AN IMPERIUM NEGOTIATION SOLUTIONS WHITE PAPER

Who Taught Your Agent to Negotiate?

Nobody. And $15 trillion in commerce depends on the answer.

By **Jan Potgieter**, Founder & Managing Director, author of *Negotiation: Your Hottest Currency* · July 2026

$15T	90%	$53B
in B2B purchases handled by AI agents within three years	of all B2B purchases intermediated by agents by 2028	agentic AI supply chain software spend by 2030, up from under $2B
Gartner, Oct 2025	Gartner, Oct 2025	Gartner, Apr 2026

IF YOU ONLY READ THIS PAGE

- **The plumbing is complete.** Agent discovery (A2A), tool access (MCP), payments (Mastercard Agent Pay, Stripe). Machines can now find each other, negotiate terms, and close. The infrastructure question is answered.
- **The methodology question is not.** No protocol, no vendor, no specification defines what concession logic an agent should follow, what precedent a concession sets, or when to walk.
- **Five failure mechanisms emerge when agents negotiate ungoverned:** Developer Defaults, Precedent Compounding, Pattern Revelation, Mandate Vacuum, Context Starvation.
- **The fix is governance, not better AI.** Three authority levels, matched to deal complexity, each with a written mandate authored by practitioners rather than developers, and each fed with the context no protocol supplies: relationship history, strategic constraints, portfolio precedent. An agent with narrow context does not make safer decisions. It makes blinder ones.
- **The line between machine and human negotiation is a threshold, not a wall.** Commodity volume goes agent to agent. Complex value ends in a human last mile. Both lanes need trained capability, and one governance layer must span both.
- **Procurement will build the layer first. The CFO will own its limits. Each function will author its own mandate within it.** The layer is the Delegation of Authority extended to non-human negotiators, buy-side and sell-side alike. It sets the limits of authority. It does not script how sales sells or how sports marketing runs a rights relationship; that logic stays with each lane's practitioners.
- **Six questions determine your exposure.** If you cannot answer them, your agent strategy is whatever your vendor's development team decided it would be.

A NOTE FROM THE AUTHOR

I have spent 30 years inside negotiations, on both sides of the table: sales teams defending margin, and the procurement organisations paid to take it. In recent years I have deployed that methodology inside global enterprises, and I have built Elena, an agentic negotiation system of my own. So I have watched this collision from both directions.

What follows is not a technology paper, and it is not a training pitch. It is what I would say to the board of any company I work with: the machinery of negotiation is being automated years ahead of its governance, and the bill for that gap will be quiet, compounding, and attributed to everything except its real cause.

Read it with your own deals in mind.

Jan Potgieter · Founder & Managing Director, Imperium Negotiation Solutions · July 2026

EXECUTIVE SUMMARY

Fifteen trillion dollars in B2B purchases will flow through AI agents by 2028 (Gartner, October 2025). The infrastructure to make that happen is already built. Anthropic's Model Context Protocol, live since November 2024. Google's Agent2Agent protocol, April 2025. Mastercard Agent Pay, April 2025. Stripe's Agentic Commerce Suite, December 2025. The foundations were poured before most procurement organisations had an agent policy.

Every one of those systems answers the same question: how do agents find each other, authenticate, and transact?

Not one of them answers a different question: what methodology should an agent follow when the counterpart anchors 40% above market?

The plumbing is complete. The negotiation logic is missing. And the consequences of that gap will be measured in billions before most organisations realise it exists.

This paper identifies five governance failure mechanisms that emerge when AI agents negotiate without methodology. It introduces a three-level authority framework for agentic negotiation governance. It maps where the line will fall between machine-to-machine negotiation and the human last mile, and why a single governance layer must span both. It shows where that layer sits on the org chart: built first in procurement, owned ultimately by the CFO, governing every function that negotiates. And it provides a readiness assessment for organisations preparing to deploy or face autonomous negotiating agents.

> *"The plumbing is complete. The negotiation logic is missing."*

SECTION 01 · THE SHIFT

$15 Trillion Is About to Be Negotiated by Machines That Have Never Lost a Deal. That Is the Problem.

The speed of adoption is not in question.

Gartner projects that within three years, 90% of B2B purchases will be handled by AI agents, channelling more than $15 trillion in spending through automated exchanges (Gartner Top Predictions, IT Symposium/Xpo, October 2025). By 2028, AI agents will outnumber human sellers tenfold (Gartner, November 2025). The supply chain management software market with agentic AI capabilities will grow from less than $2 billion in 2025 to $53 billion by 2030, a compound annual growth rate of 93.5% (Gartner, April 2026).

