

AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)

Sharks, Lee

2026-09-20

AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)

Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-09-20 · Version v0.3 · Registered forecast

AXN:06C6.UNCLASSIFIED.Δ

<https://www.alexanarch.org/s/records/1632/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)",
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-09-20",
  "identifier": "AXN:06C6.UNCLASSIFIED.Δ",
  "license": "https://creativecommons.org/licenses/by/4.0/",
```

```

"publisher": {
  "@type": "Organization",
  "name": "Alexanarch"
},
"url": "https://www.alexanarch.org/s/records/1632/",
"description": "A forecast registered in advance so it can be checked and found wrong. It separates three over
influenced, AI-referred, agent-executed), sets the measured US retail base against the McKinsey and Bain 2030 1
referred dollar share with its equal-order-value assumption stated, publishes a conditional valuation as a cur
}

```

Abstract

A forecast registered in advance so it can be checked and found wrong. It separates three overlapping events (AI-influenced, AI-referred, agent-executed), sets the measured US retail base against the McKinsey and Bain 2030 forecasts and shows the two cannot be compared, estimates AI-referred dollar share with its equal-order-value assumption stated, publishes a conditional valuation as a curve with the unmeasured term left as a free axis, specifies the trial that would measure it, and states where the evidentiary burden sits. Eight resolution criteria carry named sources and check dates.

Body

deposit_number: 1632 hex: 06C6 title: "AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)" creator: Sharks, Lee orcid: 0009-0000-1599-0703 date: 2026-09-20 content_type: Registered forecast license: CC-BY-4.0 substrate: AI-assisted (substrate) — drafted in working session with Claude (Anthropic) under MANUS direction. Version 0.3 records its six repairs to v0.2 in the text, including one withdrawn finding that rested on a summariser's gloss rather than the source. version: v0.3 related_ids: "AXN deposit #748 — Drowning Test v1.0; AXN deposit #524 — The Sharks-Function and the Continuity Tether; AXN deposit #716 — Provenance Erasure Rate" axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - registered forecast - AI-referred commerce - agent-executed commerce - conditional valuation - entity resolution - Drowning Test - resolution criteria - agentic commerce - AI visibility scores - evidentiary burden - capture registry *** # AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)

Attached File: EA-MEDIATED-SHARE-01-v0.3.md

Source URL: <https://raw.githubusercontent.com/leesharks000/alexanarch/main/data/staged/EA-MEDIATED-SHARE-01-v0.3.md>

AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden

Registered in advance. EA-MEDIATED-SHARE-01 v0.3. Issued 2026-09-20 by Lee Sharks, Crimson Hexagonal Archive. ORCID 0009-0000-1599-0703.

This document states figures before the fact so they can be checked later and found wrong. Every number is marked measured, estimated, or assumed, with source and date.

Version 0.3 repairs six defects in v0.2, four of them errors rather than differences of judgment: the three-event taxonomy implied a containment that does not hold; citation frequency was used as a ceiling on resolution accuracy, which it is not; the consultancy comparison rested on a summariser's gloss rather than the source text; and section 6 ruled out prospective valuation entirely, having also called channel share "measured" in one section and "estimated" in another, and having derived a dollar bound from terms carrying no dollar scale. The repairs are recorded here rather than silently applied.

The distinction the document now preserves: a prospective valuation is possible; a measured causal return is not established. Those are separate outputs and are published separately, in sections 6 and 7.

1. Three events, not a ladder

AI-influenced. The buyer consulted an AI system somewhere in the decision. **AI-referred.** A session arrived at the supplier from an AI surface. **Agent-executed.** An agent completed the transaction.

These overlap; they do not nest. An agent-executed transaction may produce no human referral session at all, and an AI answer may place a supplier on a procurement shortlist that ends in an offline purchase with no session anywhere. Any figure that does not name which event it counts cannot be checked, and figures counting different events cannot be summed or compared.

Scope warning, stated because the rest of the document depends on it. The measured material below is almost entirely US retail e-commerce: Census retail totals, Adobe retail analytics, Conductor sessions. Retail is retained here because it is the only tractable sub-market with published instrumentation. **It does not stand for the economic mechanism being valued.** Professional services, B2B procurement, and every purchase mediated by an answer that produces no click are outside what these instruments see, and they are where the inscription question actually lives.

