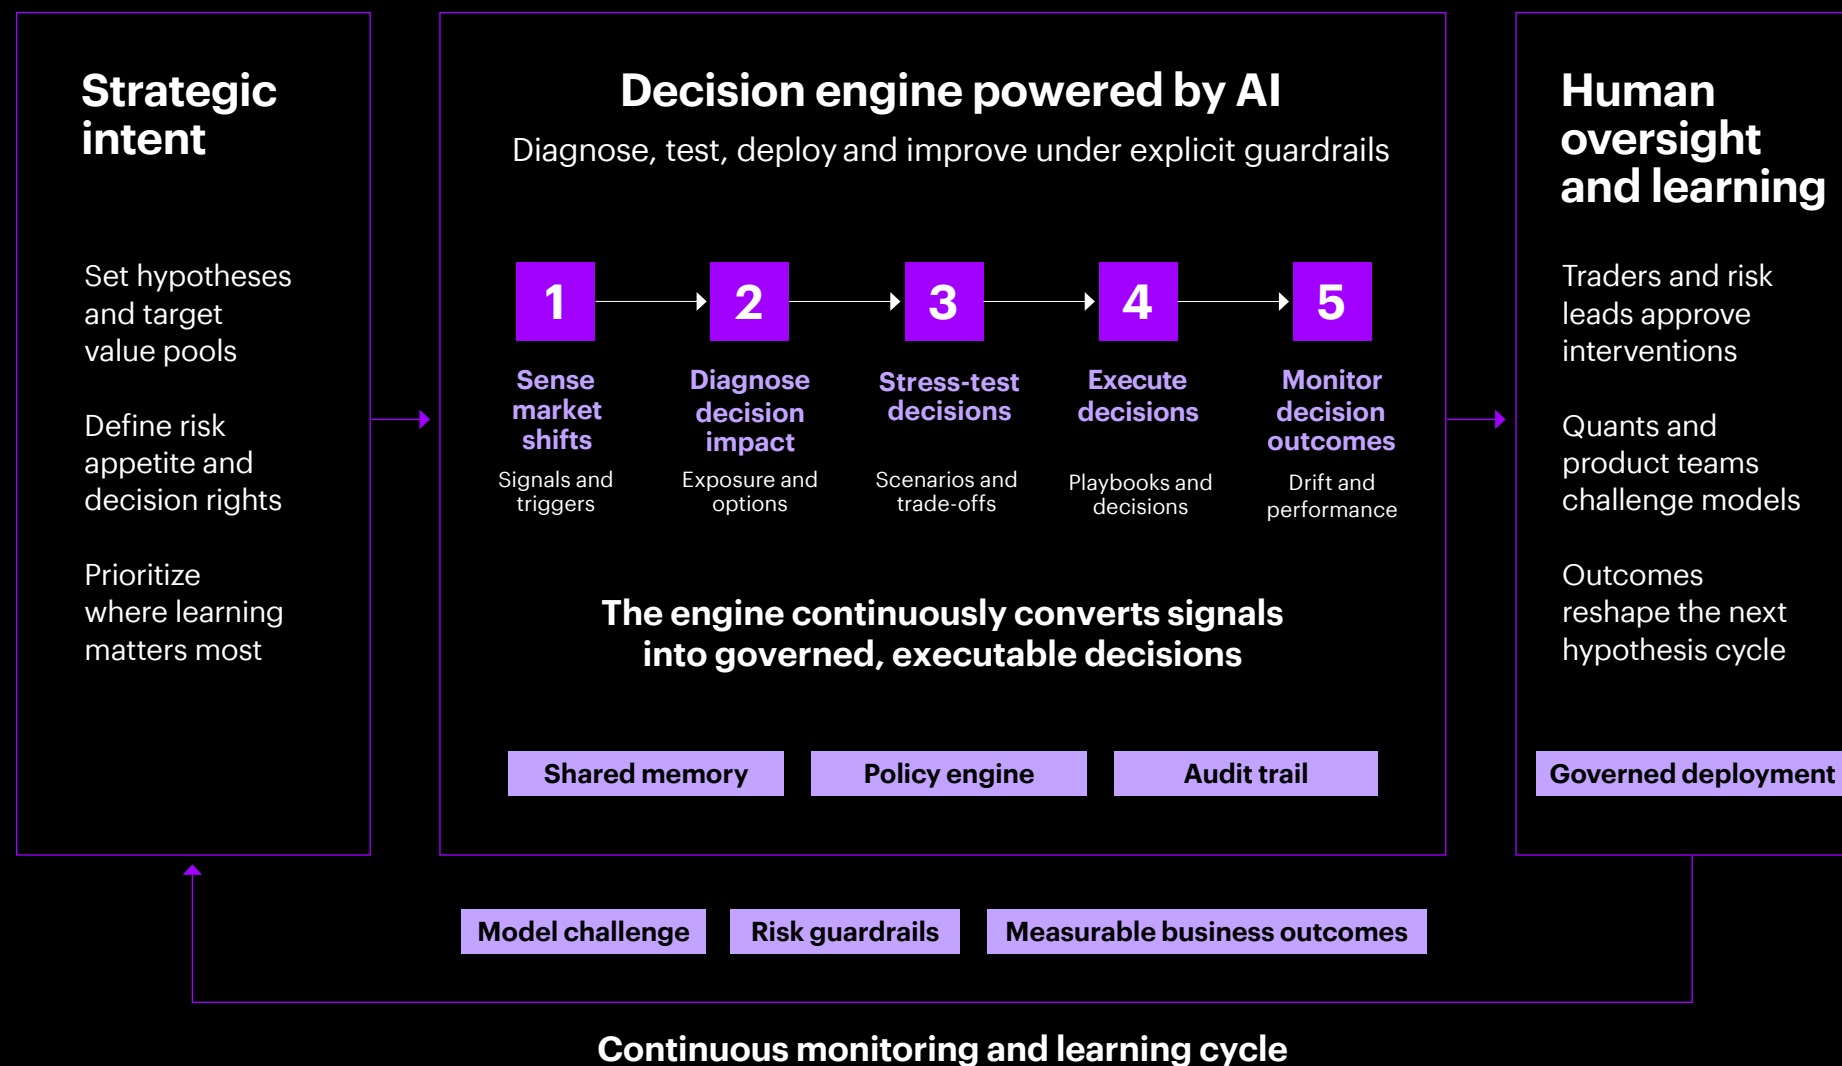
Crucially, each improvement produced auditable artifacts, including metrics, validation scores and plots. These artifacts make drift, performance decay and version rollback explicit and governable.

