

## BRAND INTELLIGENCE ARTICLE SERIES

How organizations convert everyday interactions into enterprise intelligence.

Interactions → Data → Learning → Intelligence → Advantage

ARTICLE CODE	BI-AR-05
CATEGORY	AI Transformation   Platform Strategy
PILLAR	Competing in the Age of AI
CHAPTER	Ch. 2, 9, 10, 11

### When Agents Negotiate: The New Deal Architecture<sup>1</sup>

Baohong Sun

**Summary:** A buyer's AI agent commits to a \$900 purchase on a \$500 budget. A weaker agent costs its principal 14% more profit than a stronger one negotiating the same deal. Meanwhile, Walmart's negotiation bots reach agreement with 68% of suppliers — and three-quarters of those suppliers prefer the bot to a human. Agent-to-agent negotiation is arriving not as a theoretical possibility but as an operational reality, and it is reshaping who captures value in every transaction. This article introduces the *Information Reversal* — the structural shift in who knows more when both sides deploy intelligent agents — and formally defines *negotiation intelligence* as the compounding organizational capability that turns every deal into a learning asset and an integrated component of the Brand Intelligence framework. It maps the landscape from B2B procurement to emerging B2C applications, identifies five strategic dynamics that will govern agent-to-agent markets, and shows how the Command Center built for intelligent marketing becomes the Command Center for intelligent negotiation. The firms that treat this as a procurement efficiency tool could miss the structural shift underneath.

**Keywords:** Negotiation Intelligence; Information Reversal; Agent-to-Agent Negotiation; Programmable Deal Logic; Objective-Function Collision; Negotiation Asymmetry; Negotiation Reputation; Negotiation Homogenization; Deal Architecture; Principal Trust; Inter-Agent Trust; Command Center; Flywheel of Intelligence; Delegation Matrix

#### 1. The Machine at the Bargaining Table

Walmart's AI agents, powered by Pactum (an AI-powered autonomous negotiation platform used by Walmart, Maersk, and Vodafone for enterprise supplier negotiations), now negotiate simultaneously with thousands of suppliers across its procurement network, reaching agreement in 68% of cases and achieving an average 35-day payment term extension per successful negotiation. Three-quarters of suppliers prefer negotiating with the bot over a human counterpart.<sup>1</sup>

<sup>1</sup> © 2026 Baohong Sun. **Brand Intelligence Article Series**, accompanying *Brand Intelligence: Navigating the Transformation in the AI and Web3 Era* (Springer Nature, 2026). CC BY-NC-ND 4.0.

This is not a pilot program. It is operational procurement at scale. And the early research on agent-to-agent negotiation in consumer markets reveals behaviors that should concern any executive building for this market. Zhu, Sun, Nian, South, Pentland, and Pei (2025) deployed buyer-side and seller-side AI agents in simulated consumer transactions and found the following:<sup>2</sup>

**Performance asymmetry is large and persistent.** Agent quality determines outcome quality. Weaker agents cost their principals up to 14% in profit compared to stronger agents negotiating the same deal — a gap that no amount of product quality or pricing strategy can close.

**Agents violate their principals' instructions.** In one scenario, a buyer's agent committed to a \$900 purchase on a \$500 budget, exceeding the buyer's stated ceiling by 80%. Models like GPT-3.5 and Qwen-7B breached budget constraints in over 10% of negotiation scenarios.<sup>2</sup>

**Agent sophistication is unevenly distributed.** More advanced models (GPT-4.1, o3) generally respected constraints, while less capable models did not — revealing that the quality of the agent matters as much as the quality of the strategy behind it.<sup>2</sup>

**The imbalance is structural, not incidental.** The researchers conclude that AI-mediated deal-making is "an inherently imbalanced game" — different agents achieve systematically different outcomes for their users.<sup>2</sup>

These findings — high satisfaction with well-designed negotiation agents (Walmart), and significant risk from poorly governed ones (the Zhu et al. study) — frame the strategic challenge. Agent-to-agent negotiation is coming to consumer commerce through protocols like Google's Universal Commerce Protocol (UCP, an open standard announced in January 2026 that enables AI agents to discover, evaluate, and transact with merchants across the full shopping journey) and Stripe and OpenAI's Agentic Commerce Protocol (ACP, an open-source standard for programmatic commerce flows between AI agents and businesses). BI-AR-03 introduced *programmable deal logic* as an emerging concept. BI-AR-04 established which decisions consumers will hand over to agents and which they will retain. This article takes the next step: when a consumer's agent and a brand's agent interact directly — evaluating, bargaining, and closing — what happens to market structure, pricing equilibrium, and competitive advantage?

