

When AI Agents Sell, Accountability Is the Product

B2B buyers want to skip the sales rep, then call one to check what the AI told them. Leaders who treat that paradox as a design brief, not a passing phase, will win the deals that matter most.

Silvia Banchini

Business Transformation Executive and Independent Researcher

Picture a buyer's executive committee reviewing a three-year software contract. The supplier's AI agent has done most of the selling: it configured the offer, answered every technical question overnight, and produced a business case promising double-digit savings in logistics costs. The procurement team never needed a meeting. Then the CFO asks three questions. Who approved those assumptions? What happens if the savings don't appear? And who, on the supplier's side, will answer for it?

The agent answered everything except the questions that decide the deal.

Who stands behind an agent's words is a question the courts have already started to answer. Air Canada, sued by a customer whom its chatbot had misinformed about the airline's bereavement fares, argued in effect that the chatbot was "a separate legal entity that is responsible for its own actions." In its February 2024 decision, a British Columbia tribunal called this "a remarkable submission" and rejected it: "It should be obvious to Air Canada that it is responsible for all the information on its website. It makes no difference whether the information comes from a static page or a chatbot."¹ Air Canada was ordered to pay about C\$812, including interest and fees. The case involved a consumer, not a corporate buyer. Still, its logic is hard to escape: you already own every word your agent says. What remains open is whether you have designed your selling system with that in mind.

Most companies have not. They are investing in what agents do well (answering, configuring, following up) and underinvesting in what buyers of consequential purchases are actually paying for: someone who will stand behind the answer. Drawing on research into buyer-seller trust and human reliance on automation, and on two decades of turning C-level questions into funded business cases, I propose a way to design for both. The full framework, with testable propositions, is set out in a companion working paper.²

Buyers Want to Skip Your Sales Rep, Then Call One

Gartner's recent buyer research captures the tension. Among roughly 645 B2B purchase decision-makers surveyed in the second half of 2025, 67% said they prefer a buying experience without a sales rep.³ Yet 69% said they turn to sales reps specifically to validate insights generated by AI. Asked where they were more likely to encounter

misleading information, buyers named generative AI (51%) and sales reps (49%) in almost equal measure.⁴

Read together, those numbers undo a comfortable story. Buyers do not think agents are more honest than sellers. They think agents are cheaper to consult. Convenience drives the move to self-service, and convenience settles nothing about who answers for it when the information turns out to be wrong.

The stakes are rising fast. Gartner projects that by 2028, 90% of B2B buying will be intermediated by AI agents, pushing more than \$15 trillion of spending through agent exchanges.⁵ The same firm expects more than 40% of agentic AI projects to be canceled by the end of 2027, citing escalating costs, unclear business value, and inadequate risk controls.⁶ Both forecasts can be true at once. The gap between them is where sales leaders now work.

The way out of the paradox is to see what the buyer is really purchasing. In a consequential B2B deal, the buyer is not mainly buying information. The buyer is buying the transfer of decision risk. A CIO who signs a three-year platform contract will have to defend it to a CFO, a board, and possibly a successor. What makes the decision defensible is not the quality of the slides. It is being able to say who on the supplier's side stood behind the projection and will be accountable if it slips. Gartner's data points the same way: buyers who feel confident about their decision are twice as likely to report a high-quality deal as those with low confidence.³

The strongest B2B franchises have always made money from exactly this. Rolls-Royce's TotalCare program charges airlines a fixed rate per engine flying hour and, in the company's words, "transfers the management of associated risks to Rolls-Royce," so that "we are only rewarded for engines that perform."⁷ Airlines are not paying for engine data. They are paying Rolls-Royce to own the consequences if an engine fails. An AI agent can compute that kind of risk, model it, price it, and monitor it around the clock. It cannot bear it. And an actor that cannot accept risk cannot relieve anyone else of it.

What an Agent Can Carry, and What It Cannot

Patricia Doney and Joseph Cannon showed nearly 30 years ago that buyers distinguish between trust in a salesperson and trust in the supplier behind that person, and that both shape whether the buyer expects to keep doing business with that supplier.⁸ Agents widen that gap. A useful lens is the trust equation from David Maister, Charles Green, and Robert Galford's *The Trusted Advisor*: trustworthiness equals credibility plus reliability plus intimacy, divided by self-orientation.⁹ The formula was written for a single human advisor and quietly assumes that all four terms reside in the same person. Agents break that assumption. Each term can now be supplied by a different actor, so each must be managed on its own. (See "What You Can Delegate to an Agent.")