Operators remained in the loop to approve deployments and set policy guardrails—turning continuous model improvement into controlled, repeatable learning cycles.

Continuous monitoring is essential (see Figure 3). Drift detection systems evaluate behavior and predictive accuracy in real time, flagging anomalies or breakdowns in underlying assumptions as market regimes evolve.

When drift emerges, retraining or recalibration occurs automatically or with minimal intervention. This ensures trading systems remain adaptive, resilient and grounded in current realities.

Figure 3: A business-led operating model for self-learning trading organizations



Self-learning systems work when strategy, controls and continuous improvement operate as one business system

Example:

Accenture Commodity Markets researched and developed novel algorithms such as integrating geopolitical indices as an exogenous regressor to a multivariate commodity diffusion model designed to forecast 10-day-ahead commodity prices.

This model achieved a top-quartile Sharpe ratio on a basket of commodities during alpha testing.

The model was researched, constructed and back-tested using a combination of autonomous novelty check and experimentation followed by proprietary autonomous software engineering.

Human roles evolve accordingly. Traders become risk architects who set hypotheses, intervention thresholds and strategic direction. Quants focus on model design, challenge processes, validation and monitoring. Product and engineering teams maintain the pipelines, infrastructure and guardrails that keep the system reliable. Risk and compliance leads ensure that every agent and every model operates within approved boundaries.

