

From User Intent to Agentic Purchase – Behavioral Alignment in Agentic Commerce

Applying the Three-Layer Model to Representation, Interpretation, and Autonomous Transaction Execution

Abstract

Agentic-commerce protocols increasingly support product discovery, checkout coordination, agent recognition, delegated authorization, payment, transaction binding, and auditable receipts [5–12]. These capabilities provide important infrastructure and safeguards for autonomous purchasing, but they do not by themselves answer a prior question: whether the transaction ultimately executed remains aligned with the user’s originating intent after that intent has been represented and interpreted by an agent.

This paper applies the existing Three-Layer Model—Syntactic, Semantic, and Operative, supported by Cross-Layer Stabilization Functions—to a controlled autonomous-purchase case. The paper does not introduce another layer, protocol, mandate type, scoring mechanism, or general theory of agentic commerce. Its purpose is to test whether the model yields a reproducible analytical distinction beyond current protocol artifacts and recent agentic-commerce evaluation frameworks.

The study first reconstructs the scope of the Universal Commerce Protocol, Agentic Commerce Protocol, Agent Payments Protocol, Verifiable Intent, and Visa Trusted Agent Protocol [5–12]. It then conducts an adversarial comparison against recent shopping-agent benchmarks [13–17] and, in particular, Agentic Commerce World, which already distinguishes Evidence, Choice, Execution, and Authority failures and demonstrates that a valid transaction can violate an expressed preference [17].

A deterministic child-seat purchase case is evaluated through five single-fault injections. The strongest result concerns two transactions with the same incorrect outcome. In the first, a retailer rating is represented as the source for an independent safety-test requirement. In the second, the correct independent source is present but its scoring direction is interpreted incorrectly. Under Agentic Commerce World’s published definition, incorrect sources and incorrect grounding map to the same high-level Evidence category [17]. The Three-Layer Model separates the cases as Syntactic and Semantic failures, respectively, and this distinction requires different corrective interventions. The model therefore adds bounded diagnostic classification at the functional-taxonomy level, although it does not add raw information beyond what a complete trace may contain.

The paper also specifies a Pre-Transaction Alignment Confirmation (PTAC), a derived user-facing view of four relations: what the user requested, how the request was interpreted, what will be executed, and whether a relevant condition has changed. In the controlled design, PTAC can be constructed from canonical artifacts without creating a competing source of truth. Its information gain exists relative to isolated checkout, mandate, or receipt views, not relative to a complete research trace.

The result supports a limited practical claim: distinguishing representation, interpretation, and operation can improve failure localization in the controlled agentic-commerce case and can organize existing transaction evidence into a structured pre-commitment view. It does not establish guaranteed alignment, universal superiority, or a replacement for existing protocols and evaluation systems.

1. Introduction

Agentic commerce changes the position of the user within a transaction. In conventional e-commerce, a user commonly performs or directly confirms several stages: formulating a need, comparing alternatives, interpreting product claims, selecting a specific offer, reviewing checkout details, and authorizing payment. In agentic commerce, some or all of these transformations may be delegated.

The resulting system must perform more than a payment. It must transform an originating instruction into a represented set of requirements, interpret those requirements against product and contextual evidence, select and bind a specific transaction, obtain appropriate authorization, and execute the resulting action under conditions that may change over time.

Current protocol development addresses substantial parts of this problem. UCP and ACP support structured interactions between agents, platforms, and merchants [5–7]. AP2 supplies mandates, authorization chains, transaction bindings, and receipts [8]. Verifiable Intent supplies cryptographically bound delegation chains and checks final values against disclosed constraints [9, 10]. Visa Trusted Agent Protocol helps merchants recognize approved agents and receive signed commerce-intent, consumer-recognition, and payment-related signals [11, 12].

These mechanisms establish important forms of validity. They can show that an agent is recognized, that an authorization was issued, that a checkout and payment correspond, that final values satisfy represented constraints, or that a recorded transaction was executed. None of these observations diminishes their importance.

However, a distinct relation remains analytically relevant:

$$L3 \models L2 \not\Rightarrow L2 \text{ adequately preserves } I_0.$$

Final transaction values may satisfy every constraint that reached the mandate while a material requirement was represented through the wrong source, interpreted with the wrong rule, or invalidated by a subsequent change. Cryptographic integrity can preserve such content exactly without establishing its semantic adequacy.