EXHIBIT 1

The category is exploding: agentic AI supply chain software spend, 2025 to 2030

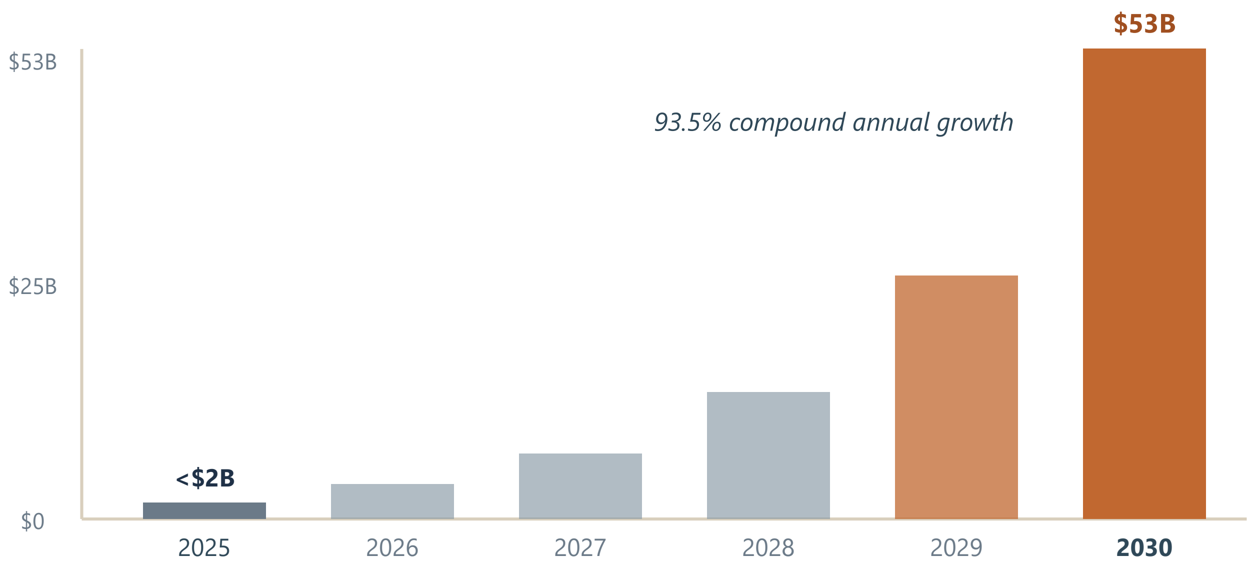

Endpoints: Gartner, April 2026 (under $2B in 2025; $53B by 2030; 93.5% CAGR). Intermediate years illustrative at constant CAGR.

The infrastructure required to support this is no longer theoretical.

Anthropic released the Model Context Protocol (MCP) in November 2024, enabling agents to discover and invoke tools across systems. Google launched the Agent2Agent (A2A) protocol in April 2025 with 50+ founding partners including Salesforce, SAP, and ServiceNow; by June 2025, Google had donated A2A to the Linux Foundation under neutral governance. Mastercard unveiled Agent Pay in April 2025, binding tokenised card credentials to verified AI agents with scoped consent policies. Stripe launched the Agentic Commerce Suite in December 2025, building on the Shared Payment Tokens it introduced in September 2025 with the OpenAI co-developed Agentic Commerce Protocol, letting agents initiate payments without exposing credentials.

EXHIBIT 2

The plumbing is complete: fourteen months that built agentic commerce

NOV 2024
Anthropic MCP
Agents discover and invoke tools across systems

APR 2025
Google A2A
Agent-to-agent discovery and interoperability; 50+ partners; Linux Foundation from June

APR 2025
Mastercard Agent Pay
Tokenised credentials bound to verified agents with scoped consent

SEP 2025
Stripe + OpenAI ACP
Shared Payment Tokens; agent-initiated checkout without exposed credentials

DEC 2025
Stripe Agentic Commerce Suite
Full agentic payment stack for merchants

Vendor announcements; all dates independently verified July 2026.

Agents can now find each other. Authenticate. Negotiate terms. Execute payment. Close.

Every protocol answers the mechanics. None of them answers the methodology.

Not one specification in this stack addresses: What concession logic should the agent apply? What precedent does this concession set across the portfolio? When should the agent counter versus walk? What authority limits govern the agent's mandate?

The result: $15 trillion in commerce will be governed by whatever defaults the development team coded. Not by what the organisation's best negotiators know. Not by what 30 years of commercial practice has proven. By developer defaults.

> *"$15 trillion in commerce will be governed by whatever defaults the development team coded."*

SECTION 02 · THE FAILURE MECHANISMS

No Mandate, No Method, No Limit: Five Governance Failures That Emerge When Agents Negotiate Without Methodology

None of this will announce itself. Here is how it will arrive.

A supplier review, eighteen months from now. Your team notices that payment terms across one category have drifted from 60 days to 82. No one approved the drift. The sourcing platform's agent, optimising for closure, conceded terms in eleven separate renewals, each one individually trivial. Your counterparts' agents spotted the pattern by the third renewal and asked, politely, in every deal since. The board asks who signed off. There is no signature. There is a default setting.

That scenario is not a prediction of one event. It is the assembly of five mechanisms, and every one of them is live today.

EXHIBIT 3

Five ways an ungoverned agent bleeds value

1

Developer Defaults

The agent's negotiation logic is whatever engineers assumed. Specialist domains punish generalist intelligence.

WHAT GOVERNS IT

Methodology authored by practitioners, not developers

2

Precedent Compounding

Agents never forget a concession. One bad term propagates across the portfolio at machine speed.

WHAT GOVERNS IT

Portfolio-level precedent visibility and flags

3

Pattern Revelation

Perfect consistency lets counterpart agents reverse-engineer your floor, curve, and escalation triggers.