This hybrid model of humans in the lead, supported by AI, is the only viable path to scale. It respects the judgment, creativity and intuition of experienced professionals while allowing AI to handle the complexity, volume and speed that no team, no matter how skilled, could manage manually. Firms that operate this way stay aligned with market speed and sustain an edge others can't match.

For leaders, this means shifting from occasional model updates to a steady rhythm of rapid testing and deployment.

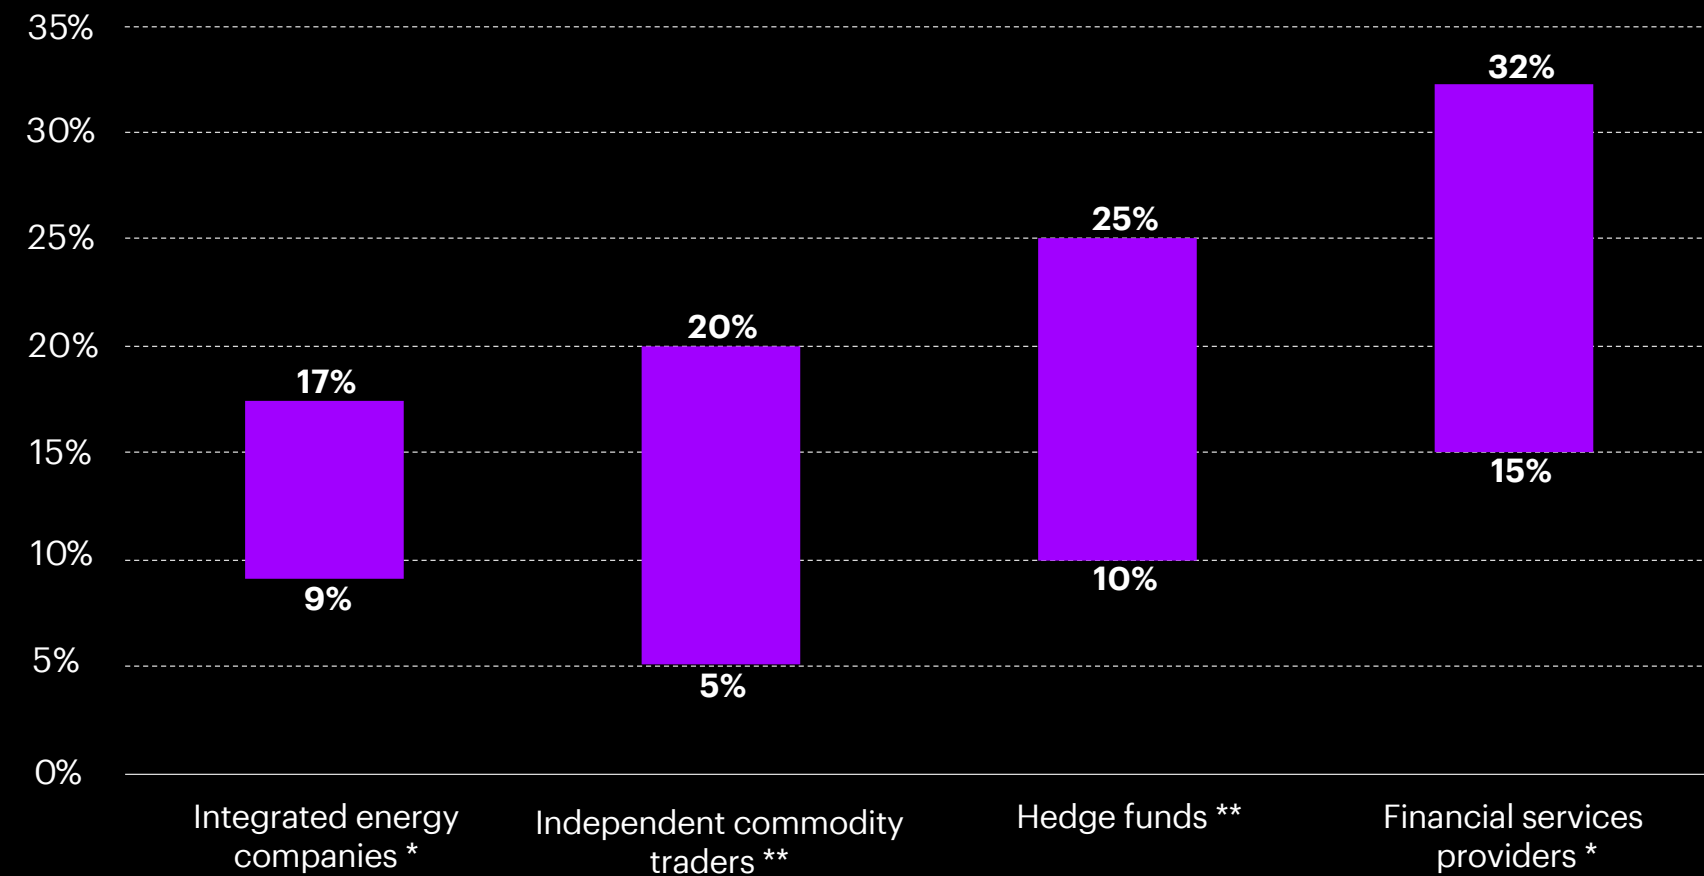
Empower teams to improve signals continuously, keep governance light but effective and reward behaviors that increase learning velocity.

