

Enterprise procurement teams increasingly turn to AI tools—especially Large Language Models (LLMs) and retrieval-augmented generation (RAG)—to support supplier negotiations. Yet, trusting these complex systems remains the single biggest bottleneck for widespread AI adoption in procurement decision support. Firms like Garrett Motion, PMI, and Pandora illustrate both the promise and pitfalls of leveraging AI for procurement insights.

Trust: The True Bottleneck in Enterprise AI Adoption

Despite improvements in AI accuracy and capabilities, many enterprises hesitate to fully integrate AI into critical supplier negotiation workflows. Why? Because trust lags behind. Accuracy gains—even if impressive—do not automatically translate into user confidence.

- **Black box excuses:** "The model just decided that" and similar hand-wavy explanations remain common.
- **Random output variance:** Users wonder why the same input sometimes yields different answers.
- **Moving goalposts:** Success metrics get adjusted post-deployment, eroding trust.

Garrett Motion recently evaluated a new AI-powered procurement assistant and highlighted these trust issues. While the AI could produce insightful negotiation tactics, procurement managers struggled with inconsistent outputs and a lack of transparent reasoning behind recommendations.

Explainability, Consistency, and Business-Logic Alignment

Reducing AI risk in supplier negotiations means shifting from mere accuracy to explainability, consistency, and aligning AI outputs with domain-specific business logic.

Explainability: Don't Just Tell Me, Show Me

What would I show in an audit? That's the question every enterprise buyer should ask. If an AI tool recommends accepting a pricing concession from a supplier, the system must provide:

- **Risk flags explanation:** Why is the AI confident in that recommendation? Which data points and historical contract clauses influenced this advice?
- **Source verification:** Which documents, supplier communications, or market benchmarks does this reasoning trace back to?

PMI's procurement team successfully piloted a retrieval-augmented generation (RAG) setup. The system cross-referenced supplier quotes with internal compliance documents and public market data, allowing negotiators to trace [enterprise AI trust framework](#) AI suggestions back to verified sources. This transparency boosted trust dramatically—even when the AI flagged unexpected risks.

Consistency: The Case Against Random Output Variance

In supplier negotiations, inconsistent AI suggestions create confusion. Imagine running the same negotiation scenario twice and receiving conflicting advice.

Pandora's supply chain analysts encountered this problem during an AI-driven price optimization rollout. They demanded stricter controls to ensure that AI outputs remained consistent for identical inputs, which led to improvements in model versioning and prompt engineering.

Aligning AI with Business Logic

AI models are generalists by design. Procurement decision support tools must be customized to incorporate company-specific rules, risk tolerance, and negotiation strategies.


Increasingly, companies use frameworks that enforce business rules either as constraints during generation or as post-processing filters. Guaranteeing that AI recommendations never violate key procurement policies is critical to maintaining trust and reducing risk.

Why Accuracy Gains Do Not Automatically Create Trust

A fascinating paradox is that improving raw accuracy does not guarantee increased trust. Even perfect accuracy on historical data doesn't assuage concerns about:

1. **Hallucinations and fabricated certainty:** LLMs are prone to generate plausible-sounding but false statements. In supplier negotiations, this could mean invented contract terms or supplier capabilities.
2. **Opaque reasoning:** If a model cannot justify why it flags a supplier as risky, users suspect the recommendation is a guess, not insight.
3. **Lack of traceability:** Users want to verify claims against verified sources to avoid compliance issues.

Therefore, investing in explainability and traceability tools often yields higher trust dividends than tweaking accuracy metrics alone.



Hallucinations and the Risk of Fabricated Certainty

One of the biggest dangers in procurement AI tools using LLMs is hallucinations—instances where the model invents facts or misattributes data.

For example, Garrett Motion discovered their AI assistant occasionally fabricated supplier certifications that did not exist. This created false confidence and near risk exposure until caught through manual audits.

Mitigating hallucinations involves:

- **Retrieval-Augmented Generation (RAG):** RAG integrates real-time document retrieval with generation, anchoring outputs to verified data sources and dramatically reducing hallucinations.
- **Risk flags explanation:** Pointing users directly to source documents where claims originate for verification.
- **Human-in-the-loop workflows:** Ensuring sensitive or high-risk recommendations are reviewed before final decisions.

Best Practices for Reducing AI Risk in Supplier Negotiations

Practice Description Example from Industry Incorporate RAG Combine LLMs with live retrieval of procurement documents and market data to ground outputs. PMI uses RAG to anchor AI suggestions against verified supplier contracts and compliance documents. Enforce Explainability Require AI tools to provide traceable risk flags with source citations. Garrett Motion audits the tool outputs for clear explanations tied to supplier records. Ensure Consistency Standardize model versions and prompt frameworks to prevent output variance on identical inputs. Pandora's procurement team enforced this to avoid conflicting negotiation advice. Embed Business Logic Design AI workflows that integrate company-specific procurement policies and risk tolerances. All three companies customize AI constraints to align with their internal negotiation rules. Human-in-the-loop Review Retain human oversight on all AI-driven purchasing decisions, especially those flagged as high risk. Garrett Motion combines AI insights with final approval from experienced procurement specialists.

Conclusion

Reducing AI risk in supplier negotiations is less about chasing perfect accuracy and more about building trust through explainability, consistency, and business-logic alignment. Companies like Garrett Motion, PMI, and Pandora demonstrate that the adoption of advanced AI tools such as LLMs and retrieval-augmented generation (RAG) must be accompanied by rigorous source verification and clear risk flags explanation.



Enterprises that invest upfront in transparency, traceability, and controlled AI integration will reap procurement decision support benefits with minimal exposure to hallucinations and fabricated certainty. In the end, trust is the currency that enables AI to move from pilot projects to indispensable negotiation partners.