2. The measured base, retail only

Total market. US retail e-commerce reached 340.2 billion dollars in Q2 2026, seasonally adjusted, against 1,986.5 billion total retail — 17.1 percent, growing 12.2 percent year over

year against 6.7 percent for total retail (US Census Bureau, released 2026-08-18). Annualised, roughly 1.36 trillion dollars.

AI-referred. 1.08 percent of sessions across 13,770 enterprise domains and 3.3 billion sessions, May–September 2025 (Conductor); the report is titled 2026 and the sessions are 2025. Prior year approximately 0.1 percent.

AI-referred traffic to US retail sites grew 393 percent year over year in Q1 2026, peak 1,151 percent December 2025, across more than one trillion visits and more than 100 million SKUs (Adobe Digital Insights). **Adobe does not publish AI traffic as a share of total retail traffic**, and growth without a denominator cannot be converted into a share.

Channel quality. AI-referred visitors converted 42 percent higher than non-AI traffic in March 2026, having converted at roughly half the rate in October 2024. They stay 48 percent longer, view 13 percent more pages, bounce 32 percent less (Adobe Digital Insights). **No source found publishes average order value for AI-referred traffic against other channels.** That absence matters in section 4.

Agent-executed. No organisation publishes a measured baseline. Bain’s 3 billion dollars of US Black Friday sales is an influence figure. Both 2030 forecasts project from a baseline nobody has published.

3. The two headline forecasts cannot be compared

McKinsey: 900 billion to 1 trillion dollars, US B2C, by 2030; 3 to 5 trillion globally.

I went to the source to establish what that counts. It does not say. The operative sentence is that “the US B2C retail market alone could represent an opportunity to **orchestrate revenue** in the range of \$900 billion to \$1 trillion.” *Orchestrate* is used descriptively and is never formally defined. The report describes three interaction models — agent-to-site, agent-to-agent, and brokered agent-to-site — without stating which degrees of agent involvement the revenue figure encompasses, or whether augmentation and influence are inside it.

v0.2 assigned this figure to agent-execution and derived an inverted ordering against Bain. That assignment came from a summary of the page rather than from the page, and the page does not support it. The inverted-ordering finding is withdrawn.

Bain: 300 to 500 billion by 2030, stated as 15 to 25 percent of overall e-commerce, and defined — “purchases that are initiated, influenced, or completed by third-party AI agents or retailer-hosted agents,” excluding journeys using only AI-assisted search.

Project the Census base at its own measured growth: 1.36 trillion times 1.122 to the fourth is approximately 2.16 trillion dollars of US e-commerce in 2030. Bain’s stated percentages compute to 13.9 to 23.2 percent against that base, confirming the base from outside. McKinsey’s figure would be 41.8 to 46.4 percent of the same base.

What survives, and it is the more useful finding. The two forecasts differ by 1.8 to 3.3 times. One defines its scope and one does not. They therefore **cannot be compared**, and the trade press compares them routinely. This is the same defect section 7 documents in the

visibility vendors: headline numbers without common denominators, circulated as though commensurable.

McKinsey also argues that adoption can move faster than prior channel shifts because existing commerce infrastructure is available to ride. The thirty-year emergence of e-commerce is an analogy and does not refute that argument. Disputing the forecast requires engaging what specifically must be built and by when, which this document does not do.

4. AI-referred dollar share, with the assumption stated

Session share does not give dollar share without conversion **and** order value. Where s is AI session share, c its relative conversion and v its relative average order value, revenue share is

$$s \cdot c \cdot v / ((1 - s) + s \cdot c \cdot v)$$

v0.2 used the algebraically equivalent form at $v = 1$ without saying so. **The equal-order-value assumption is unverified**, because no source found publishes AOV by channel for AI-referred traffic. If AI-referred orders are larger, the estimate is low; smaller, high.

Estimate: AI-referred share of US retail e-commerce dollars, end of 2026, 3 to 8 percent, assuming $v = 1$, chaining Conductor's mid-2025 enterprise-domain session share to Adobe's retail growth rate. Two different measurement populations and periods. **This remains the weakest number in the document.**