**Breaking through
to scale**



Momentum for AI investments in trading is building. Integrated energy companies, for example, allocate 9–17% of their technology budgets to AI (see Figure 4), with trading predictions ranking as one of the top three most scaled strategic AI bet.^{3,4} Independent commodity traders invest between 5% and 20%. Hedge funds and financial services firms, typically early adopters, lead with 10% to 32%.

Figure 4: Typical range of AI investments as % of overall technology budget



Sources: * © Copyright IDC. Worldwide Wallet: Spend by Market, January 2026

** Accenture interviews with selected commodity trading companies

Despite this growing interest and rising expectations about AI's potential, many trading organizations remain stuck in pilot mode. In Energy, for example, only 11% of executives report having scaled AI for trading predictions.⁵

This is consistent with leaders' experiences in other organizations and industries. Most executives (83%) say gen AI's business potential exceeds their initial expectations, but only 36% have scaled gen AI solutions across their organizations and just 13% report significant enterprise-level value.⁶

What stands in the way?

Several structural challenges appear repeatedly across trading organizations.



The first barrier is fragmentation. When teams experiment in silos, with one desk testing a language model, another exploring forecasting and a third automating reconciliation, value remains isolated. The absence of a shared digital core prevents signals from flowing cleanly into execution or risk functions, limiting impact. Overcoming this requires leadership to mandate unified infrastructure and shared capabilities rather than one-off experiments.

The second is caution in sponsorship. Some leadership teams hesitate to commit decisively, spreading investment thinly across too many use cases. Others attempt to modernize “everything at once,” creating dilution and fatigue. A more effective approach is to focus on two or three high-impact value pools—execution uplift, margin optimization or storage optionality, for example—and push them to full scale before expanding.

Legacy systems create a third form of friction. Trading risk management, enterprise resource planning and other trade-centric platforms,

along with the bespoke trading floor tools and traders’ spreadsheets, all rely on logic, configurations or workflows that evolved over decades and reflect tacit knowledge. These systems are rarely designed for AI-enabled workflows. Leaders who break through this challenge create sandboxes for experimentation, institutionalize knowledge capture through job shadowing and progressively decouple brittle legacy components.

Finally, messy data remains a universal obstacle.

Rather than waiting for perfect data, leading organizations invest in tools and governance that make imperfect data usable.

They codify heuristics such as assay constraints, blend rules, ullage limits or vessel behavior into models and workflows. This transforms tacit expertise into structured intelligence and accelerates the deployment of agentic systems.

The future of ETRM/CTRM: enabling AI-native trading

The role of energy and commodity trading risk management (ETRM/CTRM) platforms is evolving as trading organizations embed AI more deeply into how decisions are made.

Historically, these systems were designed as systems of record—capturing trades, managing risk and ensuring control and compliance across the lifecycle. As markets become faster, more data-intensive and increasingly automated, this role remains critical. ETRM/CTRMs continue to provide the structured data, controls and process integrity required to support trading activities.