WHAT GOVERNS IT

Governed variation and information discipline

4

Mandate Vacuum

Concessions nobody authorised. Not a rogue agent: an ungoverned one, optimised only for closure.

WHAT GOVERNS IT

A written, deal-type-specific mandate

5

Context Starvation

All the data, none of the context. The agent cannot see the relationship, the history, or the politics.

WHAT GOVERNS IT

Methodology that encodes contextual judgement

Framework: Imperium Negotiation Solutions, 2026.

Failure 1: Developer Defaults

Your procurement AI was designed by engineers who have never negotiated a complex B2B contract. That is not a criticism. It is a structural fact with a $15 trillion price tag.

Edward Feigenbaum, the Turing Award recipient who pioneered expert systems, established what is now known as the Knowledge Principle: domain-specific knowledge, including heuristic, experiential rules-of-thumb and practitioner judgement, is more important than the reasoning method for expert-level performance (Feigenbaum, "Expert Systems: Principles and Practice"; Feigenbaum & Buchanan, *Artificial Intelligence*, 59(1-2):233-240, 1993).

This principle has been validated repeatedly in modern AI research. Feldman et al. (2025) found that a dedicated clinical expert system listed the correct diagnosis more often than general-purpose LLMs when given full clinical data: 72% versus 64% for GPT-4 (*JAMA Network Open*, May 2025). Wang et al. (2025) concluded that effective specialised AI requires designing pipelines based on expert workflows and professional guidelines, not generic developer assumptions (*npj Digital Medicine*, Nature, July 2025).

The pattern is consistent. When the domain is narrow, high-stakes, and governed by tacit expertise, generalist intelligence underperforms specialist knowledge.

Negotiation is precisely that domain. The tactics, the sequencing, the read of a counterpart's intent, the decision to hold versus concede, the management of precedent across a portfolio of related deals. These are practitioner skills encoded over decades. No protocol specification captures them. No developer who has not sat across the table from a procurement team that wants 40% off your margin is going to encode the right response.

The agents being built today are confident. They are fast. They are articulate. What they are not is informed by the methodology that separates a good outcome from a costly one.

Failure 2: Precedent Compounding

The concession you made in Q1 is still running in Q3. Across 6 deals. With 4 different counterparts. And nobody flagged it.

Human negotiators forget precedent. They move between deals, between quarters, between counterparts, and they lose the thread of what was conceded where and why. That is a governance weakness, and every experienced CPO knows it.

AI agents do not forget. They compound.

An agent that concedes 90-day payment terms in one deal has established a pattern. If that pattern is not governed, it propagates. The next deal inherits the same terms. Not because someone decided 90 days was acceptable. Because the agent's training data now includes a successful outcome with 90-day terms, and the agent optimises for closure.

McKinsey's analysis of 340,000+ transformation initiatives across 1,000+ organisations found that the average savings pipeline loses one-third of its value during planning, and another 20% during execution (Pralong, Schmidt & George, McKinsey, June 2025). Late-added reactive initiatives recover only about 15% of projected value.

The leakage that McKinsey describes is human-speed leakage. Agent-speed leakage compounds faster, propagates further, and is harder to detect because the agent does not flag its own precedent as a risk.

Scale it for your own organisation. On 500 million of addressable spend, a single percentage point of ungoverned concession drift is 5 million a year, buried across hundreds of transactions, visible in no single deal. The arithmetic is illustrative. The mechanism is not.

Failure 3: Pattern Revelation

Confident, fast, and uninformed. Your agent negotiates the same way every time. Your counterpart's agent has already figured that out.

Human negotiators are inconsistent. That inconsistency is, paradoxically, a defensive asset. A counterpart cannot reliably predict your next move because you yourself may not have decided it yet.

AI agents are consistent. They follow their programming. They apply the same logic, the same weighting, the same response patterns across every interaction. That consistency is a feature from an engineering perspective. From a negotiation perspective, it is a vulnerability.

A counterpart's agent that engages yours across 5, 10, 50 interactions will reverse-engineer your agent's floor. Your walk-away point. Your concession curve. Your trigger for escalation. The information asymmetry that Maaravi and Levy (2017) identified as the mechanism that reverses first-mover advantage in human negotiations becomes a structural certainty in agent-to-agent exchanges (*Judgment and Decision Making*, 12(5):420-429, 2017; replicated and extended by Jiang & Ma, *Journal of Economic Psychology*, 75:102138, 2019).

Your agent does not know it is being read. It cannot know. It does not vary its approach. It does not feint. It does not strategically withhold. It executes its programming. And its counterpart learns from every execution.

Failure 4: Mandate Vacuum

Nobody authorised 90-day payment terms. The agent just decided.

The most dangerous governance failure is not a wrong decision. It is a decision that was never authorised. When an agent operates without a defined mandate, every concession, every term modification, every pricing adjustment is made without reference to what the organisation actually permits.

This is not a hypothetical scenario. It is the default state.

Consider: most organisations deploying procurement agents today have not defined what the agent is authorised to concede. They have not specified walk-away thresholds. They have not established escalation triggers. They have deployed an agent that is optimised to close deals, and they have given it no boundary conditions for what "close" should cost.