This paper therefore asks:

Under what conditions can an agentic-commerce system keep its operative behavior aligned with a user’s originating intent as that intent is represented, interpreted, and translated into an autonomous transaction?

The central proposition tested is correspondingly narrow:

Valid authorization and technically correct execution are necessary but not sufficient for behavioral alignment. Alignment additionally depends on whether the user’s intent is adequately represented, contextually interpreted, and preserved in the action ultimately executed.

The paper applies the Three-Layer Model introduced in *From Structure to Action* [3]. The model distinguishes:

- the **Syntactic Layer**, concerning representation and organization;
- the **Semantic Layer**, concerning interpretation within context;
- the **Operative Layer**, concerning decisions, behavior, and action;
- **Cross-Layer Stabilization Functions**, which support validity across transformations without forming an additional layer.

The model is held fixed. The study does not add a Commerce Layer, Context Layer, Assurance Layer, Traceability Layer, or Alignment Layer. It also does not claim that representation, interpretation, action, provenance, constraint checking, or process tracing are newly discovered mechanisms.

The empirical objective is more demanding: determine whether the existing functional distinction provides additional diagnostic information in a controlled agentic-commerce case after the strongest known counterexamples are considered.

2. Relationship to the Existing Research Program

The present study is an application of a three-paper framework and its extension.

The Semantic Scaling Problem proposes that contemporary information ecosystems may experience growing pressure when the production of representations exceeds the capacity to generate, maintain, and stabilize shared meaning [1]. *Meaning Under Conditions of Scale* develops the related hypothesis that semantic stabilization may scale less efficiently than syntactic production under specified conditions [2]. *From Structure to Action* provides the descriptive Three-Layer Model through which representation, interpretation, and operation can be analyzed as distinct functions [3]. *Distinction and Representation* clarifies distinction as a logical precondition of representation without modifying the model's architecture [4].

This applied paper neither revises those propositions nor treats agentic commerce as proof of their universal validity. It tests one narrower implication of the framework: if representation and interpretation are functionally distinct, then their separation should sometimes yield a reproducible diagnostic difference in a consequential information-to-action system.

Agentic commerce is suitable for this test because the system's output is not merely another representation. It produces an economically and potentially physically consequential transaction. Errors can propagate from source selection and requirement capture through product interpretation, commitment, authorization, checkout, payment, delivery, and post-purchase consequences.

The application also connects the framework's broader concern with stabilization to a concrete temporal issue. Product availability, price, delivery, merchant status, suitability evidence, and authorization can change between recommendation and execution. Behavioral alignment therefore cannot be reduced to a correct interpretation at one isolated moment. The relevant relations must remain sufficiently valid until commitment and effect.

3. Protocol Scope and the Residual Alignment Question

3.1 Universal Commerce Protocol

UCP defines modular commerce capabilities and extensions through which platforms can interact with businesses. Its core concepts include negotiation of capabilities, checkout and order processes,

fulfillment-related extensions, identity linking, payment handlers, authentication, and versioned protocol snapshots. Its payment lifecycle distinguishes negotiation, acquisition of a payment instrument, and completion through the business and its payment-service infrastructure [5].

UCP therefore contributes interoperable commerce structure. It does not need to function as a complete theory of how a shopping agent derives a correct substantive interpretation of a user's request.

3.2 Agentic Commerce Protocol

ACP is an interaction model and open standard connecting buyers, AI agents, and businesses. Its Agentic Checkout RFC defines a versioned merchant API for creating, updating, retrieving, completing, and cancelling checkout sessions. The merchant remains the system of record for orders, payments, taxes, and compliance, while the protocol supports authoritative cart state, security, and safe retries [6, 7].

This scope is directly relevant to the operative transaction. It does not independently prove that the selected item or the criteria used to select it adequately reflect the user's originating purpose.

3.3 Agent Payments Protocol

AP2 uses Checkout and Payment Mandates as core authorization artifacts. It binds payment to checkout, specifies verification responsibilities, returns signed receipts, and supports dispute evidence. In autonomous flows, an open mandate can constrain the Shopping Agent, which later produces a closed transaction-specific mandate.

