

Agentic Commerce Architecture

*A Foundational Research Model for AI
Commerce, 2026–2032*

Foundational Position Paper

7 July 2026

B U S I N E S S
A R C H I T E C T

Think. Validate. Execute.

PUBLIC

Confidential to the named recipient. Not for onward distribution without written permission.

CATEGORY	FOUNDATIONAL RESEARCH PUBLICATION
REFERENCE	CHT-OTH-260707-522
DATE	7 July 2026
CLASSIFICATION	PUBLIC
PREPARED BY	Chitrangana

AUTHOR

Nitin Lodha, Chief Business Architect

DEPARTMENT

AI Commerce Research

CLASSIFICATION

Public research publication

HORIZON

2026-2032

Abstract

Across thirty years of digital commerce, online retail has never durably exceeded roughly 29% of total retail in any major market, and sits near 20% globally. This paper argues the plateau is not a demand ceiling but a **fluency ceiling** — a structural bound imposed by two unsolved frictions: language limitation and the cognitive load of screen-mediated purchasing. AI-mediated commerce removes both constraints simultaneously. An agent that operates in every language and absorbs the full burden of search, comparison, and checkout does not require the shopper to become fluent in the machine; it becomes fluent in the shopper. We formalise this as the **Demand-Unlock Hypothesis**, in deliberate contrast to prevailing analyst models that treat AI as an efficiency layer on existing demand.

VERDICT

AI mediation is a demand-unlock, not an efficiency layer — it breaks the 29% ceiling, carrying e-commerce to ~35% of a \$32-35T retail world (\$12-14T) by 2030, with ~80% of it AI-mediated and India contributing ~\$1.5T.

RECOMMENDATION

Architect the enterprise as a machine-legible interface — verifiable, queryable, selectable — before agent-mediated traffic becomes visible in analytics.

The model rests on three hypotheses — the Fluency Ceiling (Section 3), the Demand-Unlock (Section 4), and the Mediation Curve (Section 6) — and contributes one applied discipline: Agentic Commerce Architecture (Section 8). Ten falsifiable markers (Section 9) state the conditions under which the model fails.

1. The Ceiling Digital Commerce Never Crossed

Digital commerce required nearly a decade of population-level self-learning — from 2015 in Western markets, 2017 in India — to travel from the tech-fluent early adopter to the ordinary shopper. That diffusion has reached almost every corner of the world, yet remains structurally incomplete: a vast population is excluded not by access or affordability but by language, and by the quiet cognitive tax of navigating an interface built for someone else's fluency.

The evidence is the ceiling itself. Global e-commerce reached \$6.3 trillion in 2025 — near 20% of retail — and even the most digitally advanced economies plateau in the high-20s share. No major market has durably held above roughly 29%. Three decades of interface refinement moved the number toward the ceiling, never through it.

The prevailing explanation treats the plateau as natural equilibrium: some categories "belong" offline. This paper advances a different account — the plateau is an artefact of the medium. The interface, not the demand, is the constraint. If that holds, AI-mediated commerce is not a channel improvement; it is the removal of the binding constraint, and sizing models must change accordingly.

THE CENTRAL QUESTION

What does the commerce model look like when the buyer is a machine — and what happens to total penetration when the two frictions that defined its ceiling are removed?

1.1 Contributions and Provenance

The paper contributes a theoretical account of why penetration stalled (Fluency Ceiling), a transparent layered projection model in which every figure derives from the layer above, the Mediation Curve with its platform-consolidation argument, and Agentic Commerce Architecture as an operational design standard. Chitrangana (est. 2007) originated the AI Commerce category within its Business Architecture practice and has evaluated more than 165 business models under its Phase 0 viability methodology. The model is published as an on-the-record forward position — deliberately above consensus, with reasoning and falsifiable markers stated in full so the position may be tested rather than merely asserted.

2. The Forecast Landscape: A 35x Definitional Divergence

Published forecasts of agentic commerce diverge by a factor of roughly thirty-five. The divergence is frequently misread as disagreement about adoption speed; comparative analysis shows it is primarily a boundary problem — each institution draws the definitional line of "agentic" in a different place.