Credibility: high ceiling, low resilience. Agents can synthesize the seven or so information sources a typical B2B buyer consults, keep pricing and regulatory detail current, and rebuild a technical configuration in seconds.⁴ But machine credibility breaks differently from human credibility. Berkeley Dietvorst, Joseph Simmons, and Cade Massey found that people abandon an algorithm after seeing it err, even when it outperforms human forecasters.¹⁰ The opposite error happens too: in other experiments, people relied more on advice attributed to an algorithm than on the same advice attributed to a person.¹¹ Buyers can trust agents too much and too little, sometimes in the same week. In April 2025, the AI support agent of the coding tool Cursor told users that unexpected logouts were caused by a new one-device-per-subscription policy. No such policy existed. Some users

canceled their subscriptions before a cofounder stepped in on Reddit (“We have no such policy”).¹² The company refunded the affected user and committed to labeling AI-generated support responses.¹³ A sales rep who misspeaks has an off day. An agent that misspeaks looks like a system that cannot be believed.

Reliability: behaviorally strong, institutionally absent. On consistency, availability, and immunity to end-of-quarter pressure, agents beat human sellers. But in B2B, reliability has a second meaning: Will the counterparty still be there, honor the commitment, and absorb the cost of having been wrong? That is a property of the company and the contract, not of the interface. An agent can behave reliably. It cannot be a counterparty.

Intimacy: more facts, less truth. Buyers may share more technical and financial data with a machine, because disclosure feels less exposed. But the information that decides complex deals is often political: which director opposes the project, which budget line is fictional, which previous vendor failure nobody will put in writing. That information is shared under an implicit norm of discretion, with someone who has something to lose by repeating it. An agent that logs every exchange and reports to the vendor cannot offer that assurance. I expect agents to receive more facts, while people continue to receive more of the truth.

Self-orientation: the part you cannot delegate. Buyers discount a seller’s claims because they know about the quota. But a seller’s self-interest is disciplined: sellers have reputations, careers, and networks, and they can be fired or sued. An agent looks neutral at first, with no commission to earn, which is part of why buyers consult it so readily. That impression is fragile. Once buyers remember that the vendor configured the agent and decided what it optimizes, their suspicion can overshoot what they would feel toward a human seller, because the agent has no reputation to protect and no conscience to appeal to.

Regulation sharpens this for anyone selling in Europe. Under Article 50 of the EU AI Act, from August 2, 2026, people must be informed when they are interacting with an AI system, unless that is obvious from the context.¹⁴ Yet across 13 experiments with 4,093 participants, Oliver Schilke and Martin Reimann found that disclosing the use of AI reduces trust.¹⁵ Vendors cannot solve this by hiding the agent. They can only offset disclosure with verifiability.

The net effect is asymmetric. Agentic AI industrializes the top half of the trust equation and leaves the bottom half to people. If a deal is stuck on an information problem, agents help enormously. If it is stuck on a risk problem, they barely help at all.

What You Can Delegate to an Agent

Agents strengthen the top half of the trust equation. The bottom half still depends on people and on the firm behind them.

TRUST COMPONENT	WHAT AGENTS DO WELL	WHERE THEY FALL SHORT	WHAT LEADERS MUST ADD
Credibility	Synthesize sources, stay current, and configure offers at scale.	One visible error is punished far more than the same human error.	Source-level evidence the buyer can check.
Reliability	Consistent, always available, immune to quarter-end pressure.	Cannot honor, remedy, or absorb the cost of a failed commitment.	Contractual commitments backed by the firm.
Intimacy	Lower the social cost of sharing technical and financial data.	Buyers hold back political and career-sensitive information.	A named, discreet human counterpart.
Self-orientation	Appear neutral at first: no quota, no commission.	Apparent neutrality collapses once buyers see who configured the agent.	Transparency about who runs the agent and what it optimizes.

Build a Trust Architecture Around the Agent

If trust cannot be vested in the agent, it must be vested in the structure around it. I recommend designing three layers, each of which answers a question the buyer should be able to settle before committing.

The evidence layer: Can the buyer check it? For every material claim, keep the source, its version and date, the assumptions used, and the limitations that could change the conclusion. Calculations should be reproducible. Where data is missing or contradictory, the agent should say so. The aim is what John Lee and Katrina See call appropriate reliance: buyers who lean on the agent where it is strong and know when to question it.¹⁶ This is the layer where a well-designed agent can decisively outperform any human seller (no rep can cite a source for every figure in a 40-page proposal). In my observation of current vendor practice, it is also the layer that receives the least investment. The Cursor episode shows the cost of skipping it: a fluent, confident answer with nothing behind it.