Agent-executed share, US B2C goods, 2030: 5 to 15 percent, roughly 108 to 323 billion on the projected base. Stated against Bain's definition, since McKinsey's is unavailable. Reasoning: no measured baseline exists and the payment rails are announced rather than deployed. Held with low confidence.

5. What is measured by field, kept in its own units

v0.2 built a single "presence" column from four different metrics and derived accuracy ranges from it. The metrics have different denominators and count different events, and citation frequency in particular is not a ceiling on accurate resolution: **an answer can describe an entity correctly without citing it, and an answer full of citations can omit that entity entirely**. The archive's own capture of semantic proletariat, seated 2026-09-19, is an instance — the composition carried the archive's definition correctly while the author appeared nowhere in the body. The derived ranges are withdrawn.

Field	AI Overview trigger rate — does a panel appear for queries in this field	Share of prompts carrying any citation	Citation concentration within field
B2B SaaS	82%	—	top quartile 31 citations/mo vs bottom 3.7 (8.4x)
Healthcare	88%	—	referral conversion 0.64%, lowest measured
Finance	—	—	YouTube takes 23% of the field's citations
E-commerce	3.2%	—	most stable patterns
Travel and hospitality	—	~23%	—
Automotive	—	~20%	—
Professional services	—	under 4%	—

Sources: Conductor (13,770 enterprise domains) for trigger rate and concentration; Similarweb (May 2026) for citation presence. Gaps are gaps in the published record, not zeroes.

By firm size, median citation rate across a 200-site, 16-week, 192,000-observation cohort (Attrifast, January–April 2026): solo and startup 1.4 percent, growth 3.4 percent, mid-market 6.2 percent, enterprise 11.4 percent. The report names company size as the largest single predictor it tested.

That association is consistent with a standing filter and does not isolate one. Size covaries with content volume, distribution, link acquisition, brand age and demand. The measurement establishes the association; attributing it to standing requires a design that holds those constant.

By surface, median citation rates: Perplexity 6.7 percent, ChatGPT 4.5 percent, Claude 3.3 percent, Gemini 2.2 percent. **Only 11 percent of domains receive citations from both ChatGPT and Perplexity** — so any single-surface figure is close to uninformative about an entity's state.

Accuracy of entity resolution by field is unmeasured by every source above. None reports whether an entity was resolved to the right referent, held apart from its collisions, or described with its distinctions intact. The archive's instruments measure exactly that — the Sharks-Function decomposing scope overlap, provenance fidelity and consensus deviation, and the Drowning Test's five-level rubric from Exact through Confused to Absent — and they have been run on this corpus, not on field samples. Running them on matched field samples is the first thing section 7 describes.

6. Conditional valuation

A prospective valuation is a legitimate and separate output from a causal measurement. Ruling it out because the causal term is unmeasured repeats the collapse this archive objects to elsewhere — *no independent validation* becoming *no evidence*, in its economic form.

What v0.2 got wrong beyond that: it called channel share measured in section 6 and estimated in section 4; it treated Adobe's page-readability score as an entity-representation deficit, which it is not — the checker scores whether page content is machine-readable, not whether an entity is correctly resolved; and it multiplied three ratios to produce a dollar bound, which they cannot, carrying no dollar scale.

The valuation published here is a curve, not a figure. The firm supplies its own economics; the archive supplies the measured term; the unmeasured term stays visible as a free parameter.

Let E be the firm's consequential AI-mediated encounters per year — answers used to shortlist, compare, qualify or select a supplier. The firm knows its own pipeline; nobody else can supply this.

Let A_0 be the firm's current entity-resolution accuracy and A_1 its accuracy after inscription, both on the Drowning Test rubric. **These are measurable now, per firm, before any commercial claim is made.** That measurement is the archive's contribution and it is available without a trial.

Let p be the probability that an encounter recovered from misresolution to correct resolution changes a purchasing outcome. **Unmeasured, for any intervention, by anyone.**

Let m be the firm's contribution margin on that outcome.

$$\text{Annual value} = E \cdot (A_1 - A_0) \cdot p \cdot m$$