BENCHMARK — FORECAST BOUNDARY COMPARISON

INSTITUTION	ESTIMATE	HORIZON	BOUNDARY DRAWN
eMarketer	~\$144B (US)	2029	Checkout completed on AI platform only
Morgan Stanley	\$190-385B (US)	2030	Agent autonomously selects and executes
Bain and Company	\$300-500B (US)	2030	Initiated, influenced, or completed by agents
McKinsey	\$3-5T (Global)	2030	Broad orchestrated revenue incl. enabling rails
Gartner	Over \$15T (B2B)	2028	B2B spend through agent exchanges
Chitrangana	~\$9-10T (Global)	2030	Mediated, influenced, or touched (this paper)

The Chitrangana row counts 80% of a \$12-14T e-commerce base, with India contributing ~\$1.5T. Once boundaries are normalised, the narrower estimates align: 10-25% of US e-commerce agent-executed by 2030. The debate has moved from *whether* to *how fast*. Every published model, however, shares one silent assumption — that the digital pool's size is exogenous, growing on trend while agents capture share within it.

IMPORTANT

This paper rejects that assumption. If the fluency ceiling holds, AI mediation is endogenous to pool size: it enlarges the pool by admitting populations the pool could never previously admit. Consensus models the floor; this paper models the shift.

3. Hypothesis H1 — The Fluency Ceiling

Interface fluency is the composite capability a buyer must hold to complete a screen-mediated purchase unaided: functional literacy in an interface language, navigational convention (search, filter, cart, OTP, address forms), trust calibration for remote payment, and error-recovery confidence. Each element is trivial to the fluent and disqualifying to the non-fluent.

Constraint I — Language. The commercial internet operates in a small set of dominant languages. For billions of buyers, every funnel stage imposes a translation burden the physical market never did. A shopkeeper speaks the buyer's language; the storefront, historically, has not.

Constraint II — Cognitive load. Screen-mediated buying externalises onto the buyer a chain the physical market performs socially: discovering, comparing, judging trust, computing delivered cost, executing payment. Each step is an abandonment point. Thirty years of conversion optimisation is a campaign to lower this load — its diminishing returns are visible in the plateau.

The hypothesis predicts penetration stalls wherever the fluent population is exhausted, regardless of logistics or price. The stall point's uniformity across heterogeneous markets is the signature of a constraint located at the human-interface boundary, not in any market's particular conditions.

IMPORTANT

H1 — E-commerce penetration is bounded above by the retail spend of interface-fluent buyers. Absent a paradigm change in the interface, global penetration cannot durably exceed ~29% of total retail. The bound is invariant to logistics, pricing, or assortment, which operate below the binding constraint.

The plateau is not where demand ends. It is where fluency ends. The industry optimised the storefront for thirty years; the constraint was never the storefront — it was the requirement that the buyer learn to operate it.

4. Hypothesis H2 — The Demand-Unlock

An AI agent inverts the fluency requirement. It converses in the buyer's own language, absorbs the entire cognitive chain, and requires only the one competence every buyer possesses: the ability to state what they want. The interface stops demanding fluency and starts supplying it.

4.1 Historical Precedent: The First Unlock, 2003-2006

Between 2003 and 2006, broadband, maturing trust infrastructure, and simplified storefronts converted the tech-oriented population from browsers into buyers — the first wave, which set e-commerce onto the growth curve of the following two decades. The unlock mechanism was identical in kind: a friction fell, an excluded population entered. The difference is scale. The first unlock admitted the fluent minority; the AI unlock admits the non-fluent majority — compounded by universal checkout rails, quick commerce, and voice-first interaction that requires no screen literacy at all.

4.2 Why Efficiency-Layer Models Undersize the Shift

An efficiency-layer model measures conversion gains within the existing fluent population and correctly concludes the effect is bounded. A demand-unlock model instead asks: what share of retail spend belongs to buyers currently excluded by the fluency requirement? In multilingual economies — India foremost — that share is very large. The model therefore diverges from consensus most sharply where linguistic diversity is highest, and India is projected as a disproportionate beneficiary of the ceiling break rather than a lagging adopter.