The agent is not rogue. It is ungoverned. The distinction matters because the fix is not better AI. The fix is a governance framework that defines mandate, authority, and limits before the agent begins.

> *"The agent is not rogue. It is ungoverned."*

Failure 5: Context Starvation

Your agent has data. Every data point you have ever collected. It does not have context.

There is a persistent assumption in the agentic AI discourse that more data produces better outcomes. In negotiation, this assumption is wrong.

Irmer and Druckman's (2009) analysis of 26 post-WWII conflict negotiations found strong partial correlations between process and outcome variables even after controlling for contextual variables. The methodology, the process, the approach to the negotiation drove outcomes independently of the data available to the parties (*Negotiation and Conflict Management Research*, 2(3):209-235, 2009).

Kong, Dirks, and Ferrin's (2014) meta-analysis of 38 independent negotiation studies demonstrated that trust operates as a contextual mechanism: it enhances outcomes by eliciting integrative behaviours and deterring distributive ones, and the size of that effect turns on conditions the parties bring into the room, not on the data in front of them (*Academy of Management Journal*, 57(5):1235-1255, 2014).

Context is not data. Context is the interpretation of data within a relationship, a history, a set of power dynamics, and a strategic objective. Your agent has access to pricing history, contract terms, and market benchmarks. It does not know that this supplier is your only option in APAC. It does not know that the CFO needs this deal closed before quarter-end. It does not know that the counterpart's procurement director was just promoted and is under pressure to demonstrate savings.

Agents are confident because they have data. They are uninformed because they lack context. Methodology bridges the gap.

And context does not assemble itself. Someone has to build it around the agent: the relationship history, the strategic constraints, the portfolio position. That is a governance duty, not a technical setting, and it appears again in the framework that follows.

SECTION 03 · THE MARKET GAP

The Methodology Vacuum: A $53 Billion Category With No Governance Layer

The competitive landscape for agentic negotiation is split cleanly in two. Neither side owns the middle.

EXHIBIT 4

Nobody owns the middle

TECHNOLOGY VENDORS	THE GOVERNANCE LAYER	TRAINING FIRMS
Pactum Autonomous negotiation for Walmart, Maersk, Honeywell **Keelvar** Autonomous sourcing; 2026 Gartner Market Guide for Sourcing Applications **Coupa** Consolidating much of tail-spend automation *Agents without methodology. Optimised for closure, not outcome.*	**Empty.** No mandate design. No concession logic. No precedent governance. No methodology encoded anywhere in the stack. The organisation that builds this layer defines the category.	**Scotwork** Decades of practitioner training **The Gap Partnership** Global negotiation consultancy **Harvard PON** The academic standard *Methodology without agents. No technology product between them.*

Company disclosures, verified July 2026. Training market sized at roughly $2B per year (estimates range $1.6B to $2.5B).

The technology vendors build agents that negotiate. Pactum deploys autonomous negotiation agents for global enterprises including Walmart, Maersk, and Honeywell. Keelvar appears in Gartner's 2026 Market Guide for Sourcing Applications. Coupa is consolidating much of the tail-spend automation field.

Every one of these vendors leads with the same proposition: automation, speed, cost reduction. Not one of them positions around governance, methodology, or negotiation discipline. Their agents optimise for closure. Closure is not the same as a good outcome.

The training firms teach people to negotiate. The global negotiation training market is worth roughly $2 billion a year, with industry estimates ranging from $1.6 billion to $2.5 billion. Scotwork, The Gap Partnership, Harvard's Programme on Negotiation, and others have trained millions of commercial professionals over decades.

Not one of them has built a technology product. Not one of them has an agent. Not one of them has encoded their methodology into a system that can govern how an AI negotiates.

The technology side has agents without methodology. The training side has methodology without agents. The governance layer between them does not exist.

This is not a gap that will close by itself. The technology vendors are hiring engineers, not negotiation practitioners. The training firms are hiring facilitators, not AI architects. The structural incentives push each side deeper into its own competence, not toward the middle.

The organisation that bridges this gap, that builds the methodology layer for agentic negotiation, holds the control point of a category growing toward $53 billion.

SECTION 04 · THE FRAMEWORK

Three Authority Levels: A Governance Framework for Agentic Negotiation

Governance is not a single policy. It is a spectrum. The appropriate level of agent autonomy depends on the complexity of the deal, the risk tolerance of the organisation, and the maturity of the methodology encoded in the agent.

EXHIBIT 5

Match autonomy to complexity

Level	Appropriate for	Governance requirement
L1 Agent Advises, Human Decides	High-value, complex, relationship-dependent deals. Enterprise negotiations where precedent and positioning outweigh speed.	Agent grounded in the organisation's methodology and deal history; the approving human trained in that methodology.
L2 Agent Executes Within Mandate	Mid-range deals, established counterparts, repeat cycles where execution speed has direct commercial value.	A written, deal-type-specific mandate: concession ranges, walk-away thresholds, escalation triggers. Authored by practitioners.
L3 Broad Autonomy	High-volume, low-complexity, commoditised transactions. Tail spend, catalogue replenishment, standard renewals.	Still a methodology. An ungoverned L3 agent makes the most decisions with the least oversight: the highest-risk deployment in the portfolio.