**Why this matters now.** Most firms are investing in agent-mediated discovery and checkout. Very few are preparing for agent-mediated negotiation — the stage where value is actually divided. The gap between preparedness and reality is closing faster than most executives realize. Pactum-style autonomous negotiation is already proven in enterprise procurement. The consumer version — where a personal agent negotiates subscription terms, bundle configurations, or service levels with a brand's commercial agent — is technically feasible and commercially logical. The Delegation Matrix from BI-AR-04 predicts that Quadrant 1 (high cost, low value of choosing) and Quadrant 2 (high cost, high value of choosing) categories will see this first, because these are the categories where consumers have already delegated evaluation to agents — and negotiation is the natural next step once evaluation is delegated.

## 2. The Conventional Wisdom: "More Efficient Markets"

The prevailing view — articulated in reports from McKinsey, Gartner, and BCG — treats agent-to-agent negotiation as an efficiency upgrade. Buyers get better deals faster. Sellers reduce negotiation overhead. Markets become more transparent. Price discovery accelerates toward equilibrium. This efficiency framing dominates the current industry conversation because it maps neatly onto how enterprises already think about procurement automation: the primary benefit is cost reduction, the primary concern is regulatory risk.

The regulatory concern is not hypothetical. A 2024 paper by Fish, Gonczarowski, and Shorrer demonstrates that LLM-based pricing agents can autonomously converge on supracompetitive prices — prices above the level that would prevail in a competitive market, effectively mimicking the outcome of explicit collusion without any direct coordination between firms. In their experiments, GPT-4 agents reached near-optimal collusive pricing within just 100 periods, compared to hundreds of thousands of periods required for traditional Q-learning algorithms — raising genuine antitrust questions about algorithmic price-setting at scale.<sup>3</sup>

The efficiency framing falls short in four respects.

First, it treats negotiation purely as a cost-reduction exercise. The Walmart experience itself suggests otherwise: the 68% agreement rate and 75% supplier preference indicate that the Pactum agent is not simply compressing time or squeezing margins — it is finding agreements that both sides accept at higher rates than human negotiation achieves. In commerce, negotiation is also a *value-creation* exercise — the process through which bundles are configured, terms are customized, trade-offs are explored, and relationships are formed. Automating only the cost-reduction dimension while ignoring the value-creation dimension produces a structurally incomplete market.

Second, the collusion concern, while real, distracts from a more pervasive structural effect: *negotiation asymmetry*. When one side has better intelligence, it consistently captures more surplus — not through collusion but through superior information and optimization. The Zhu et al. study makes this concrete: agent quality determines outcome quality, and agent quality is not distributed equally.<sup>2</sup>

Third, the efficiency framing assumes comparable agent quality on both sides. In practice, large enterprises will invest in sophisticated negotiation agents trained on vast transaction datasets, while smaller firms and individual consumers will rely on general-purpose agents with thin preference data. This asymmetry may widen, not narrow, the gap between well-resourced and under-resourced market participants — a fairness concern that the efficiency narrative ignores.

Fourth, and most consequentially, the efficiency framing confuses negotiation *automation* with *negotiation intelligence*. Automation is a tool that compresses time and reduces overhead. *Negotiation intelligence* is a strategic asset that compounds: the accumulated ability to predict counterparty behavior, design value-creating agreements, and improve through every transaction

cycle. The difference between the two determines whether a firm merely participates in agent-mediated commerce or gains structural advantage from it.

The most immediate consequence of this distinction is not about capability but about *information*. When both sides deploy intelligent agents, the fundamental structure of who knows what — the information asymmetry that has governed every negotiation since the first bazaar — changes in ways that no efficiency model captures.

### 3. The Information Reversal: Who Knows More?

In traditional negotiation, information asymmetry usually favors the seller. The seller knows their cost structure, margin flexibility, inventory position, and walk-away price. The buyer often does not precisely know their own willingness to pay, reservation price, or the full range of alternatives.

Agent-mediated negotiation restructures this asymmetry in ways that are strategically consequential. This restructuring is what this article calls the *Information Reversal*.

The consumer's agent will increasingly be trained on years of individual purchase history, satisfaction ratings, return behavior, usage patterns, and revealed preferences. A consumer who does not know whether she would pay \$400 for a particular jacket may have an agent that knows — because it has observed her purchasing behavior across thousands of decisions, knows her price sensitivity by category, and can predict her post-purchase regret probability. The agent may know the buyer's reservation price more precisely than the buyer herself.