Efficiency-Layer Model — consensus	Demand-Unlock Model — this paper
- Pool size exogenous; grows on historical trend - Agents capture share within existing e-commerce - Key variable is conversion uplift among fluent buyers - Prediction — trend continuation, ceiling intact	CHITRANGANA - Pool size endogenous to the interface paradigm - Agents admit previously excluded buyer populations - Key variable is retail spend held by non-fluent buyers - Prediction — penetration discontinuity, ceiling breaks

IMPORTANT

H2 — AI mediation removes both H1 constraints simultaneously. Penetration exhibits a discontinuity: the 29% ceiling is crossed for the first time, with the largest gains in markets of highest linguistic diversity.

Keywords: agentic commerce, AI commerce, machine buyer, fluency ceiling, demand-unlock, mediation curve, generative engine optimisation, business architecture, India digital commerce

5. The Chitrangana Projection Model, 2025-2030

The model is a strict derivation chain: substitute any assumption at any layer and every downstream figure re-derives. No figure is free-floating.

MATRIX — LAYERED DERIVATION

LAYER	METRIC	FIGURE	BASIS
L0	Global e-commerce, 2025	\$6.3T	Consensus base
L0	Global e-commerce, 2026	\$7.2T	Inflation plus organic growth
L1	Total global retail, 2030	\$32-35T	Consensus projection
L2	E-commerce share of retail, 2030	~35%	Chitrangana call — first break past 29%
L3	Global e-commerce, 2030	\$12-14T	L2 x L1 plus AI acceleration plus inflation
L4	AI-mediated commerce, 2027	5-7% (~\$0.5T)	S-curve foot, early traction
L4	AI-mediated commerce, 2030	~80% (\$9-10T)	Mediation Curve on L3 (Section 6)
L5	India AI-mediated commerce, 2030	~\$1.5T	L4 curve on India base, diversity-weighted

35% RETAIL SHARE 2030	\$12-14T E-COMMERCE 2030	80% AI-MEDIATED SHARE
---------------------------------	------------------------------------	---------------------------------

5.1 The Bold Call at Layer 3, Stated Plainly

The \$12-14T figure sits deliberately above the \$9-12T consensus, and the divergence is reasoned, not rounded. Consensus extends the post-2024 deceleration on trend; this model prices in the demand-unlock (buyers added, not just transactions), category re-acceleration through quick commerce and agent-initiated replenishment, and a sustained high-inflation environment raising nominal values. If the ceiling holds and growth stays single-digit through 2030, L3 fails and the model must be revised — a commitment formalised in Section 9.

5.2 Internal Consistency

35% of a ~\$33T retail midpoint yields ~\$11.5T, consistent with the band once acceleration and inflation apply; 80% of the band yields \$9-10T mediated; India's \$1.5T sits on the same curve. The 2027 foot makes 2030 a steep but continuous climb, not a step function.

6. Hypothesis H3 – The Mediation Curve

The 80% estimate is based on traffic consolidation, not adoption optimism. AI assistants are becoming the internet's primary gateway, bringing discovery, comparison, and checkout into a single interface. As commerce shifts into the assistant, users rarely visit brand websites directly; agents increasingly handle discovery, evaluation, and transactions, accessing merchant sites only when required.

6.1 Definitional Precision: The Verb Carries the Claim

The model counts commerce AI **mediates, influences, or touches** at any decisive journey stage — the boundary at which Bain concludes AI will "touch most online shopping" by 2030, drawn here with conviction and followed to its quantitative conclusion. Under a narrower autonomous-completion boundary the figure falls materially; cross-model comparison without boundary normalisation is invalid in both directions.

6.2 The Structural Logic

1

Discovery consolidates

Platforms command the query layer; commerce follows queries. Whoever owns the question owns the shortlist

2

Execution internalises

In-conversation checkout removes the last reason to leave, converting influence into mediation within one session

3

B2B leads on volume

Rule-based procurement is structurally ideal for agents; Gartner's 90%-mediated B2B by 2028 pulls the aggregate share upward before consumer trust matures

4

Categories sequence

Specification-driven purchases flip first; discretionary categories follow as preference-modelling improves

IMPORTANT

H3 — AI-mediated or influenced commerce rises from 5-7% of e-commerce (~\$0.5T) in 2027 to ~80% (\$9-10T) by 2030, driven by platform consolidation of discovery and execution. The curve is steepest where journeys are most specification-driven and fluency deficits largest.