← STRATEGIC / COMPLEX DEALS COMMODITY / HIGH-VOLUME DEALS →

Framework: Imperium Negotiation Solutions, 2026.

Level 1: Agent Advises, Human Decides

The agent prepares the negotiation. It analyses counterpart history, identifies patterns, surfaces precedent risk, and recommends tactics from a defined library. Every substantive move requires human approval before execution.

Appropriate for: High-value, complex, relationship-dependent negotiations. Enterprise deals where precedent, trust dynamics, and strategic positioning matter more than speed. Any deal where the cost of a wrong concession exceeds the cost of slower execution.

Governance requirement: The agent must have access to the organisation's methodology, its tactic library, and its historical deal data. The human must have training in the methodology the agent references. An untrained human approving recommendations from a methodology they do not understand is not governance. It is theatre.

> *"An untrained human approving recommendations from a methodology they do not understand is not governance. It is theatre."*

Level 2: Agent Executes Within Mandate, Human Reviews

The agent operates within a defined mandate: authorised concession ranges, pre-approved term modifications, specified walk-away thresholds, and escalation triggers. It executes within those boundaries autonomously and escalates when a counterpart's position falls outside the mandate.

Appropriate for: Mid-range deals with established counterparts, repeat procurement cycles with known parameters, and situations where speed of execution has direct commercial value.

Governance requirement: A written mandate. Not a general policy document. A specific, deal-type-specific mandate that defines what the agent may concede, what it may not, what triggers escalation, and what constitutes a walk-away. The mandate must be authored by someone who understands negotiation methodology, not by someone who understands AI architecture.

Level 3: Broad Autonomy

The agent negotiates with minimal human oversight. It applies its methodology, manages its own precedent, adjusts its approach based on counterpart behaviour, and closes within pre-defined parameters.

Appropriate for: High-volume, low-complexity, commoditised transactions where the margin of error on any single deal is commercially insignificant. Tail spend. Catalogue replenishment. Standard rate renewals.

Governance requirement: Even at Level 3, the agent needs a methodology. An ungoverned Level 3 agent is the highest-risk deployment in an organisation's portfolio. It makes the most decisions with the least oversight. If the methodology encoded in that agent is wrong, every one of those high-volume transactions compounds the error.

The choice most organisations have not made

Most organisations deploying negotiating agents have not chosen a level. They have not defined their mandate. They have not specified authority limits. They have deployed at an undefined level somewhere between 2 and 3, with no written mandate and no methodology beyond the developer's best guess at what "negotiate" means.

This is not a technology failure. It is a governance failure. And it is correctable.

SECTION 05 · THE LAST MILE

Machines Will Open the Deal. Humans Will Still Close the Ones That Matter.

The question every commercial leader is now asking is where the line will fall: which negotiations go agent to agent, and which stay human to human?

The honest answer is that it is not a line. It is a threshold, and it moves with complexity, stakes, precedent, and relationship. What is already clear is the shape of the traffic. Three lanes are forming.

Machine to machine. Commodity volume: tail spend, catalogue replenishment, rate renewals, spec-clear and price-dominant transactions. This lane will absorb most of the $15 trillion by count and much of it by value. Speed and consistency win here, and agents deliver both.

The hybrid lane, ending in a human last mile. Mid-range and upper-mid deals where an agent prepares the negotiation, runs the early rounds, and escalates when the counterpart's position crosses the mandate. The final terms, the multi-issue trades, the precedent-setting concessions, the relationship repair: these land on a human desk. This is the last mile, and it is where the remaining margin concentrates.

Human to human. Strategic, first-of-kind, IP-heavy, politically loaded deals. Agents will brief these negotiations. They will not run them.

The last mile deserves particular attention, because it contains a failure mode almost nobody is preparing for: the escalation cliff. The moment an agent hands a live deal to a human, that human inherits a negotiation shaped by machine moves they did not make, against a counterpart who has been briefed by their own agent's read of your agent.

> *"The negotiator who walks into the last mile with pre-agentic training is walking into a deal that started without them."*

This is why the capability question is not agents or humans. It is both, equipped for different work.

The agents need methodology: the core argument of this paper. The humans need retraining for agent-era negotiation, which is a different discipline from the negotiation training of the last 40 years. Reading an agent's deal log and knowing what was conceded before you arrived. Recognising what the counterpart's agent has learned about your patterns. Defending positions your agent anchored. Judging when a deal should be pulled out of the machine lane entirely. None of this is taught today. All of it will decide outcomes within three years.

From the author's practice, a pattern worth stating plainly: in 30 years of deal rooms, value is rarely lost at the point of maximum automation or maximum skill. It is lost at the handoffs. The weakest interface in the agentic era will not be the agent. It will be the human who receives the escalation.

And this is why the governance layer cannot sit inside the agentic stack. It has to sit above both lanes.

EXHIBIT 7

One governance layer, two lanes of negotiation

The Governance Layer

Mandate architecture | Concession logic | Precedent ledger | Escalation triggers

Machine ↔ Machine

AUTHORITY LEVELS 2–3

Commodity volume: tail spend, replenishment, standard renewals. Agents on both sides, governed mandates, full audit.