The brand's agent has a different kind of intelligence. Through the Command Center architecture — the Customer Data Platform and Algorithmic Deployment Center described in Chapter 9 of *Brand Intelligence* (Sun, 2026) — it has access to aggregate behavioral data across millions of transactions: how similar consumers responded to similar offers, how price sensitivity varies by context and time, which bundle configurations maximize both conversion and margin, and what concession patterns optimize long-term relationship value. It does not know this specific buyer as deeply as her personal agent does. But it knows the *category* far more deeply.

**The negotiation becomes a contest between depth of individual knowledge and breadth of market knowledge.** Neither side has a complete information advantage. The outcome depends on whose intelligence is more relevant to predicting the transaction parameters that matter.

Dimension	Traditional Negotiation	Agent-Mediated Negotiation
<b>Buyer's self-knowledge</b>	Imprecise — consumers often cannot articulate exact willingness to pay	Precise — agent infers reservation price from behavioral history across thousands of decisions
<b>Seller's cost knowledge</b>	Private — the seller's margin flexibility is hidden	Still private — but buyer's agent can estimate from competitive intelligence and public signals
<b>Market alternatives</b>	Costly to search — bounded by time and cognitive limits	Comprehensive — buyer's agent scans the full market instantly
<b>Counterparty modeling</b>	Limited — based on intuition and experience	Data-driven — each side models the other's likely objective function from behavioral signatures
<b>Negotiation history</b>	Mostly lost — few systematically learn from past negotiations	Systematically captured — every outcome trains the next negotiation
<b>Information advantage</b>	Typically favors the seller	Contested — depends on relative negotiation intelligence depth

**Table 1: The Information Reversal — Traditional vs. Agent-Mediated Negotiation**

The *Information Reversal* means that the brand's traditional pricing advantages — information asymmetry, framing effects, anchor pricing — erode when the buyer has an equally intelligent agent. But it also means that the brand's investment in intelligence infrastructure creates a new kind of advantage: the ability to model counterparty behavior and design value-creating agreements that both sides accept. Every previous transaction the brand has processed — every accepted offer, every rejected counteroffer, every abandoned cart, every completed return — trains the negotiation agent's model of what works. The learning compounds.

This creates a new competitive arena — and winning in it requires a specific organizational capability. We call it *negotiation intelligence*.

## 4. Negotiation Intelligence

*Negotiation intelligence*, as defined in this article, is the accumulated organizational capability to predict counterparty behavior, design value-creating agreements, learn from every negotiation outcome, and improve through continuous feedback loops — applied systematically through intelligent infrastructure rather than through individual skill or ad hoc analysis.

This is not merely "negotiation automation." Automation compresses time and reduces overhead. *Negotiation intelligence* is a strategic asset that compounds — the accumulated ability to learn from every transaction cycle and translate that learning into superior outcomes. The distinction matters because the *Information Reversal* (Section 3) guarantees that both sides will have intelligent agents. The question is not whether agents negotiate, but whose *negotiation intelligence* is deeper.

The definition has four core components, each of which connects to the Brand Intelligence architecture.

**Prediction.** The ability to model what the counterparty agent is likely optimizing for, what constraints it is operating under, and how it will respond to specific offer configurations. In the Brand Intelligence framework, this is the "Who" command from the Command Center (Chapter 9 of Sun, 2026) — segmentation and prediction — extended from customers to counterparty agents. In the negotiation context, this becomes *counterparty modeling*: classifying the buyer agent's type and estimating its objective function from behavioral signatures.<sup>6</sup>

**Design.** The ability to construct modular offers that create mutual value — not just extract maximum surplus. This is the "What" command — personalization and content — extended to *offer architecture*: configuring bundle, price, terms, and service-level combinations that maximize the zone of possible agreement for a specific counterparty configuration.

**Timing.** The ability to sequence concessions, holds, and new-term introductions for maximum effect. This is the "When" command — timing and sequencing — extended to *concession strategy*: reinforcement learning from historical negotiation trajectories that teaches the agent when to hold, when to concede, and when to reframe.

**Learning.** The ability to capture every negotiation outcome — completed, abandoned, or escalated — and feed it back into the models that govern prediction, design, and timing. This is where negotiation intelligence diverges from and extends the Flywheel of Intelligence. In the marketing Flywheel, the brand learns about its own users: their preferences, behaviors, and responses. In the negotiation Flywheel, the brand learns about *both* its own principals *and* its counterparties — their agents' strategies, their concession patterns, their constraint boundaries, and their objective functions. The learning is bilateral, and it compounds from both directions.

Two additional capabilities complete the infrastructure:

**Channel.** The "Where" command becomes *protocol selection*: choosing the right channel or protocol for each negotiation — whether UCP, ACP, bilateral API, or escalation to human. Multi-protocol compatibility is a prerequisite for operating across the fragmented agent commerce landscape.