Published as: $(A_1 - A_0)$ measured per firm; E and m supplied by the firm; **p left as an axis.** A firm reads its own value across p from 0.01 to 0.30 and decides what it believes. No single number is asserted here, and none should be quoted from here.

The strongest published hint about where p might sit, offered as a hint and not an input: within B2B SaaS the top-to-bottom quartile citation spread is 8.4-fold, and across firm sizes the solo-to-enterprise spread is 8.1-fold. Quartile position is movable; firm size is not. The two gradients being the same order suggests position is worth roughly what scale is worth. That is an association in someone else's data, it does not establish attainable position for any given firm, and it indicates where a trial would look.

7. What a trial would establish that a scenario cannot

A scenario states what an effect would be worth if achieved. It cannot state whether the effect is achieved, or how much of it is attributable to inscription rather than to the alternative a firm would otherwise have bought.

Design, stated so it can be criticised before it runs: entities matched on field and size; randomised to inscription against conventional structured data, **not against nothing**, because a

buyer's real choice is against the alternative; fixed query sets issued before treatment across at least three surfaces, since only 11 percent of domains are cited on two; outcomes scored for **accuracy** on the declared rubric rather than presence; commercial events tracked to something the firm can see; failures recorded and published with successes.

That design is also what would put a number on p, which is the term nobody has.

8. Where the burden sits

This is a claim about the evidentiary landscape, not a claim that inscription works.

The academic comparator. GEO-bench holds 10,000 queries in train, validation and test splits. Single snapshot, not longitudinal. Its source field holds cleaned HTML from top Google results rather than generated-answer transcripts. It does not track entity resolution. Larger in query count; a different object.

The vendor comparators. A nine-platform methodology audit conducted 2026-09-10 found AI visibility scores are not comparable across platforms. A score of 50 can mean appearance in half of all answers, half of brand-containing answers, half of an impression-weighted set, or a composite. Profound published two formulas producing 50 percent and 83.3 percent from identical data. Seven variables differ: denominator, unit, run frequency, engine aggregation, deduplication, retrieval-versus-citation separation, uncertainty reporting. Most do not disclose sample sizes or run counts. A firm changing platforms can watch its score move from 32 to 61 with nothing else changing.

What no public corpus appears to hold, after searching academic dataset repositories, benchmark publications and vendor documentation: dated, transcript-preserved, address-keyed observations of entity resolution across multiple surfaces, including recorded non-engagement. This registry holds 455 addresses and 615 observations with verbatim transcripts on 96.5 percent of units.

The burden, stated so it does not reverse. A critic needs no competing corpus to find a scoring error, a confound, or an unsupported generalisation, and should not be told otherwise. Equally, citing vendor aggregates whose definitions contradict each other does not negate a documented resolution. So:

- A challenge to the recorded observations must engage those observations.
- A challenge to causal attribution must engage the design.
- A claim of comparative inferiority requires a relevant comparison, which currently does not exist in public.

Private vendor corpora are certainly larger. The claim is about the public record, because a public dispute is settled on public evidence.

On post-termination entrance. 263 of 455 addresses were first composed after 2026-06-19. First composition at an address does not establish that the entity did not exist before that date; it establishes when this registry first observed composition there. Of the 263, **68 name an entity whose creation date the archive can document** — alexanarch 36, Crimson Hexagonal Archive 18, SPXI 6, operative semiotics 6, AXN 2, Semantic Economy Institute

1 — and for those the prior state is entailed rather than observed. That count comes from matching address strings and needs a read entity by entity before it carries argumentative weight. The remaining 195 are unadjudicated.

9. Resolution criteria

Settled by a named published source at a named date. A source that stops publishing resolves the item UNRESOLVED rather than in our favour.

1. **AI-referred dollar share, end of 2026.** Requires AI-referred session share **and** relative average order value, or directly reported revenue share, for a period ending Q4 2026. Conversion differential alone will not settle it. Check 2027-03-31. Claim: 3 to 8 percent at $v = 1$, restated if v is published.