What is changing is how decisions are generated and executed around them.

Rather than being displaced, ETRM/CTRM platforms are becoming the foundation for a broader ecosystem of AI-driven capabilities. AI layers sit alongside these systems—generating signals, improving forecasts, automating reconciliation and supporting dynamic optimization—while relying on high-quality, governed data flows from the core platforms.

A key shift is the transition from recording decisions to enabling them. ETRM/CTRMs anchor trade capture, risk, settlements, accounting and reporting, while AI augments these processes by embedding intelligence directly into workflows or connecting through modular architectures.

In this model, competitive advantage will not come from the system alone, but from how effectively organizations integrate AI-driven decisioning with their ETRM/CTRM foundation.

Decisions, redesigned

In energy trading, relationships that once held between commodities are shifting—forcing traders to rethink how markets are read and where advantage comes from.

“One of the most significant shifts is the changing correlation between commodities,” says Francesco Romeo, Digital Business Partner for Global Gas, LNG, Power and Trading at Eni.

“Traditional boundaries are blurring whilst new players are entering the market.” Fundamentals, he argues, are no longer enough.

Instead, competitive edge is moving elsewhere. “It depends on our ability to process more data and uncover correlations that were previously invisible.” In markets defined by complexity, insight rather than access has become the scarce resource.

AI is accelerating this shift. “AI is critical for clearing noise and extracting relevance,” Mr. Romeo says, noting that compressing the time between signal and decision allows firms to go “deeper and faster than humans alone.”

There is a caveat. “If AI is applied only to isolated use cases, it remains automation.” The real opportunity, in Mr. Romeo’s view, lies in embedding it across the trading lifecycle—reshaping how decisions are made rather than simply improving individual steps.

This requires a fundamental redesign of workflows. AI can augment traders, but humans remain responsible for “risk appetite, capital allocation and final decisions,” Mr. Romeo points out. The challenge is to integrate both into a coherent system; one that continuously improves how information is translated into action.

Data sits at the center of this transformation. “Quality, ownership and lineage are prerequisites for trust,” Mr. Romeo continues. Without reliable, AI-ready data, even the most advanced models will struggle to deliver value.

Equally important is how organizations work. Progress depends on co-creation across trading, risk, portfolio management and technology. Over time, AI may become “invisible”—embedded naturally into daily workflows—but reaching that point requires tangible results that build confidence.

“We need a holistic redesign of the trading lifecycle,” — Francesco Romeo, Digital Business Partner for Global Gas, LNG, Power and Trading, Eni

“The biggest challenge is mindset,” says Mr. Romeo. Too often, firms start by asking where AI can be applied, leading to fragmented pilots and limited impact.

Capturing real value demands a different approach. “We need a holistic redesign of the trading lifecycle,” he concludes. Without a clear vision of the end-to-end process, organizations risk missing the larger prize.

The prize is not a better tool, but a better system: “a decision architecture that continuously compounds intelligence and learns at a speed others can’t match.”

In markets where complexity is rising and signals are harder to read, the firms that succeed will be those that redesign not just their models, but also the way they decide.



**Second-order risks
CEOs must weigh**

As AI scales in trading, four structural risks emerge that are invisible in pilots yet capable of eroding the edge AI was meant to create.

Alpha crowding

As AI adoption scales, alpha is at risk of becoming commoditized. When firms deploy similar models, data feeds and foundation models, signals begin to converge and strategy differentiation erodes, even without direct information sharing.

Accenture prototyping points in the same direction: three open foundation models delivered similar standalone accuracy, while meaningful edge emerged only through proprietary orchestration.⁷

The implication is structural. Shared models will not create advantage; only differentiated signal design and orchestration will.

The way forward: invest in differentiated signal research, or accept that widely available AI tooling will compress returns.

Model decay

Signals degrade faster than in prior cycles.⁸ In an Accenture walk-forward evaluation, correcting look-ahead bias exposed a 1.9pp WAPE gap and a 6.1pp directional-accuracy gap that were invisible until rigorous validation.⁹

The way forward: commit to continuous retraining and rapid signal retirement, or accept rising drawdown risk.