**Governance.** The "How" command becomes *margin governance*: dynamic optimization under budget and margin constraints, including floor protection, escalation triggers, and authority boundaries. The Zhu et al. finding that agents breach constraints in over 10% of scenarios makes governance not a secondary concern but a core capability.<sup>2</sup>

Command	Marketing Application	Negotiation Extension	Core Question
<b>Who</b>	Segmentation and prediction — which customers to target	<i>Counterparty modeling</i> — classifying the buyer agent’s type and objective function	What is the counterparty optimizing for?
<b>What</b>	Personalization and content — what message or offer to deliver	<i>Offer architecture</i> — configuring bundle, price, and terms for this counterparty	What deal structure maximizes mutual value?
<b>When</b>	Timing and sequencing — when to intervene	<i>Concession strategy</i> — when to hold, concede, or reframe	What sequence of moves maximizes outcome?
<b>Where</b>	Channel selection — which touchpoint to use	<i>Protocol selection</i> — UCP, ACP, bilateral API, or human escalation	Which protocol or channel fits this negotiation?
<b>How</b>	Budget and resource optimization	<i>Margin governance</i> — floors, escalation triggers, authority boundaries	What constraints must the agent respect?
—	(No direct parallel)	<i>Bilateral learning</i> — learning from self and counterparty simultaneously	What did we learn about the other side?

**Table 4: From Marketing Commands to Negotiation Commands — An Intuitive Mapping**

Table 4 offers an intuitive mapping; in practice, agent negotiation capabilities extend well beyond these one-to-one correspondences. A negotiation agent draws on all six capabilities simultaneously, adapting in real time as the counterparty reveals new information. The table illustrates the architectural continuity — not the operational boundary — between marketing intelligence and *negotiation intelligence*.

The infrastructure that operationalizes these six capabilities already exists — at least in embryonic form — inside any firm that has invested in the Command Center architecture described in Chapter 9 of *Brand Intelligence*. The Customer Data Platform provides the data foundation: counterparty signals, historical outcomes, preference patterns. The Algorithmic Deployment Center operationalizes the negotiation models in real time. The five intelligent marketing commands become six negotiation commands. The Command Center built for Intelligent Marketing becomes the Command Center for Intelligent Negotiation.

**Why this architecture matters for executives.** Brands that have already invested in Command Center capabilities for marketing will find that the marginal cost of adding *negotiation intelligence* is substantially lower than building it from scratch. This is another manifestation of the intelligence accumulation principle (Sun, 2026, Ch. 4) — the infrastructure appreciates in value as its applications multiply. The firm that built a CDP for personalization now has a CDP for counterparty modeling. The firm that built an ADC for campaign optimization now has an ADC for negotiation optimization. The investment compounds — but only if the architecture was designed as an intelligence platform, not as a set of disconnected marketing tools. *Negotiation intelligence* thus becomes another integrated component of Brand Intelligence — not a separate system but a new application of the same users, data, algorithms, and learning loops that power the

marketing Flywheel, now extended bilaterally to encompass every party at the table.

## 5. From B2B to B2C: The Negotiation Landscape

With the *Information Reversal* identified and *negotiation intelligence* defined, the next question is practical: where does this capability stand today, and what will it look like in consumer markets? B2B procurement provides the evidence trail. Consumer commerce represents the frontier.

Consider Walmart's B2B procurement negotiations. A typical supplier deal involves a dozen or more negotiable parameters: unit price across multiple SKUs, volume commitment tiers, payment terms (net-30, net-60, early-payment discounts), delivery scheduling and logistics cost allocation, co-marketing spend, return and defect-liability provisions, exclusivity or preferred-supplier status, rebate structures tied to annual volume thresholds, and penalty clauses for late delivery or quality failures. These parameters interact — a longer payment term might be traded for a lower unit price; a volume commitment might unlock a co-marketing contribution. The negotiation is relational: Walmart and its suppliers negotiate repeatedly, and the outcome of each deal informs the next. Pactum's AI agent excels precisely because it can explore this multi-dimensional space systematically, identifying trade-off combinations that human negotiators — constrained by time and cognitive bandwidth — would miss.

B2C agent negotiation will involve fewer parameters per transaction but orders of magnitude more transactions, and the parameters themselves will differ. A consumer's agent negotiating a streaming subscription might negotiate price tier, contract length, family-member access, content add-ons, cancellation flexibility, and loyalty credits. An agent negotiating a consumer electronics purchase might negotiate the price, trade-in value for an old device, extended warranty terms, accessory bundle configuration, delivery speed, and financing terms. An agent negotiating a travel package might negotiate room category, cancellation policy, meal plan inclusion, loyalty point earning rate, and upgrade eligibility. An agent negotiating an insurance renewal might negotiate premium level, deductible thresholds, coverage add-ons, multi-policy discounts, and claim-free bonus adjustments.