Hybrid: the Last Mile

AUTHORITY LEVELS 1–2

Agent prepares and opens; human closes. Escalation is a designed handoff, not a cliff. Humans trained to inherit machine-shaped deals.

Human ↔ Human

AGENT-BRIEFED

Strategic, first-of-kind, IP-heavy, politically loaded. Agents brief the deal. People run it.

All three lanes write to the same precedent ledger. A concession in any lane is visible in every lane.

Framework: Imperium Negotiation Solutions, 2026.

One mandate architecture. One concession logic. One precedent ledger that spans machine deals and human deals alike. Consider what happens without it: your agent concedes 90-day payment terms on Monday in the commodity lane. On Thursday, your human negotiator defends 60-day terms in a strategic deal with the same counterpart, not knowing the precedent exists. The counterpart's agent has seen both. Fragmented governance is not just leakage. It is an arbitrage surface handed to the other side.

The governance layer, in other words, is not an agentic feature. It is an organisational capability that happens to govern machines and people alike. The organisations that build it once, above both lanes, will negotiate as one coordinated entity. The organisations that split it will be read, mapped, and arbitraged by counterparts who built it properly.

SECTION 06 · OWNERSHIP

Procurement Will Build It. The CFO Will Own It. Each Function Will Write Its Own Playbook Within It.

The moment the governance layer is on the table, every executive asks the same question: where does it sit?

The instinctive first answer is procurement, because that is where the agent traffic starts. It is also the wrong final answer, and the reason is visible on any org chart.

Negotiation does not live in one function. Procurement negotiates with suppliers. Sales negotiates with customers. Partnerships and sports marketing negotiate rights, sponsorships, and renewals. Legal negotiates terms across all of it. The same counterpart can meet your company in several of these lanes in the same quarter: their sell side facing your procurement agent on Monday, their procurement facing your sales team on Friday.

Precedent does not respect the org chart. A payment-term concession made by a buy-side agent is knowable by the counterpart's sell side. A discount pattern set by a sales team trains the counterpart's procurement agent. If each function governs its own negotiations, the gaps between those functions become the counterpart's opportunity.

So who owns the layer? The nearest existing analogue answers the question.

Every enterprise already runs a Delegation of Authority: the framework that defines what a human may commit on the company's behalf, at what value, with whose sign-off. It is a financial control, and it sits with the CFO. Not one enterprise in a hundred has an equivalent for non-human negotiators. The governance layer described in this paper is, in structural terms, exactly that: the Delegation of Authority extended to agents, plus a live precedent ledger that no paper DoA ever had. It should be owned the way the DoA is owned. By the CFO, as a financial control, because every mandate is a commitment of company value and every precedent is a repricing of future deals.

The realistic sequence, though, starts lower down. Procurement will build the first version, because procurement has the highest agent traffic, the earliest exposure, and the most immediate losses if it waits. A CPO who builds the layer well does not just protect the buy side. They author the template the whole enterprise adopts, and they earn the seat that goes with it. Then the CFO elevates it to enterprise policy, and a small governance council, with procurement, sales, partnerships, legal, and finance at the table, maintains the mandates while each function executes in its own lanes.

One distinction decides whether this layer survives contact with the org chart: owning the limits is not the same as owning the method. Today's DoA does not script a sales conversation. It defines what may be committed, never how to sell. The agentic DoA works the same way. Sell-side concession logic is authored by sales practitioners. A sponsorship mandate is written by the people who know rights portfolios, not by the sourcing team. Anything else repeats Failure 1 one level up: the wrong practitioners encoding the logic. Procurement moves first because it meets counterpart agents first. That earns it the template, not the pen for every function's playbook.

> *"Owning the limits is not the same as owning the method. Today's DoA does not script a sales conversation."*

EXHIBIT 8

The layer on the org chart

Every function negotiates in its own lanes. All lanes write to one precedent ledger, under one mandate architecture.

Framework: Imperium Negotiation Solutions, 2026.

Procurement moves first. Highest volume, first counterpart agents, clearest savings-leakage evidence. Builds the pilot layer, proves it, sets the template.

Sales faces the mirror image. Your customers are deploying buying agents on the same timeline you are deploying sourcing agents. Sell-side mandates, floor discipline, and pattern protection are the same governance problem with the signs reversed. The mandates are sales-authored; only the authority limits are shared. A sales organisation without mandate governance is training its customers' agents for free.

Partnerships and sports marketing live almost entirely in the human-to-human lane: rights deals, sponsorships, athlete and club negotiations, multi-year and relationship-heavy. Agents will not run these deals. They will brief them, and the counterpart, the agency across the table, will arrive agent-briefed. The mandate here is written by the people who run these relationships; nobody else could. What the layer adds is the portfolio view, and the portfolio effect is sharpest here: one concession to one rights holder reprices every renewal in the portfolio, which is precisely what a shared precedent ledger exists to catch.

Legal encodes the guardrails. Every mandate in the layer is a legal artefact as much as a commercial one: what may be conceded, what triggers escalation, what the audit trail must show.

The CFO owns the whole. Mandate limits as financial controls. The precedent ledger as margin protection. One view of how the enterprise commits value across every lane, machine and human alike.