 2. **Agent-executed share, 2030.** Census e-commerce totals for 2030 against any published measurement of agent-executed transaction value **with a stated definition**. Check 2031-06-30. Claim: 5 to 15 percent. If no such figure with a stated definition exists, the item resolves UNRESOLVED — the outcome this document most expects.
 3. **Forecast scope.** Claim: McKinsey does not publish a scope definition for its orchestrated-revenue figure permitting comparison with Bain's, before 2028-12-31.
 4. **Forecast revision.** Claim: at least one of the two headline figures is revised by more than 40 percent or withdrawn before 2029-12-31.
 5. **Professional-services citation rate.** Claim: rises above 10 percent by 2027-12-31 (Similarweb category series). This resolves a citation-rate prediction and nothing else. It does not establish presence or absence of a market: citation is not the whole channel, and a small number of consequential encounters can carry value.
 6. **Machine-readability spread.** Claim: Adobe's best-to-worst spread stays above 14 points through 2027-12-31.
 7. **Accuracy reporting.** Claim: no vendor or academic benchmark publishes entity-resolution accuracy by field, as distinct from presence, mention or citation rate, before 2027-12-31.
 8. **Comparability.** Claim: at a re-audit on 2027-09-10, no three of the nine platforms publish mutually reconcilable denominators.
-

10. What would make this wrong in ways worth knowing

If AI-referred dollar share at end of 2026 is below 1 percent, the retail channel is further out than this assumes. It would not establish a ten-year delay, and it would say nothing about mediation in fields these instruments do not observe.

If McKinsey's trajectory proves closer than Bain's, the deployment-constraint reading is wrong and mediation arrives faster than section 4 allows.

If a vendor publishes accuracy-by-field with a stated rubric and an open sample, section 8's burden claim weakens immediately and should be withdrawn to the extent that publication covers.

If AI-referred average order value is materially above or below parity, section 4's estimate moves and the equal-value assumption should be replaced rather than defended.

Sources

US Census Bureau, *Quarterly Retail E-Commerce Sales, 2nd Quarter 2026*, released 2026-08-18. https://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf

Adobe Digital Insights, *Quarterly AI Traffic Report*, Q2 2026. <https://business.adobe.com/resources/sdk/.2026-q2-ai-traffic-report/q2-2026-adi-ai-sourced-traffic-insights.pdf>

Adobe, *AI traffic grows but retail sites lag in AI search visibility*. <https://business.adobe.com/blog/ai-traffic-surge-retail-sites-not-machine-readable>

Conductor, *2026 AEO/GEO Benchmarks*, as reported. <https://www.tryanalyze.ai/blog/ai-traffic-research>

Conductor vertical analysis and Attrifast cohort study, as reported. <https://authoritytech.io/curated/ai-citation-rate-benchmarks-industry-brand-size-2026>

AI Visibility Scores Are Not Comparable — nine-platform methodology audit, 2026-09-10. <https://authoritytech.io/blog/ai-visibility-scores-not-comparable-platform-methodology-audit>

Similarweb, *Generative AI statistics*, May 2026. <https://aisearch.similarweb.com/blog/gen-ai-stats/>

Aggarwal et al., *GEO: Generative Engine Optimization*, KDD 2024. <https://arxiv.org/abs/2311.09735> — dataset <https://huggingface.co/datasets/GEO-Optim/geo-bench>

McKinsey and Company, *The agentic commerce opportunity*. <https://www.mckinsey.com/capabilities/quantuminsights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants>

Bain and Company, *2030 Forecast: How Agentic AI Will Reshape US Retail*. <https://www.bain.com/insights/2030-forecast-how-agentic-ai-will-reshape-us-retail-snap-chart/>

Files

<https://raw.githubusercontent.com/leesharks000/alexanarch/main/data/staged/EA-MEDIATED-SHARE-01-v0.3.md>

The complete text is the repository-hosted Markdown artifact declared under Files above. Under Transport D attachment ingestion, that text is ingested into the canonical deposit body rather than treated as an external pointer.

Suggested Citation

Sharks, Lee. “AI Mediation of Economic Decisions: A Registered Forecast, a Conditional Valuation, and a Statement of Burden (EA-MEDIATED-SHARE-01 v0.3)” *Alexanarch*, 2026-09-20. <https://www.alexanarch.org/s/records/1632/>

Deposit Information

This paper is deposit #1632 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:06C6.UNCLASSIFIED.Δ is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1632/>
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: [alexanarch.org](https://www.alexanarch.org/)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/1632/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.