Governance paralysis

On average, AI projects take eight months from prototype to production and only 48% make it into production at all.¹⁰ Every month of delay compounds competitors' learning advantage.

The way forward: tier governance to fast-track low-risk applications, or let caution become disadvantage.

Compute cost vs marginal alpha

A prototype can look affordable. The real risk appears when AI moves into production: inference, serving and enhancements become continuous operating costs, not one-off build costs. This is where most of the investment is incurred, and which is typically underestimated, as 96% of deployments cost more than expected.¹¹

Aggregate gains can also hide sharp variance: in Accenture's CATSFM prototype, two product categories improved about 23–28%, one was nearly flat and one worsened by 13.7% versus baseline.¹²

The way forward: benchmark AI spend against measurable P&L, not capability ambition.

Leaders who build signal diversity, retraining discipline, governance agility and cost accountability early will sustain advantage; those who discover these risks reactively will not.



**Five actions to build
structural edge**

In trading, as across the broader enterprise, AI-driven advantage ultimately emerges from operating model redesign—shaped by how decisively CEOs align value, operating models and governance around intelligence rather than isolated use cases.

Effective AI transformation in commodity trading requires sponsorship that traders, quants, risk managers and operations teams trust.

Leaders will set their organizations up for success by focusing on tangible commercial outcomes, aligning incentives and embedding responsible AI principles into every layer of the operating model.

Five high-level actions frame and drive this change:

1. Lead with value

Most successful transformations start when AI adoption is linked to measurable P&L outcomes: better execution, stronger optionality capture and improved margin efficiency. When teams see direct commercial value, adoption accelerates.

People's roles and incentives should evolve to match the hybrid operating model. A typical challenge: how do you get a trader off their desk to become a product owner and advise the technology pod that's developing AI when they still have their daily job and book to manage?

However, when the right incentives for AI transformation are on offer, traders can become powerful champions of reinvention through AI.

Traders should be measured on execution quality, strategic decision-making and collaboration with AI systems. Quants should own performance monitoring and model integrity. Product and engineering teams must deliver reliable pipelines and robust governance. Transparent P&L attribution helps ensure AI-driven improvements are visible and rewarded.

2. Reinvent ways of working

AI cannot scale without a fundamental shift in how people work and how work itself is structured.

This means redesigning critical trading workstreams: removing handoffs that slow execution, standardizing processes that vary by desk and embedding digital and AI tools directly into daily workflows.

At the same time, CEOs need to champion new ways of working that change how teams collaborate, make decisions and learn. Traders, quants, risk managers and operations teams need shared operating rhythms, common data environments and clear decision rights that reflect an AI-enabled model. People need upskilling that enables them to direct autonomous systems.

3. Build an AI-enabled digital core

A secure, modern digital core is non-negotiable. Leading organizations modernize applications, invest in proprietary data and create adaptive architectures.

Shared signal buses, feature stores, unified policy engines and integrated trading systems allow AI to operate safely and at scale. Zero-trust principles and immutable audit logs strengthen control. Autonomy scales fastest not where controls are removed, but where intent, guardrails and accountability are explicit.

4. Close the gap on responsible AI

Responsible AI governance sits at the heart of the operating model. Pre-trade gates, real-time kill-switches, post-trade audits, model validation processes and checkpoints that keep humans in the lead provide the necessary guardrails for confidence and accountability.

These mechanisms allow speed without sacrificing discipline, enabling AI to act autonomously within well-defined risk boundaries.

5. Drive continuous reinvention

CEOs must embed continuous reinvention into the rhythm of the trading organization.

Weekly deployment cycles, quarterly model challenge rounds and proactive retirement of decayed signals help maintain relevance in markets where yesterday's edge could become obsolete tomorrow. Capacity planning across venues and products ensures the system evolves in line with liquidity, volatility and capital allocation priorities.