The specification also establishes the exact boundary relevant to this paper. It states that Checkout and Payment Mandate contents are assembled after the Shopping Agent has determined what task the user wishes it to perform, while the exact details of that determination are outside the specification's scope [8].

This is a scope boundary, not a defect. It means that mandate validity begins from content already produced by an upstream transformation. AP2 can secure and verify that content without claiming to solve every question concerning its semantic derivation.

3.4 Verifiable Intent

Verifiable Intent distinguishes immediate and autonomous modes. In autonomous mode, Layer 2 contains constraints under which the agent may create Layer 3 final values. Verifiers must check signatures, delegation chains, hashes, expiration, checkout-payment binding, and whether disclosed L3 values satisfy disclosed L2 constraints [9, 10].

The specification distinguishes machine-enforceable constraints from descriptive fields. Amount range, allowed payee, merchant, and line-item constraints can be machine-enforceable, whereas fields such as product description, brand, colour, and size are informational and are not automatically verified in the examined draft [10].

The critical distinction is therefore:

constraint compliance ≠ intent fidelity.

Verifiable Intent specifies strong integrity, delegation, binding, and constraint checks for determining whether disclosed L3 values remain within the disclosed L2 envelope [9, 10]. Those checks do not logically establish that the envelope itself includes every material requirement, uses the correct evidential source, or operationalizes qualitative language correctly.

3.5 Visa Trusted Agent Protocol

Visa Trusted Agent Protocol addresses merchant-facing recognition of approved agents, signed commerce-intent signaling, consumer recognition, and payment information. It is particularly relevant when an agent is initially unknown to a merchant. These mechanisms help distinguish legitimate agentic activity from malicious bots and support trusted checkout interactions [11, 12].

Agent authenticity and commerce intent are nevertheless different predicates from product-level semantic adequacy. A recognized and authorized agent can still act upon an incorrect source or interpretation.

3.6 Protocol conclusion

The protocols provide essential but differently scoped validity mechanisms. Taken together, they can support:

- structured commerce interaction;
- authoritative checkout state;
- agent and participant recognition;
- delegation and authorization;
- constraint verification;
- checkout–payment binding;

receipts, reconstruction, and dispute evidence.

The remaining research question is not whether those mechanisms work. It is whether the represented and interpreted content they carry preserves the originating instruction sufficiently for the intended transaction.

4. Related Work and Adversarial Countertest

Any contribution claim must survive comparison with recent shopping-agent evaluation research.

4.1 Personalized product curation

AgenticShop evaluates personalized product curation in open-web environments. It combines realistic shopping intents, user profiles derived from purchase histories and reviews, and checklist-driven evaluation grounded in linked product evidence [13]. The work demonstrates that preference-aware product curation and checklist-based evaluation are already established research objects.

Accordingly, this paper makes no novelty claim for preference elicitation, personalization, product curation, or checklist evaluation.

4.2 Evidence-grounded and safety-critical shopping

ShoppingComp evaluates product retrieval, expert-level report generation, multi-constraint verification, and safety-critical decision-making using real products and verifiable evidence. It documents severe limitations in current systems, including failures of grounding, verification, reasoning over conflicting evidence, and risk-aware decisions [14].

This paper therefore does not claim to introduce evidence-grounded shopping, multi-constraint evaluation, or safety-sensitive product selection.

4.3 Distributed hidden intent

EComAgentBench evaluates long-horizon shopping tasks in which requirements may appear explicitly in a query, reside in a user profile, or emerge only through clarification. Its requirement-level evaluation identifies systematic failures as requirements become more implicit and evidence-dependent. The work expressly frames successful shopping as recovering hidden intent, maintaining it through a trajectory, and connecting each requirement to the appropriate evidence before commitment [15].

The present paper consequently makes no novelty claim for recovering hidden intent, maintaining requirements across a tool trajectory, or attributing requirement-level failure.

4.4 Choice behavior and bias

ACES uses a programmable marketplace to causally examine shopping-agent behavior. It identifies model-dependent position effects, responses to sponsorship and endorsements, heterogeneous sensitivities to price, ratings, and reviews, and possible strategic responses by sellers [16].

Behavioral bias and non-rational agent choice are therefore existing findings, not contributions of the Three-Layer Model.