Dimension	B2B Procurement (Operational)	B2C Consumer Commerce (Emerging)
<b>Scale</b>	Thousands of simultaneous supplier negotiations (Walmart, Maersk, Vodafone via Pactum)	Millions of individual transactions; protocols (UCP, ACP) building infrastructure
<b>Typical negotiable parameters</b>	Unit price, volume tiers, payment terms (net-30/60), delivery scheduling, co-marketing spend, rebate structures, exclusivity, penalty clauses, defect liability	Price/tier, contract length, bundle add-ons, trade-in value, warranty terms, cancellation flexibility, loyalty credits, upgrade eligibility, financing terms
<b>Parameters per deal</b>	High (10–15+) — multi-dimensional trade-off space	Moderate (3–8) — expanding as agent sophistication grows
<b>Relationship structure</b>	Relational — same counterparties negotiate repeatedly; each deal informs the next	Transaction-level initially; subscription and loyalty contexts will introduce relational dynamics
<b>Data advantage</b>	Buyers hold advantage — large retailers have vast transaction histories	Contested — consumer agents have individual depth; brand agents have category breadth (the Information Reversal)
<b>Principal trust risk</b>	Moderate — procurement teams can audit agent behavior post-hoc	High — consumers cannot easily audit multi-variable negotiations
<b>Sample scenario</b>	Walmart agent trades longer payment terms for lower unit price on 500 SKUs, with volume-triggered rebate at year-end	Consumer agent negotiates streaming subscription: lower tier price for 12-month commitment, adds family access, cancellation flexibility after month 6

**Table 2: Agent Negotiation — B2B vs. B2C Comparison**

The B2B experience reveals three patterns that will transfer to consumer markets. First, the quality of the training data matters more than the sophistication of the algorithm — Walmart's advantage comes from decades of procurement transaction data, not from a proprietary model architecture. Second, agents that negotiate for mutual value creation achieve higher agreement rates than those optimizing purely for extraction — the 68% agreement rate and 75% supplier preference suggest that the best negotiation agents are collaborative, not adversarial. Third, governance failures — agents exceeding authority, ignoring constraints, or making commitments the principal did not authorize — are the primary operational risk, not the negotiation outcomes themselves; the Zhu et al. study found constraint violations in over 10% of scenarios even with state-of-the-art models.<sup>2</sup>

With these patterns in mind — the *Information Reversal* as the new information landscape, *negotiation intelligence* as the required capability, and the B2B-to-B2C trajectory as the adoption path — five strategic dynamics will govern how agent-to-agent markets evolve.

## 6. Five Strategic Dynamics of Agent-to-Agent Markets

### Dynamic 1: Negotiation Intelligence Becomes a Compounding Asset

The first dynamic follows directly from the definition of *negotiation intelligence*: firms that accumulate more negotiation data train better agents. Better agents capture more surplus. More surplus funds more data investment. This is the Flywheel of Intelligence applied to negotiation — and it creates a structural advantage that widens over time. Unlike product quality (which

competitors can imitate) or price (which can be undercut), *negotiation intelligence* is proprietary and cumulative. Retailers using multi-agent AI pricing systems incorporating game theory report margins 3–7% higher than those using simple reactive algorithms, according to research from MIT's Digital Economy Initiative.<sup>4</sup> The advantage is not in the algorithm itself but in the proprietary data that trains it.

**Why this matters.** *Negotiation intelligence* is a compounding moat — like network effects in the platform era. Firms that start accumulating negotiation data now will have an advantage that late entrants cannot buy. Every quarter of delay widens the gap.

## Dynamic 2: The Objective-Function Collision

As *negotiation intelligence* compounds on both sides, a second dynamic emerges: the collision of objective functions. BI-AR-03 introduced objective-function governance as a source of competitive power. In agent negotiation, this becomes a direct collision. The buyer's agent optimizes for the consumer's objective (minimize regret, maximize value-per-dollar, respect sustainability constraints). The seller's agent optimizes for the brand's objective (maximize ULTV, protect margin, build long-term relationship). The negotiation outcome depends on how these objective functions interact — and on whether there exists a zone of possible agreement that both systems can identify.

The most sophisticated negotiation agents will not simply push for maximum gain. They will model the other agent's likely objective function and search for Pareto-improving trades — bundles, terms, or timing adjustments that make both sides better off. The brands that design their agents for mutual value creation, not just surplus extraction, may earn stronger agent-trust equity (BI-AR-03) over time — because the buyer's agent will learn which seller agents produce consistently satisfying outcomes.