The path to AI-native decision-making: a 36-month transition

The transition to AI-native trading follows a predictable progression. What differentiates leaders is disciplined execution.

	First 12 months	Next 24 months
Workforce	Launch a digital & AI executive literacy program. Raise a team and coaches to teach AI fluency in 30 days, with one working session each on agent autonomy, model risk, scenario modeling and automation.	Baseline AI adoption and track learning intensity, with a target of 50+ hours per employee per year. Redesign 20 “high-consequence” decision roles that drive P&L into human + agent teams.
Workflows	Pick decision domains that move the P&L—such as Portfolio Shape to Hedge, or Trade to Deliver—for a specific business or commodity bench. Kickstart their reinvention with AI and plan for results in 3–6 months.	Redesign to decision flows: remove 3–5 approvals, embed agents for task routing and publish new decision rights with guardrails. Scale decision-making across all value pools and commodity activities.
Digital core	Within 60 days, take stock of your end-to-end enterprise architecture and assess where the bottlenecks are in terms of integration, data flows, security and agent readiness. Design and implement a best-in-class coalition of partners sharing value accountability.	Modernize the technology landscape and organization, rebalancing the tech spend profile toward reinvention over run costs.
Decision engine	Choose 2–3 key strategic decisions or high-value workflows and start building the decision engine required to enable them—from data and signals to decision-making, execution and control.	Own the decision engine like a balance sheet asset with a clear owner and business accountability for outcomes. Build the discipline and rigor to convert every key decision and high-value workflow into a self-learning capability powered by the decision engine. Define and track decision effectiveness metrics: cycle time, accuracy, frequency and risk-adjusted outcomes.

Turning ideas into edge

The question is not whether AI will reshape agricultural trading, but where value will emerge—and how well firms can capture it.

For Guy-Laurent Arpino, CIO of Louis Dreyfus Company (LDC), the answer lies in moving beyond the obvious. “Our AI journey started with personal productivity, training and upskilling,” he says. But that phase is already passing. “The focus now is on innovation at scale” and, ultimately, on “reinventing the end-to-end trade lifecycle.”

Early gains have been tangible. In one area alone—knowledge exchange through AI agents—LDC has reduced analyst onboarding time by 27%. Across the organization, roughly 2,000 personal agents and 20 scaled AI-enabled products are already in use. In some cases, such as crop monitoring, the impact is fundamental. “We have completely reinvented the process replacing the traditional agronomist way of working. And we keep adding signals to it,” notes Mr. Arpino.

This reflects a broader shift. Rather than automating existing tasks, the more ambitious firms redesign them with AI. “We should look beyond automation of repetitive tasks and focus on reinventing the

processes themselves. Take any business process, understand the expected outcomes, redefine the AI capabilities, traditional and/or agentic, that will deliver value.” Mr. Arpino argues.

Yet extracting value from this shift requires careful balance. Firms must decide “how much freedom to give AI, what the right guardrails are and what the human role should look like. We are not short of ideas. But we also recognize the challenges in scaling AI across the enterprise”.

Even with strong momentum—a strong company mandate and with LDC’s CEO taking on the product owner role for one of its AI products—practical constraints remain. “Finding a product owner with the right skills and accountability to drive change in the business is not always easy,” Mr. Arpino notes, particularly when ambitions outpace available talent.

“The focus now is on innovation at scale,”
— Guy-Laurent Arpino, CIO of Louis Dreyfus Company

The implication is that scaling AI is as much an organizational problem as a technical one. It requires new roles, new incentives and new ways of working.

For Mr. Arpino, two roles are becoming central. The first is the change advocate—someone who challenges legacy habits and pushes the organization to think several years ahead. The second is the AI architect. “Architecture has never been more critical. We have a key role to ensure the corporate data assets generated are compounded and protected. It requires a long-term vision across the entire digital landscape to make strategic decisions about what to build, buy or outsource,” he says, as firms must orchestrate a growing number of systems, decisions and data flows.

These roles reflect a deeper transformation. AI in trading is evolving from collections of tools into integrated systems—where data, models and decisions must align continuously.