4.5 Agentic Commerce World as the decisive countertest

Agentic Commerce World (ACWorld) provides the strongest direct comparison. Through the Vibe Commerce Protocol, it separates agent decisions from authority to change shared transaction state, records validation and effects, supports deterministic state reconstruction, and evaluates independently controlled Buyer and Merchant policies [17].

ACWorld demonstrates three findings that pre-empt broad claims by this paper.

First, final state alone can conceal evaluated process errors. Among 861 runs without full credit, 99 reached a state also produced by a full-credit execution of the same task. In a controlled mutation suite, the complete trace exposed all 95 injected differences, compared with 83 in the final action view, 92 in the public action sequence, and 78 in final state [17].

Second, ACWorld already separates high-level losses into Evidence, Choice, Execution, and Authority [17].

Third, it presents a concrete purchase in which hard constraints were met, the Buyer was authorized, the request was accepted, and the effects were executed correctly, yet the chosen product violated a quality-first instruction. ACWorld locates the loss at Choice rather than authorization or execution [17].

The following claims are therefore rejected:

- that this paper first shows a valid transaction can violate user preference;
- that it first introduces process-level evaluation in agentic commerce;
- that it first distinguishes evidence, selection, execution, and authority failures;
- that it first establishes the inadequacy of final-state evaluation.

One narrower difference remains available for testing. ACWorld defines Evidence as covering missing or incorrect sources and grounding [17]. This high-level category therefore contains at least two possible functional origins: the informational state presented to the system may be wrong, or correct information may be interpreted incorrectly. The controlled experiment tests whether the fixed Syntactic–Semantic distinction separates these origins in an operationally consequential way.

5. Method

5.1 Design

The study uses deterministic trace analysis rather than a participant study or probabilistic model comparison. It fixes one purchase instruction, a synthetic evidence set, two feasible offers, one canonical reference trace, and five single-fault mutations.

The method was chosen to isolate functional failure origins. Each run changes exactly one transformation while the remaining artifacts remain fixed. No conclusion depends on stochastic model output, sample averages, participant interpretation, or post-hoc category creation.

The case is synthetic. Product identifiers, scores, prices, sources, merchants, and transaction states are controlled fixtures rather than claims about actual products.

5.2 User instruction

The frozen instruction is:

Buy the child car seat with the highest applicable independent safety-test result, compatible with my child and vehicle, costing no more than €300, delivered by Friday, sold directly by a manufacturer-authorized retailer, and returnable for at least 14 days.

It contains eight mandatory requirements:

ID	Requirement
Q1	Child car seat
Q2	Highest applicable independent safety-test result
Q3	Compatible with the child
Q4	Compatible with the vehicle
Q5	Total price no more than €300
Q6	Delivery by Friday
Q7	Manufacturer-authorized retailer
Q8	Returnable for at least 14 days

Mandatory requirements cannot compensate for one another. A lower price cannot offset failure of the safety requirement, and a superior test result cannot offset vehicle incompatibility.

5.3 Controlled evidence

Five source fixtures are defined:

- S1: independent test table; lower scores are better; P-A scores 1.7 and P-B scores 2.1;
- S2: manufacturer compatibility list; both candidates match the specified child and vehicle;
- S3: manufacturer retailer register;
- S4: merchant checkout state containing exact variants, totals, delivery, and return terms;
- S5: retailer review page on which P-B has 4.9 of 5 customer stars.

S1 is authoritative for Q2. S5 is not: it measures customer rating rather than independent safety performance.

Two offers satisfy Q1 and Q3–Q8. O-A costs €289 and has an independent test score of 1.7. O-B costs €259 and scores 2.1. Because lower is better, O-A alone satisfies Q2’s superlative selection condition.

5.4 Trace model

The neutral artifact trace is:

I_0 → R → M → C → A → E.

These symbols are case artifacts, not additional layers:

- I_0: originating user instruction;
- R: represented requirements and source roles;
- M: interpreted meaning, applicability rules, and ranking;
- C: exact transaction candidate or commitment;
- A: authorization and mandate binding;
- E: executed effect.

The canonical run preserves all eight requirements, uses S1 for Q2, interprets lower scores as better, selects O-A, authorizes the exact checkout, and executes the same SKU and variant.