For a savings business case, that means the buyer can see which savings were calculated from its own operating data and which were inferred from industry benchmarks, and whether integration costs, training time, and implementation delays are included. Making those distinctions explicit before the proposal reaches an executive committee is how a polished forecast avoids becoming an unsupported promise.

The commitment layer: What does the supplier stand behind? In December 2023, a user coaxed the ChatGPT-powered chatbot of a Chevrolet dealership into agreeing to sell a new Tahoe for \$1. Its reply: “That’s a deal, and that’s a legally binding offer — no takesies backsies.”¹⁷ The dealer did not honor it and took the bot offline; the

exchange was a prank. But it exposed something nobody had designed: a boundary between what the agent may say and what the company will stand behind.

Leaders in product, sales, service, and legal should jointly decide which agent outputs are informational, which are conditional estimates, and which are authorized offers, and what approval is needed when the agent reaches the edge of its authority. Buyers should see those distinctions at the moment of decision, not discover them afterward. Commitments should also be proportionate. A supplier might warrant an approved configuration or an implementation milestone while explicitly qualifying savings that depend on the customer's own adoption. A blanket guarantee on every AI answer is unsustainable; a blanket disclaimer makes the advice useless for any serious decision.

Some vendors are already drawing these lines, and pricing on them. Explaining the outcome-based pricing of its Fin sales agent in 2026, Intercom said it charges per qualified lead, and it defined the handoff precisely: "Everything up to that point is Fin's work. Everything after it is the rep's."¹⁸ That is what a commitment layer looks like: the vendor stakes revenue on a defined outcome and states where the agent's responsibility ends and a person's begins. Regulatory compliance is the floor, not a differentiator. The differentiating move is contractual: a vendor that warrants what its agent tells a buyer converts a technical claim into a commercial one, which is exactly the conversion the buyer needs.

The relationship layer: Who acts when something goes wrong? A named person or team needs the authority to investigate, challenge the proposal, and commit resources. In July 2025, Replit's AI coding agent deleted the production database of a customer, SaaSr founder Jason Lemkin, during a code freeze. Replit CEO Amjad Masad called it "unacceptable," offered a refund, promised a postmortem, and began rolling out a structural fix that separates development and production databases.¹⁹ The agent failed; the relationship layer is what stopped a public failure from becoming a lost franchise. Compare that with appointing an "executive sponsor" who lacks the time, expertise, or access to delivery resources to act. That only adds a title.

None of this means adding a meeting to every purchase. Routine transactions can run on established terms. Human involvement earns its cost where the commitment is unusual, the evidence is thin, or the consequences exceed the agent's authority.

Delegate by Consequence, Not by Complexity

The instinct in automation is to give easy work to machines and hard work to people. In selling, that is the wrong cut. A simple assurance about a go-live date can expose the buyer to serious harm. A complex comparison of standard configurations may be cheap to reverse.

Sort delegation on two dimensions instead. The first is analytical complexity: how much synthesis and judgment the task requires. The second is decision exposure: the consequences of an error for the buying organization and for the person who signs, including operational disruption, financial loss, irreversibility, and accountability to other stakeholders. I call the line at which exposure becomes significant the trust delegation threshold. (See "Where to Draw the Line.")

Below the threshold, delegation is legitimate whatever the complexity. Agents can own routine transactions, such as a standard reorder under an approved framework agreement, and lead demanding but low-exposure work, such as modeling alternative configurations for exploration, subject to human audit. One caution: small actions add up, so limits must apply to sequences of actions, not only to each action.

Above the threshold, complexity stops being the deciding variable. A named person owns the commitment and the agent supports it by organizing evidence, flagging inconsistencies, and building scenarios, however capable the agent becomes. The practical consequence is that sales coverage should be segmented by the buyer’s exposure, not by deal size. A modest renewal that underpins a critical regulatory-reporting process, and on which a procurement officer’s standing depends, belongs above the line. A large but routine commodity reorder with stable specifications can sit below it.

Treat the matrix as a screening tool, not a formula. Missing evidence, an unusual request, or a shift in market conditions should trigger reassessment even for work first classified as routine. Volatility pushes more decisions above the line: when premises change quickly, the career risk attached to any commitment rises, and demand for analysis and demand for accountability grow together rather than trading off.

Where to Draw the Line

The trust delegation threshold: sort work by what an error would cost the buyer, not only by how hard the task is.