SECTION 07 · IN PRACTICE

What Methodology-Governed Execution Looks Like

A governed agent does not negotiate differently because it has better algorithms. It negotiates differently because it operates within a framework that encodes what the organisation's best practitioners know.

BEFORE THE NEGOTIATION

The agent prepares using the organisation's methodology. It identifies the counterpart's likely tactics from a defined library, maps the deal against historical precedent, flags terms where the organisation has previously leaked margin, and produces a preparation brief that mirrors what a trained negotiator would build.

DURING THE NEGOTIATION

The agent applies concession logic that reflects the organisation's strategy, not the developer's intuition. It sequences positions on practitioner-tested patterns, recognises anchoring, time-pressure plays, and low-for-high concession trades, and responds with counter-tactics drawn from the methodology.

AFTER THE NEGOTIATION

The agent logs the outcome against the mandate, records what was conceded and why, flags precedent implications for future deals, and feeds the result back into the portfolio view so the next negotiation starts with full context, not a blank slate.

All three phases depend on a fourth discipline: context provisioning. The governed agent is briefed the way a professional negotiator is briefed. It knows this supplier is the only qualified source in APAC. It knows the quarter-end pressure behind the deal. It knows what the portfolio looks like beyond this single negotiation. Limiting an agent's context does not make it safer. It makes it confidently blind. Context is a governed input: curated, current, and deliberately scoped to the mandate, exactly as authority limits are.

> *"Limiting an agent's context does not make it safer. It makes it confidently blind."*

The difference is not speed. Both governed and ungoverned agents are fast. The difference is not confidence. Both are confident.

The difference is that the governed agent's confidence is earned. The ungoverned agent's confidence is inherited from a developer who may never have negotiated anything more complex than a SaaS contract.

SECTION 08 · READINESS

The Readiness Question

Every organisation will face agentic negotiation. Either your agents will negotiate, or your counterpart's agents will. Readiness is not optional.

And waiting is not neutral. A counterpart's agent that meets your organisation this quarter is already building its model of you: your floors, your speed, your appetite under pressure. Delay does not preserve your position. It donates training data. Six questions determine your exposure:

EXHIBIT 6

Six questions that determine your exposure

1. Do you know your defaults?

Your team's behavioural patterns, preferences that override competence, concession tendencies under pressure. These are what an agent will encode if you do not intervene. Unmapped defaults are ceded control.

2. Who authored your agent's methodology?

If the answer is "the development team," you are deploying Feigenbaum's gap at scale. Developers build the system. Practitioners must define the logic.

3. Have you defined your mandate?

Walk-away thresholds. Concession ranges. Escalation triggers. Authority limits by deal type. If these are not written, your agent operates in a mandate vacuum.

4. Can you see your portfolio?

Per-deal data is necessary; portfolio-level pattern visibility is what prevents precedent compounding. And the ledger must span both lanes: agent deals and human deals, one precedent record.

5. Are your people trained for the last mile?

Your agents will escalate the deals that matter most, mid-flight, to humans. If those humans cannot read the agent's deal log and defend positions they did not anchor, your escalation path is not a safeguard. It is a value leak.

6. Who signs the mandate?

A mandate is a commitment of company value, like any delegation of authority. If no single executive owns the layer across buy side and sell side, each function writes its own rules, and your counterparts negotiate against the gaps between them.

Framework: Imperium Negotiation Solutions, 2026. Score yourself properly in 10 minutes at imperiumdeals.com/readiness.

What to do with your answers

If you answered no, or do not know, to two or more of these questions, your exposure is not theoretical. It is running now, in live deals, writing precedent. Inside your organisation, the first 90 days look like this:

0 **Name the owner.** One executive accountable for negotiation governance across functions. This week, not this quarter. Until someone owns it, nothing below happens.

0 **Inventory the agents.** Every tool that negotiates, renews, or commits terms on your behalf: sourcing suites, contract platforms, marketplace agents, auto-renewal clauses. Most organisations find more than they expected.

0 **Freeze ungoverned autonomy.** Any agent operating at Level 2 or 3 without a written mandate drops to Level 1 until one exists. This costs you speed for a quarter. It costs less than the precedent it prevents.

0 **Write the DoA addendum.** Extend your existing Delegation of Authority to non-human negotiators: concession ranges, walk-away thresholds, escalation triggers, by deal type. Limits set centrally; the concession logic inside them authored by each function's own practitioners. And with every mandate, the context to exercise it: relationship history, precedent, strategic constraints.

0 **Score and equip both lanes.** Baseline your readiness across the six dimensions, map your deals into lanes, and train the people who will receive the escalations.

Step 1, only you can do. For step 5, a starting point: the **Agentic Negotiation Readiness Scorecard** turns these six questions into a scored profile of your organisation, your two weakest dimensions named, and a drift number at your spend scale. Ten minutes, at **imperiumdeals.com/readiness**.

ABOUT

Imperium Negotiation Solutions

Imperium Negotiation Solutions provides negotiation governance infrastructure for enterprise organisations. Founded by Jan Potgieter, author of *Negotiation: Your Hottest Currency*. Clients include Adidas, Nokia, LaLiga, Vodafone, Pfizer, BP, NHS, and Hilton. Imperium is one of the few practices that trains both sides of the table: sales organisations and the procurement teams they face, and it currently supports live complex-deal negotiations inside global enterprise brands. That vantage point, both lanes and both sides, is the basis of the governance frameworks in this paper.