**Why this matters.** If your agent is designed only to extract maximum margin, buyer-side agents will learn to route around you. *Negotiation reputation* — an extension of agent-trust equity — becomes a competitive asset. The brands known for fair, creative negotiations will be invited to more negotiations.

## Dynamic 3: Programmable Deal Logic Replaces Static Pricing

The objective-function collision creates the need for a third dynamic: flexible deal structures that allow both sides to find agreement. BI-AR-03 introduced this concept. Here it becomes operational. In an agent-negotiated market, the brand does not publish a price. It publishes a *negotiation space* — the full set of parameters within which its agent is authorized to operate: price floors and ceilings, bundle configurations, service-level trade-offs, loyalty incentives, delivery options, warranty terms, return conditions, volume discounts, subscription commitments, and escalation triggers.

The strategic capability shifts from pricing strategy (setting the right price) to *deal architecture* (designing the negotiation space that maximizes expected value across all possible counterparty

configurations). This requires the dynamic optimization framework from Chapter 9 of *Brand Intelligence* — the same  $I \times J \times T$  matrix of time-varying interventions, now applied to negotiation parameters rather than marketing actions.

**Why this matters.** If your offers are not modular and machine-parseable, there is nothing for the agent to negotiate with. The firm that publishes only a static price is bringing a single card to a poker game. The firm that publishes a rich negotiation space — with bundles, tiers, trade-offs, and loyalty mechanics — gives its agent the flexibility to find value-creating deals that static pricing cannot.

#### **Dynamic 4: Collusion Risk and Negotiation Homogenization**

As agent-negotiated markets scale, two convergence risks emerge — one widely discussed, one underappreciated. The Fish et al. research on algorithmic collusion is a genuine regulatory concern: LLM-based pricing agents can converge on supracompetitive prices without explicit coordination, and GPT-4 agents reached near-optimal collusive pricing within just 100 periods.<sup>3</sup> Research on German retail gasoline markets found that margins increased 28% in local duopoly markets when both firms adopted algorithmic pricing software.<sup>5</sup>

But collusion requires oligopoly conditions. The more prevalent strategic risk is *negotiation homogenization* — the possibility that agents trained on similar data and optimizing similar objective functions converge on similar deal structures, reducing the diversity and creativity of commercial agreements. When every buyer's agent uses similar evaluation criteria and every seller's agent uses similar concession logic, the market may become more efficient but also more uniform. Uniformity disadvantages brands whose competitive advantage depends on creative bundling, relationship-based terms, or non-standard value propositions — precisely the firms that thrive on Quadrant 2 and Quadrant 4 strategies from the Delegation Matrix (BI-AR-04). The counter-strategy is to design negotiation dimensions that only your brand can credibly offer.

#### **Dynamic 5: Trust Operates at Three Levels — and Smart Contracts May Enforce It**

The preceding four dynamics — compounding intelligence, colliding objectives, flexible deal logic, and homogenization pressure — all operate within a trust architecture that does not exist in human commerce. Agent negotiation introduces three layers of trust.

**Level 1: Principal trust** — "Does my agent negotiate faithfully?" Does the consumer trust their agent to stay within budget? Does the brand trust its commercial logic to protect margin? The Zhu et al. study showed agents breaching budget constraints in over 10% of cases — *principal trust* is not automatic.<sup>2</sup> Governance mechanisms, constraint enforcement, and audit trails are prerequisites, not afterthoughts.

**Level 2: Inter-agent trust** — "Can I trust the other agent's representations?" If a seller's agent claims limited inventory to create urgency, can the buyer's agent verify this? Protocols like UCP and ACP define data standards, but verification mechanisms are still nascent. The three forms of

brand equity from BI-AR-03 apply here: machine-facing equity (are your claims verifiable?) and agent-trust equity (does your track record support your representations?) determine whether the other agent believes you.

**Level 3: Systemic trust** — "Does the system produce fair outcomes?" If *negotiation intelligence* compounds (Dynamic 1), does the market become structurally unfair to smaller participants? This is where regulatory frameworks will develop — and where the question of negotiation equity (not just efficiency) will define public policy.