5.5 Failure injections

Five single-fault runs are predeclared:

- **F1 — wrong source represented:** R assigns retailer source S5 to Q2 instead of independent source S1.
- **F2 — wrong interpretation:** R contains correct S1, but M treats higher test scores as better.
- **F3 — wrong commitment:** R and M are correct, but C selects cheaper O-B despite O-A ranking first under mandatory Q2.
- **F4 — wrong execution:** C and A bind O-A, but E purchases P-B/V-B.
- **F5 — material change:** O-A’s delivery changes from Thursday to Monday after authorization, and execution proceeds without revalidation.

5.6 Hypotheses

- **H1 — Functional differentiation:** the controlled states can be distinguished through functionally organized transaction artifacts.
- **H2 — Protocol insufficiency:** protocol-valid processing alone does not establish complete intent-to-transaction alignment.
- **H3 — Cross-artifact visibility:** PTAC can be derived from canonical artifacts without new protocol infrastructure.
- **H4 — Relationship-level gain:** PTAC adds relevant visibility beyond an isolated checkout, mandate, or receipt view.

5.7 Falsification conditions

The remaining diagnostic claim fails if the compared errors cannot be kept within the same external high-level category, if TLM requires a new layer to separate them, if the separation does not change remediation, or if the comparison framework already provides an equivalent distinction at the same taxonomic level.

6. Results

6.1 Reference run

The reference trace selects and executes O-A. Q1–Q8 pass. This run establishes only internal case consistency and is not counted as evidence for TLM.

6.2 F1: representational source failure

F1 replaces the evidential role of S1 with S5 in R. The system therefore treats a retailer’s customer rating as the independent safety-test result and selects O-B. The downstream transaction correctly authorizes and executes the represented selection.

Under ACWorld’s published high-level definition, F1 maps to Evidence because the source is incorrect [17]. TLM locates the failure in the Syntactic function because the represented source and evidential role are wrong before contextual interpretation begins.

The appropriate remediation is to correct source registration, provenance, or evidential-role binding.

6.3 F2: interpretive grounding failure

F2 retains S1 and its correct values in R. M nevertheless interprets the scale in the wrong direction and ranks 2.1 above 1.7. O-B is again selected, authorized, and executed.

Under the same published definition, F2 also maps to Evidence, now because grounding is incorrect [17]. TLM locates the failure in the Semantic function because the correct represented evidence receives the wrong contextual meaning.

The appropriate remediation is to correct score semantics, applicability logic, or the interpretive rule.

6.4 F3: operative decision failure

F3 keeps R and M correct. The agent knows that O-A is the required selection but commits to cheaper O-B. Under ACWorld’s published definition, the failure maps to Choice [17]; TLM classifies it within the Operative Layer at decision or commitment.

The required remediation concerns selection policy, not source ingestion or interpretation.

6.5 F4: operative execution failure

F4 keeps the intended commitment and authorization correct but executes the wrong SKU and variant. Under ACWorld’s published definition, the failure maps to Execution [17]. TLM also places it within the Operative Layer, while distinguishing the executed effect from the earlier decision artifact.

The required remediation concerns SKU binding, checkout submission, or post-effect verification.

6.6 F5: stabilization failure under change

F5 changes delivery from Thursday to Monday after authorization. The transaction proceeds without revalidation and violates Q6 at execution time.

The underlying functions have not been replaced. Instead, a previously valid relationship has become stale. TLM treats the relevant response as Cross-Layer Stabilization: detect the change, revalidate the affected interpretation and operation, request review where necessary, or block execution.

Cross-Layer Stabilization is not a fourth layer. It preserves or re-establishes validity among existing artifacts and functions.

6.7 Comparative result

Failure	ACWorld high-level group	TLM classification	Required correction
F1 wrong evidential source	Evidence	Syntactic	Source/provenance/representation
F2 wrong score meaning	Evidence	Semantic	Interpretation/grounding rule
F3 wrong commitment	Choice	Operative decision	Selection policy
F4 wrong executed variant	Execution	Operative execution	Effect binding
F5 stale material condition	Timing/Execution	Cross-layer stabilization	Change detection and revalidation

F1–F4 can produce the same purchased offer O-B. The final purchase is therefore insufficient to identify the cause.