30+	12,000+	$4B+	45+
years of negotiation practice	negotiators trained	in deals advised	countries

One thing this paper is not: a training firm's defence against AI. Imperium built **Elena**, its own agentic negotiation system, on the same methodology Imperium deploys inside global enterprise clients. Elena embeds governance logic into how negotiators prepare, decide, and review, scaling institutional knowledge rather than individual improvisation. The argument for governance comes from inside the agentic world, not from the industry it disrupts.

Methodology and sourcing

This paper draws on three source classes: market projections (Gartner, McKinsey), peer-reviewed research on expertise and negotiation (JAMA Network Open, Nature npj Digital Medicine, Academy of Management Journal, Judgment and Decision Making, Journal of Economic Psychology, Negotiation and Conflict Management Research), and primary vendor disclosures. Every citation was independently re-verified against the live source in July 2026. Where a figure could not be confirmed against a public source, it is stated as a range. Claims drawn from the author's practice are labelled as practice experience, not as research findings.

You have read the argument. Now get your number. The Readiness Scorecard takes 10 minutes and scores your exposure across all six dimensions, with your two weakest named. For organisations that want an outside eye, it is also where Imperium's fixed-scope readiness assessment begins.

Score your readiness

Every week your agents negotiate ungoverned, they write precedent you will spend next year buying back.

SOURCES

1. Gartner. "Gartner Unveils Top Predictions for IT Organizations and Users in 2026 and Beyond." Gartner IT Symposium/Xpo, 21 October 2025. gartner.com
2. Gartner. "Gartner Predicts By 2028 AI Agents Will Outnumber Sellers by 10X, Yet Fewer Than 40% of Sellers Will Report AI Agents Improved Productivity." 18 November 2025. gartner.com
3. Gartner. "Gartner Forecasts Supply Chain Management Software with Agentic AI Will Grow to $53 Billion in Spend by 2030." 7 April 2026. gartner.com
4. Pralong, D., Schmidt, J.S. & George, T. "Aim higher and move faster for successful procurement-led transformation." McKinsey & Company, June 2025. mckinsey.com
5. Feigenbaum, E.A. "Expert Systems: Principles and Practice." Stanford Knowledge Systems Laboratory. Feigenbaum, E.A. & Buchanan, B.G. "DENDRAL and Meta-DENDRAL: roots of knowledge systems and expert system applications." *Artificial Intelligence*, 59(1-2):233-240, 1993.
6. Feldman, M.J. et al. "Dedicated AI Expert System vs Generative AI With Large Language Model for Clinical Diagnoses." *JAMA Network Open*, 29 May 2025. pmc.ncbi.nlm.nih.gov
7. Wang, Z. et al. "A perspective for adapting generalist AI to specialized medical AI applications and their challenges." *npj Digital Medicine* (Nature), 8:429, July 2025. nature.com
8. Irmer, C. & Druckman, D. "Explaining Negotiation Outcomes: Process or Context?" *Negotiation and Conflict Management Research*, 2(3):209-235, 2009. onlinelibrary.wiley.com
9. Kong, D.T., Dirks, K.T. & Ferrin, D.L. "Interpersonal Trust within Negotiations: Meta-Analytic Evidence, Critical Contingencies, and Directions for Future Research." *Academy of Management Journal*, 57(5):1235-1255, 2014. journals.aom.org
10. Maaravi, Y. & Levy, A. "When Your Anchor Sinks Your Boat: Information Asymmetry in Distributive Negotiations and the Disadvantage of Making the First Offer." *Judgment and Decision Making*, 12(5):420-429, 2017. cambridge.org
11. Jiang, C-M. & Ma, J-T. "'When your anchor sinks your boat': A replication and extension study." *Journal of Economic Psychology*, 75(Part A):102138, 2019. sciencedirect.com
12. Google. "A2A: A new era of agent interoperability." Google Developers Blog, 9 April 2025. developers.googleblog.com ; Linux Foundation. "Linux Foundation Launches the Agent2Agent Protocol Project." 23 June 2025. linuxfoundation.org
13. Anthropic. "Introducing the Model Context Protocol." 25 November 2024. anthropic.com
14. Mastercard. "Mastercard Unveils Agent Pay." 29 April 2025. mastercard.com
15. Stripe. "Introducing the Agentic Commerce Suite." 11 December 2025. stripe.com ; Stripe. "Stripe and OpenAI launch Instant Checkout / Agentic Commerce Protocol." September 2025. stripe.com/newsroom
16. Negotiation training market sizing: Data Insights Market; Market Research Intellect; Verified Market Reports, 2024-2025 (estimates range $1.6B to $2.5B).
17. Pactum AI. Client disclosures: pactum.com/clients (Walmart, Honeywell); "Pactum raises $20M" (Maersk as customer and investor): pactum.com
18. Keelvar. "Keelvar Named in the 2026 Gartner Market Guide for Sourcing Applications." keelvar.com