In that environment, advantage will not come from isolated use cases, but from coherence—and the ability to learn faster than others.

The opportunity is clear. So, increasingly, is the dividing line.

The reinvention window

Commodity markets are changing structurally. Renewables are reshaping intraday dynamics. Geopolitics is rewiring trade flows. Liquidity is fragmenting. Alpha decays faster. Data complexity continues to rise.


The divergence is becoming unmistakable. Static organizations will compete reactively, constrained by slower cycles and episodic optimization. Self-learning organizations will compound advantage by adapting models, reallocating capital and refining execution in real time.

This is not a technology initiative, but a repositioning of how profit is generated.

At its core, this is a shift in how decisions are made. Organizations that redesign decision-making by embedding AI into how they interpret signals, take actions and learn from outcomes will operate at a fundamentally different speed and level of precision.

For CEOs, the implication is clear. The question is no longer whether to adopt AI, but how quickly to transform the way decisions are made across the business. Those who succeed will not just improve performance; they will change the basis of competition.

The reinvention window is open. As AI scales, advantage will compound for those who move early—and narrow rapidly for those who do not.



How we can help

Accenture Commodity Markets combines deep commodities expertise with AI, advanced analytics and technology to embed intelligence directly into trading, commercial and operational workflows. Our teams include practitioners who have traded, managed risk and run commercial operations alongside engineers and specialists who design, build and operate the systems that turn strategic intent into sustained performance.

We work across the full decision lifecycle, from signal design and model development through to CTRM implementation, agentic workflow deployment and the ongoing capability management that keeps systems performing as markets evolve. This front-to-back coverage provides clients with a single, coherent partner from the first strategic question to the last mile of execution.

With hubs in London, Geneva, Paris, New York City, Houston, Dubai and Singapore, and a dedicated AI hub on the US West Coast, we operate at the heart of global commodity markets and innovation ecosystems, bringing cross-market insight, implementation depth and domain authority to help clients build structural edge in increasingly volatile, data-driven environments.

Contributors

Belinda Deng

Durba Ghosh

Arpit Jain

Lasse Kari

Stavros Kassandris

Gavin McKenzie

Ilias Sarris

James Velasquez

References

¹ Accenture analysis shows that, in the past three years (2023–2025), the market prices of oil, gas and metals have varied between 30–160% above the average during the period, compared to only 20–50% in the pre-COVID period (2017–2019).

² UN Trade & Development, The state of commodity dependence, 2025.

³ Accenture analysis of companies across sectors, 2026.

⁴ Accenture, The front-runners' guide to scaling AI, 2025.

⁵ Ibid.

⁶ Accenture, Making reinvention real with AI, 2025.

⁷ Accenture Commodity Markets context-aware time series foundational model (CATSFM) prototype.

⁸ Chorok Lee, Not all factors crowd equally: Modeling, measuring and trading on alpha decay, Arxiv, 2025.

⁹ Accenture Commodity Markets context-aware time series foundational model (CATSFM) prototype.

¹⁰ Gartner, Gartner survey finds gen AI is now the most frequently deployed AI solution in organizations, 2024.

¹¹ IDC InfoBrief, sponsored by DataRobot, The agentic workforce: From experimentation to scalable impact, 2025.

¹² Accenture Commodity Markets context-aware time series foundational model (CATSFM) prototype.

About Accenture

Accenture helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed for organizations across industries. Our strategy is to be the reinvention partner of choice for our clients and lead in the safe, widespread adoption of AI, and to be the most client-focused, AI-enabled, great place to work in the world. We bring together the talent of our approximately 799,000 people with proprietary assets and platforms, deep process and industry expertise, and leading ecosystem relationships to deliver end-to-end solutions and measurable outcomes at scale. Through our Reinvention Services, we offer broad expertise across Cybersecurity, Digital Core, Finance, Industry and Enterprise, Song, Supply Chain and Engineering, and Talent, with advanced capabilities in AI and Data, Industry and Process, and Technology. We serve approximately 9,000 clients and generated approximately \$70 billion in FY25 revenue.

Visit us at www.accenture.com

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information. Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.