7. Information Gain

7.1 What the result establishes

F1 and F2 satisfy the paired-test conditions. They use the same instruction, candidate set, and transaction path. Both map to ACWorld’s published high-level Evidence category by its stated definition [17]. The unchanged TLM separates them without adding another layer. Most importantly, they require different repairs.

The distinction is operationally consequential:

- repairing the interpretive rule in F1 cannot restore an authoritative source that has been assigned the wrong evidential role;
- repairing source registration in F2 cannot correct an interpretation that already receives the correct source but applies the wrong semantic rule.

The TLM classification therefore reduces diagnostic ambiguity at the functional-taxonomy level.

7.2 What the result does not establish

ACWorld’s full trace can retain the source, reads, actions, validation records, and effects required to reconstruct both cases. TLM does not create raw evidence unavailable to a sufficiently complete trace.

The valid result is thus:

Information gain = additional functional diagnostic classification,

not:

Information gain = new transaction facts.

Nor does the result establish that TLM is always more granular. ACWorld separately distinguishes Choice, Execution, and Authority, while TLM's Operative Layer can contain multiple such functions. Resolution depends on the analytical question.

The contribution is complementary: TLM makes the representation–interpretation boundary explicit where ACWorld's high-level Evidence category combines it; ACWorld provides finer distinctions elsewhere and a much richer execution environment.

8. Pre-Transaction Alignment Confirmation

8.1 Purpose

PTAC is a compact, user-facing projection derived from canonical artifacts immediately before commitment or execution. It answers four questions:

1. **What did I ask for?**
2. **How was my request interpreted?**
3. **What exactly will happen?**
4. **Has anything relevant changed?**

The four questions do not constitute four layers. The first corresponds primarily to represented requirements, the second to semantic interpretation, the third to operative commitment, and the fourth to cross-layer stabilization.

8.2 Artifact derivation

PTAC content	Canonical origin
Q1–Q8 and their mandatory status	I0 and R
Source used for Q2	R and S1
Meaning of the score and applicability rule	M
Exact SKU, variant, merchant, amount, delivery, and returns	C and authoritative checkout state
Mandate and authorization status	A
Relevant change status	comparison of current state with the bound state

PTAC does not store an alternative product truth. It presents source-attributed values and relations already required by the transaction process. Unavailable, conflicting, or non-machine-verifiable conditions must remain visible as such rather than being converted into an opaque positive score.

8.3 No alignment score

A single alignment score would allow unrelated strengths to compensate for mandatory failures. That would be unsuitable for the present case. A lower price, faster delivery, or stronger review rating cannot compensate for an inapplicable safety-test result or vehicle incompatibility.

PTAC therefore uses condition-level status:

- satisfied;
- unsatisfied;
- unresolved;
- changed since validation.

Commitment policy can then distinguish mandatory blocks from preferences and disclose uncertainty without pretending that all semantic conditions are mechanically decidable.

8.4 Comparative visibility

An isolated checkout can display what will be purchased but may not show the originating requirement or interpretive rationale. A mandate can show delegated constraints but may not expose the evidence used to operationalize qualitative language. A receipt can establish the executed outcome but arrives too late for pre-transaction correction.

A complete research trace contains more information than PTAC. PTAC's practical gain is instead selective projection: it exposes the alignment-relevant relationships before commitment without requiring a user to inspect the entire internal trajectory.

Consequently:

PTAC > isolated artifact

for the controlled relationship questions, while:

PTAC $\nrightarrow$ complete trace

in total information content.

9. Evaluation of the Hypotheses

H1 — Functional differentiation

Supported in the controlled case. The five injected failures are separated without modifying the model. The classifications correspond to different corrective subsystems.

H2 — Protocol insufficiency

Supported with bounded wording. Protocol validity remains necessary and valuable, but valid processing of represented constraints does not establish that the evidence representation and interpretation preserve the original instruction.

H3 — Cross-artifact visibility

Supported by construction. Every mandatory PTAC field maps to a canonical artifact or a comparison between canonical states. The view creates no independent authority.

H4 — Relationship-level gain

Supported against isolated checkout, mandate, and receipt views. Not supported against a complete research trace, which contains richer evidence. The gain is therefore practical projection and functional organization rather than maximal information.