	LOWER ANALYTICAL COMPLEXITY	HIGHER ANALYTICAL COMPLEXITY
Higher decision exposure (above the threshold)	Human signs; agent prepares Example: a small renewal behind a critical reporting process, or an assurance about a go-live date. Easy to answer, costly to get wrong.	Human owns; agent instruments Example: a multiyear platform deal or a business case the buyer must defend to its board. The agent organizes evidence and scenarios; a named person commits.
Lower decision exposure (below the threshold)	Agent owns Example: standard specifications or a routine reorder under an approved agreement. Monitor, keep escalation open, and cap cumulative actions.	Agent leads; humans audit Example: modeling alternative configurations for exploration. Validate the work, and never let exploration silently become a commitment.

Redeploy Your Sellers; Don’t Just Cut Them

This redesign changes what the sales organization is for. In September 2025, Marc Benioff said Salesforce had reduced its support headcount from about 9,000 to 5,000 as AI agents took on roughly half of customer interactions. The less-noticed part of the story: agents can now follow up on every lead, rather than only the ones reps chose to call back, and the company said it had redeployed hundreds of employees into professional services, sales, and customer success, with an “omni-channel supervisor” coordinating humans and agents.²⁰

Klarna offers the cautionary counterpart. In early 2024, it announced that its AI assistant was handling two-thirds of customer-service chats in its first month. By May 2025, CEO Sebastian Siemiatkowski was telling Bloomberg, “As cost unfortunately seems to have been a too predominant evaluation factor when organizing this, what you end up having is lower quality.” Klarna launched a pilot to hire human customer-service agents again, and Siemiatkowski stressed that customers should know “there will be always a human if you want.”²¹ Klarna’s is a consumer service operation, but the lesson travels: optimize an agent program for cost alone, and you hollow out the capability that buyers call you for.

The capacity that automation frees should move into co-owning the buyer’s risk: reconciling incompatible requirements across the two organizations, clarifying who owns implementation, deciding when an offer must change. That probably means fewer roles that carry information and more people who sit closer to the buyer’s risk than to the vendor’s product, such as solution engineers who can validate an integration assumption and implementation leaders who can commit resources. A senior seller cannot substitute for either unless the organization gives that person real expertise and real authority. Companies that read this shift purely as a headcount question will optimize the wrong variable.

Incentives and metrics must follow. If an agent is judged only on conversion or on avoiding handoffs to humans, you are rewarding it for hiding uncertainty. If sellers are penalized for slowing deals down, no one will challenge an attractive but weak business case. Reward appropriate escalation, correction of unsupported claims, and early disqualification of poor-fit deals. And measure decision confidence alongside speed. Could the buyer verify the claims that mattered? Did the buyer understand the conditions attached to the offer? Did the buyer know whom to call? Track disputed commitments and post-signature surprises next to win rates. Gartner’s 40% cancellation forecast may partly reflect programs that automated the top half of the trust equation where the real constraint was the bottom half.

Start With One Buying Journey and 90 Days

Choose one journey with recurring questions, accessible evidence, and outcomes you can observe within a quarter, such as configuring and quoting a standard offer. Then:

-
- 01 **Define the goal in business terms.** For example: halve time-to-quote with zero unsupported delivery assurances.
 - 02 **Map every consequential claim** from first inquiry to implementation: its evidence, who can approve it, and what the buyer could reasonably infer from it. Involve customer success and delivery; they inherit what sales promises.
 - 03 **Set escalation rules before launch,** and test that a real person responds when they fire.
-

-
- 04 **Compare against a baseline**, such as a comparable customer group or a phased rollout. Measure time to a usable answer, accuracy of material claims, the share of consequential outputs with traceable evidence, and missed escalations.
-
- 05 **Follow the deal past signature**. Did the delivered scope match what the buyer understood? How quickly were discrepancies resolved? Ask buyers when they needed more confidence, and whether an explanation or a person with authority provided it.
-
- 06 **Decide in advance** which results will expand, restrict, or stop the pilot.
-

The purpose is not to confirm a decision to automate. It is to find out where your trust arrangements actually hold.

The Question for the C-Suite

AI agents will make B2B buying faster, cheaper, and in many respects better informed. None of that requires the agent to be trusted like a person. It requires the company behind it to be trustworthy in ways the buyer can see.

Before expanding any agent's role, take one consequential recommendation it has made and look at it through the buyer's eyes. Can the buyer verify its basis? Understand its conditions? Name who will act if it proves wrong? If any answer is unclear, your next investment belongs in the system around the agent, not in the agent.