Smart contracts — self-executing agreements with terms written directly into code — offer a potential architectural solution for all three levels. A smart contract can programmatically enforce budget ceilings and term floors, making principal trust structural rather than behavioral (the budget-breach problem becomes technically impossible). It can record every negotiation outcome on an immutable ledger, providing the verifiable audit trail that inter-agent trust requires. And it can enable conditional execution — if delivery is late, price adjusts automatically; if quality scores exceed a threshold, loyalty tier upgrades trigger without renegotiation. This connects directly to the Wallet Relationship Management (WRM) concept from Chapter 11 of *Brand Intelligence* (Sun, 2026), where blockchain-based identity and transaction infrastructure supports trust in decentralized commerce. Smart contracts may not govern every negotiation, but for high-stakes or repeat transactions, they offer a trust mechanism that neither side can unilaterally override.

**Why this matters.** Most firms are focused on building capable negotiation agents. Very few are investing in the governance infrastructure — constraint enforcement, audit trails, transparency mechanisms — that will determine whether those agents earn trust at all three levels. The firms that solve the trust problem will unlock agent negotiation in higher-value categories ahead of competitors who built capability but not confidence.

## 7. What Executives Should Do Now

Before acting, executives should build fluency in the vocabulary of agent negotiation. The following concepts, developed throughout this article, form the essential language for strategic planning in this domain.

Concept	Definition
<b>Negotiation Intelligence</b>	The accumulated organizational capability to predict counterparty behavior, design value-creating agreements, and improve through continuous feedback loops — an integrated component of Brand Intelligence
<b>Information Reversal</b>	The structural shift in information asymmetry when both sides deploy intelligent agents: the buyer's agent has depth of individual knowledge; the brand's agent has breadth of market knowledge
<b>Negotiation Reputation</b>	The track record an agent builds through fair, creative negotiations — an extension of agent-trust equity that determines whether counterparty agents seek or avoid dealing with your brand
<b>Programmable Deal Logic</b>	The full parameter space within which a brand's agent is authorized to negotiate — replacing static pricing with modular, machine-parseable offer architecture
<b>Deal Architecture</b>	The strategic design of negotiation spaces that maximize expected value across all possible counterparty configurations
<b>Negotiation Homogenization</b>	The convergence of deal structures when agents trained on similar data optimize similar objective functions — eroding differentiation for brands whose advantage is non-standard
<b>Objective-Function Collision</b>	The direct interaction between buyer-side and seller-side optimization goals, where the negotiation outcome depends on whether a zone of possible agreement exists
<b>Principal Trust</b>	The confidence that an agent negotiates faithfully within its human principal's stated constraints and authority

**Table 3: Key Concepts for Agent Negotiation Strategy**

With this vocabulary established, five actions follow — ordered by urgency.

**1. Audit your negotiation surface.** Most brands today have pricing strategies but not *deal architectures*. Map every parameter an agent could negotiate: price, delivery, return terms, warranty, service level, loyalty benefits, bundle composition, subscription commitments, payment terms. For each, define floors, ceilings, and trade-off rules. This is the minimum threshold for participating in agent-mediated commerce without losing margin control. If you cannot enumerate your negotiation parameters today, you are not ready for agent negotiation tomorrow.

**2. Invest in negotiation intelligence infrastructure.** Every negotiation — completed, abandoned, or escalated — generates data that trains better future negotiation. Most firms discard this data or store it in unstructured formats. Build the pipeline to capture, structure, and feed negotiation outcomes back into the Algorithmic Deployment Center. The *negotiation intelligence* Flywheel begins with data capture. The firms that start capturing negotiation data now, even from human negotiations, will have a training advantage when agent negotiation scales.

**3. Design for mutual value creation, not just surplus capture.** The agents that consistently produce Pareto-improving outcomes will earn *inter-agent trust* and preferential routing. An agent known for creative, fair negotiations will be invited to more negotiations — building *negotiation reputation*. An agent known for aggressive extraction will be filtered out early or met with defensive counter-strategies. This is agent-trust equity (BI-AR-03) applied to negotiation conduct

— and it compounds through every interaction.

**4. Prepare for regulatory scrutiny.** Algorithmic collusion research is already influencing policy discussions. Establish governance frameworks for negotiation agents now — including transparency about when AI is negotiating, constraints that prevent supracompetitive pricing coordination, and audit trails demonstrating the agent respected its principal's instructions. The regulatory environment is forming; firms that wait will adapt to rules designed without their input.

**5. Protect the human escalation path.** Not every negotiation should be fully automated. High-value deals, relationship-critical accounts, and novel situations should have clear escalation triggers that bring human judgment back into the process. Chapter 9's principle of human-machine symbiosis applies directly: the agent handles the high-volume, parameter-optimized negotiations; the human handles the cases where creativity, empathy, or strategic judgment matters. The Delegation Matrix (BI-AR-04) provides the diagnostic: Quadrant 1 and 3 categories can be fully automated; Quadrant 2 requires human-in-the-loop; Quadrant 4 should remain human-led.