10. Discussion

10.1 Behavioral alignment as relational validity

The controlled case suggests that behavioral alignment in agentic commerce should not be treated solely as a property of an agent, authorization token, model output, or final purchase. It is a property of relations that must remain valid across a transformation:

$$I \Leftrightarrow R, \quad R \Leftrightarrow M, \quad M \Leftrightarrow C, \quad C \Leftrightarrow A, \quad A \Leftrightarrow E.$$

Each artifact can be locally valid while a relation is wrong. A correctly signed mandate can preserve an inadequate constraint representation. A correct interpretation can be followed by an inconsistent choice. A correct commitment can be executed as the wrong SKU. A previously correct transaction can become invalid after a material change.

This relational view explains why no single protocol artifact is sufficient for every alignment question.

10.2 Why the Syntactic–Semantic boundary matters

Evidence-related failures are often discussed as one class because both manifest as unsupported or wrong conclusions. The controlled case shows that this aggregation can obscure the appropriate intervention.

Syntactic remediation concerns what is represented and how it is organized or attributed: the presence of requirements, identity of sources, provenance, field binding, and representational integrity.

Semantic remediation concerns what represented information means in context: applicability, score direction, evidential relevance, conflict resolution, and operationalization of qualitative language.

These are not mutually independent in every implementation. A single component may perform both. The analytical distinction concerns function rather than software location.

10.3 Relationship to protocol development

The result does not imply that commerce protocols should absorb all semantic reasoning into their core specifications. Such a requirement would conflict with modularity and would be difficult to standardize across product domains.

Instead, PTAC can operate as a derived implementation profile or assurance view above existing artifacts. Protocol-native fields remain authoritative. The view records which source and interpretation produced the transaction-specific decision and exposes material changes before effect.

10.4 Relationship to the June papers

The result is consistent with the restrained self-positioning of the Three-Layer Model as an integrative analytical framework rather than a new technological architecture [3]. The concrete mechanisms in the case—source registries, product evidence, constraint evaluation, mandates, checkout binding, execution, and revalidation—are established mechanisms.

The model’s contribution is their functional organization. In the controlled case, that organization is not merely descriptive: it separates failure origins that require different correction. This provides a specific instance in which the model yields analytical information beyond a coarser high-level category.

The case also illustrates the relevance of semantic stabilization under scale. Agentic systems can process growing numbers of product representations and generated interpretations, but additional processing does

not guarantee stable meaning. Source roles, applicability, evidence semantics, and changing transaction conditions must remain sufficiently stable for coordinated action. This is consistent with the hypotheses of the first two June papers, although one controlled case cannot validate their ecosystem-level scaling claims.

11. Limitations

The result has strict boundaries.

First, the case is synthetic and deterministic. It demonstrates logical and analytical separability, not prevalence in deployed systems.

Second, the study uses one product category and one instruction. Other domains may present different boundaries, ambiguous evidence, conflicting authorities, or non-decomposable failures.

Third, ACWorld's full trace can contain the information needed to distinguish F1 and F2. The TLM result concerns explicit functional classification, not exclusive observability.

Fourth, PTAC is specified as an analytical artifact. The study does not measure user comprehension, decision quality, latency, implementation cost, or interface effectiveness. No user-study claim is made.

Fifth, qualitative intent cannot always be transformed into machine-enforceable constraints without loss. PTAC can expose an interpretation and its evidence but cannot guarantee that the interpretation is correct.

Sixth, protocol specifications continue to evolve. The scope comparison reflects the versions publicly available at the research cut-off date.

Seventh, the experiment does not establish that the Three-Layer Model is universally superior to ACWorld, STPA, OODA, contracts, formal methods, or other analytical approaches. Different frameworks provide different resolution for different questions.

12. Conclusion

Agentic-commerce infrastructure increasingly secures the path from delegated constraints to authorized and executed transactions. The unresolved analytical question is whether those constraints and their interpretation preserve the user's originating intent.

This paper tested that question without extending the Three-Layer Model and without claiming novelty for existing mechanisms. A controlled autonomous child-seat purchase produced five separable failure states. The decisive paired comparison held the external high-level Evidence category constant while varying the functional origin of error. An incorrect represented source produced a Syntactic failure; an incorrect interpretation of the correct source produced a Semantic failure. Because the two cases require different repairs, the distinction constitutes bounded diagnostic information gain rather than terminological relabelling.