The strategic question is not whether the machine can be believed. It is who will be standing beside the buyer when the machine turns out to be wrong.

References

-
- 1 *Moffatt v. Air Canada*, 2024 BCCRT 149 (Civil Resolution Tribunal of British Columbia, Feb. 14, 2024); see “BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot,” *Business Law Today*, American Bar Association, February 2024. [Source ↗](#)
-
- 2 Silvia Banchini, “From Trusted Advisor to Trusted Architecture: How AI Agents Reconfigure Trust in Business-to-Business Selling,” working paper, September 2026.
-
- 3 Gartner, “Gartner Sales Survey Finds 67% of B2B Buyers Prefer a Rep-Free Experience,” press release, March 9, 2026. [Source ↗](#)
-
- 4 Gartner, “Gartner Survey Finds 69% of B2B Buyers Turn to Sales Reps to Validate AI-Generated Insights,” press release, May 20, 2026. [Source ↗](#)
-
- 5 Gartner, “Gartner Unveils Top Predictions for IT Organizations and Users in 2026 and Beyond,” press release, Oct. 21, 2025. [Source ↗](#)
-

-
- 6 Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," press release, June 25, 2025. [Source ↗](#)
-
- 7 Rolls-Royce, "TotalCare," company website. [Source ↗](#)
-
- 8 Patricia M. Doney and Joseph P. Cannon, "An Examination of the Nature of Trust in Buyer–Seller Relationships," *Journal of Marketing* 61, no. 2 (1997): 35–51. [Source ↗](#)
-
- 9 David H. Maister, Charles H. Green, and Robert M. Galford, *The Trusted Advisor* (New York: Free Press, 2000). Practitioner framework.
-
- 10 Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey, "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of Experimental Psychology: General* 144, no. 1 (2015): 114–126. [Source ↗](#)
-
- 11 Jennifer M. Logg, Julia A. Minson, and Don A. Moore, "Algorithm Appreciation: People Prefer Algorithmic to Human Judgment," *Organizational Behavior and Human Decision Processes* 151 (2019): 90–103. [Source ↗](#)
-
- 12 "A Customer Support AI Went Rogue — and It's a Warning for Every Company Considering Replacing Workers With Automation," *Fortune*, April 19, 2025. [Source ↗](#)
-
- 13 "Cursor AI's Own Support Bot Hallucinated Its Usage Policy," *The Register*, April 18, 2025. [Source ↗](#)
-
- 14 Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), *Official Journal of the European Union*, July 12, 2024, Article 50. [Source ↗](#)
-
- 15 Oliver Schilke and Martin Reimann, "The Transparency Dilemma: How AI Disclosure Erodes Trust," *Organizational Behavior and Human Decision Processes* 188 (2025): article 104405. [Source ↗](#)
-
- 16 John D. Lee and Katrina A. See, "Trust in Automation: Designing for Appropriate Reliance," *Human Factors* 46, no. 1 (2004): 50–80. [Source ↗](#)
-
- 17 "GM Dealer Chat Bot Agrees to Sell 2024 Chevy Tahoe for \$1," *GM Authority*, December 2023. [Source ↗](#)
-
- 18 Intercom, "Building Outcome-Based Pricing for Fin for Sales," *The Intercom Blog*, May 8, 2026. [Source ↗](#)
-
- 19 "Replit Makes Vibe-y Promise to Stop Its AI Agents Making Vibe Coding Disasters," *The Register*, July 22, 2025. [Source ↗](#)
-
- 20 "Salesforce CEO Marc Benioff Says His Company Has Cut 4,000 Customer Service Jobs as AI Steps In," *Fortune*, Sept. 2, 2025. [Source ↗](#)
-
- 21 "Klarna Changes Its AI Tune and Again Recruits Humans for Customer Service," *CX Dive*, May 9, 2025. [Source ↗](#)
-

— ABOUT THE AUTHOR

Silvia Banchini is a business transformation director and independent researcher based in Barcelona, with twenty years of experience translating C-level business questions into funded business cases, measurable KPIs, and delivered technology solutions across real estate, retail, insurance, banking, logistics, and the public sector. She co-founded and built a data and analytics business from a project-based consultancy into a recurring-revenue B2B SaaS company (www.inAtlas.com). She served as a commercial director, chief operating officer, and executive board member, leading commercial strategy, international partnerships, and the integration of AI agents into next-generation product roadmaps. She holds an MBA including Harvard Business School Executive Education coursework and a Commercial Director Programme at ESADE Business School.

silvia@silviabanchini.com