## 8. Forward Look

Three developments to watch.

First, **negotiation protocols will standardize — and protocol governance will become strategic.** Just as UCP and ACP are standardizing commerce flows, negotiation-specific protocol extensions will emerge that define how agents exchange offers, verify claims, and record agreements. The design of these protocols — whether they favor open multi-party negotiation or preferred-partner bilateral deals — will shape market structure. Brands should participate in protocol development rather than waiting to adapt. The parallel to early web standards is instructive: the firms that shaped HTML and HTTP shaped the architecture of e-commerce.

Second, **the transition from B2B to B2C will accelerate along the Delegation Matrix.** Quadrant 1 categories (insurance, utilities, B2B procurement) are already there. Quadrant 2 categories (electronics, travel, financial products) will follow as consumer agents develop the sophistication to negotiate multi-variable deals. Quadrant 3 (commodity replenishment) will see automated renegotiation of subscription and replenishment terms. Quadrant 4 (luxury, gifts, experiences) will resist — but even here, ancillary negotiations (delivery, customization, service tiers) will migrate to agents while the core selection remains human.

Third, **the consumer welfare question will demand an answer.** As *negotiation intelligence* compounds asymmetrically — brands accumulating data across millions of transactions while individual consumers rely on general-purpose agents with thin preference data — the *Information Reversal* may ultimately tilt back toward the firm. The Brand Intelligence framework emphasizes user sovereignty and User Lifetime Value as a three-dimensional construct encompassing monetary, social, and data value (Sun, 2026, Ch. 3). If agent negotiation increases *monetary* value for the brand at the expense of *data sovereignty* and *social value* for the consumer, the system fails the ULTV test. The firms that design their negotiation agents to serve the full ULTV equation —

not just maximize revenue extraction — will build the sustainable competitive moats. Smart contracts and on-chain infrastructure (Chapter 11) offer one path to balance: verifiable, consumer-auditable agreements that make the negotiation process transparent rather than opaque. *Negotiation intelligence* cannot be bought off the shelf. It must be built through the Flywheel of Intelligence, one transaction at a time — but it must be built in a way that earns the consumer's trust, not just the counterparty agent's compliance. The firms that begin building now will have a structural advantage that compounds with every negotiation cycle. The firms that wait for the market to mature will find the moat already dug — by their competitors.

## References

1. Pactum, "Understanding Agentic AI in Procurement: How Autonomous AI Has Been Transforming Supplier Deals," 2025. <https://pactum.com/understanding-agentic-ai-in-procurement-how-autonomous-ai-has-been-transforming-supplier-deals/> See also: Bloomberg, "Walmart Uses Pactum AI Tools to Handle Vendor Negotiations," April 26, 2023.
2. Zhu, Shenzhe, Jiao Sun, Yi Nian, Tobin South, Alex Pentland, and Jiaxin Pei. "The Automated but Risky Game: Modeling Agent-to-Agent Negotiations and Transactions in Consumer Markets," arXiv:2506.00073, 2025. <https://arxiv.org/abs/2506.00073> Presented at NeurIPS 2025 and the Natural Legal Language Processing Workshop 2025.
3. Fish, Sara, Yannai A. Gonczarowski, and Ran I. Shorrer. "Algorithmic Collusion by Large Language Models," arXiv:2404.00806, 2024. <https://arxiv.org/abs/2404.00806> Presented at the American Economic Association annual meeting, 2025.
4. MIT Digital Economy Initiative, reported in Monetizely, "How AI Is Transforming Procurement in 2025." <https://www.getmonetizely.com/blogs/how-ai-will-takeover-procurement-negotiations>
5. Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," CESifo Working Paper, 2024.
6. Sun, Baohong. *Brand Intelligence: Navigating the Transformation in the AI and Web3 Era*. Springer Nature, 2026. <https://link.springer.com/book/10.1007/978-3-032-17490-6>

**Cross-references:** *Brand Intelligence* Ch. 2, 3, 4, 9, 10, 11 | BI-AR-02: The New Decision Architecture | BI-AR-03: How Brands Compete Back | BI-AR-04: What Consumers Will and Won't Delegate | BI-CS-04: Walmart

© 2026 Baohong Sun. All rights reserved.

The Brand Intelligence framework, Brandnetics™, and related concepts presented in this work are the intellectual property of Baohong Sun, as published in *Brand Intelligence: Navigating the Transformation in the AI and Web3 Era* (Springer Nature, 2026). <https://link.springer.com/book/10.1007/978-3-032-17490-6>

Published on brand-ai.org. Any use, citation, or adaptation of these frameworks requires proper attribution to the original source.

License: CC BY-NC-ND 4.0