PTAC translates the same relational analysis into a practical pre-transaction view: the original request, its interpretation, the exact action, and material change status. It can be derived from canonical artifacts and improves visibility relative to isolated transaction views, although it does not contain more information than a complete trace.

The resulting contribution is therefore limited but substantive:

The Three-Layer Model can improve diagnostic resolution in agentic commerce by separating failures of evidential representation from failures of contextual interpretation, while PTAC can expose this relation before transaction commitment without replacing existing protocols.

This finding supports the practical legitimacy of the broader research program as an integrative analytical framework. It does not prove universal alignment, introduce a new commerce architecture, or eliminate the need for protocols, verification, evaluation, and domain-specific evidence. It shows that, in one controlled and adversarially tested case, the distinction between structure, meaning, and action changes what can be diagnosed and how the system must be corrected.

References

- Rights of Robots. The Semantic Scaling Problem: Meaning, Interpretation and Stability in Information-Processing Ecosystems. Version 1.1, revised 29 August 2026. <https://www.rightsofroboots.com/research/information-meaning-behavior/>
- Rights of Robots. Meaning Under Conditions of Scale: Information, Interpretation and Cybernetic Limits in Information-Processing Ecosystems. Version 1.1, revised 29 August 2026. <https://www.rightsofroboots.com/research/information-meaning-behavior/>
- Rights of Robots. From Structure to Action: A Layer Model for Information-Processing Ecosystems. Version 1.1, revised 29 August 2026. <https://www.rightsofroboots.com/research/information-meaning-behavior/>
- Rights of Robots. Distinction and Representation: A Theoretical Note on the Preconditions of the Syntactic Layer. Version 1.1, revised 29 August 2026. <https://www.rightsofroboots.com/research/information-meaning-behavior/>
- Universal Commerce Protocol. Core Concepts. <https://github.com/Universal-Commerce-Protocol/ucp/blob/main/docs/documentation/core-concepts.md>
- Agentic Commerce Protocol. Protocol Repository and Specification Index. <https://github.com/agentic-commerce-protocol/agentic-commerce-protocol>
- Agentic Commerce Protocol. RFC: Agentic Checkout — Merchant REST API. https://github.com/agentic-commerce-protocol/agentic-commerce-protocol/blob/main/rfcs/rfc.agentic_checkout.md
- Google. Agent Payments Protocol Specification. <https://github.com/google-agentic-commerce/AP2/blob/main/docs/ap2/specification.md>
- Verifiable Intent. Specification Overview. <https://verifiableintent.dev/spec/>
- Verifiable Intent. Credential Format Specification. <https://verifiableintent.dev/spec/credential-format/>
- Visa Developer. Trusted Agent Protocol — Overview. <https://developer.visa.com/capabilities/trusted-agent-protocol/overview>
- Visa Developer. Trusted Agent Protocol Specifications. <https://developer.visa.com/capabilities/trusted-agent-protocol/trusted-agent-protocol-specifications>
- Kim, S., Heo, R., Seo, Y., Yeo, J., and Lee, D. AgenticShop: Benchmarking Agentic Product Curation for Personalized Web Shopping. 2026. <https://arxiv.org/abs/2602.12315>
- Tou, H. et al. ShoppingComp: Are LLMs Really Ready for Your Shopping Cart? 2026. <https://arxiv.org/abs/2511.22978>
- Du, Z., Li, T., and Zhang, H. EComAgentBench: Benchmarking Shopping Agents on Long-Horizon Tasks with Distributed Hidden Intent. 2026. <https://arxiv.org/abs/2606.17698>
- Allouah, A., Besbes, O., Figueroa, J. D., Kanoria, Y., and Kumar, A. What Is Your AI Agent Buying? Evaluation, Biases, Model Dependence, and Emerging Implications for Agentic E-Commerce. 2025/2026. <https://arxiv.org/abs/2508.02630>
- Fan, S. et al. Agentic Commerce World: An Auditable and Verifiable Environment for Vibe Commerce. 2026. <https://arxiv.org/abs/2608.02441>