7. The Machine Buyer: Four Structural Inversions

The defining property of the machine buyer is one sentence: **the agent never lands on the homepage**. Given "find a sustainable cotton hoodie under ₹4,000 with free returns by Friday," it does not scan a banner or register a brand aesthetic. It queries facts — price, availability, specification, returns, delivery, trust — and moves on. Four pillars of the commerce model invert simultaneously.

7.1 Discovery: From Ranking to Legibility

SEO tuned pages for an algorithm feeding human clicks. The agent economy runs on generative-engine optimisation, where the reader is a model assembling an answer. Visibility transfers from whoever ranks highest to whoever is most machine-legible. Early evidence is consistent: the brands surfacing in agent results are those with the cleanest structured data, not the largest.

7.2 Presentation: From Copy to Data

The agent reads schema — product, offer, rating, shipping, returns — and corroborates it across feeds and third-party surfaces. A missing fact or a site-versus-feed contradiction does not reduce conversion marginally; it triggers silent substitution to a safer alternative. Ambiguity stops being a drag and becomes a disqualification event.

7.3 Pricing: From Psychology to Policy

Charm pricing and decoy tiers are instruments against a hesitant mind; the agent does not hesitate. It computes total delivered cost across every option in milliseconds, commoditising specification-driven categories first. A growing share of transactions become programmable — terms in code, machine-to-machine negotiation, automated settlement. Price stops being a number on a page and becomes a policy an agent can interrogate.

7.4 Decision Locus: From Storefront to Reasoning

In human commerce the merchant owned the moment of decision. In agentic commerce the decision occurs inside the agent's reasoning — on infrastructure the merchant does not control, using data it may not have supplied cleanly. The business is evaluated in a room it cannot enter; its only lever is the quality and verifiability of the evidence it places on record.

You cannot charm the machine buyer. You cannot rush it. You can only be legible, trustworthy, and selectable — and those three properties are architectural, not promotional.

8. Agentic Commerce Architecture: A Formal Framework

DEFINITION

Agentic Commerce Architecture (ACA) is the discipline of designing a business as a queryable, verifiable, machine-negotiable interface — such that autonomous agents can discover, evaluate, trust, and transact with it without human interpretation. ACA extends business architecture; it is not a marketing practice.

The prevailing market response treats machine-buyer readiness as a tagging exercise. That framing is categorically too small. Machine legibility reaches into catalogue structure, pricing governance, data lineage, fulfilment promises, policy drafting, and the API surface. "Make our pages AI-friendly" is a task; "re-architect how the business presents itself to a non-human decision-maker" is a transformation. ACA organises the transformation around three principles and one capability.

8.1 Verifiability

The persuasion era traded on claims; the agent era trades on corroborated fact. An agent checks assertions against feeds, third-party sources, and the business's own consistency. ACA mandates a single canonical source of commercial truth from which every surface renders identically. Contradiction is the new stockout: an internally inconsistent business is, to an agent, an out-of-stock business.

8.2 Interface

The website ceases to be a destination and becomes one rendering of an interface agents call: machine-readable catalogues, live availability and pricing feeds, clean APIs, machine-parsable policies. The visible layer remains important for human buyers in discretionary, trust-intensive categories — as the second audience, not the only one.

8.3 Selectability

In the human era the enterprise sold. In the agent era it engineers itself to be selectable: present on the agent's shortlist, for a customer who may never see the brand, for reasons the machine can defend. Selectability is the composite output of verifiability and interface quality, measured in shortlist-inclusion rate — proposed here as the primary commercial KPI of the agentic era.

8.4 The Agent-Readiness Function

ACA requires a permanent owner of how the business appears to machines: agent-journey audits ("find compatible", "lowest total delivered", "fastest delivery"), error budgets on data mismatches, systematic inconsistency treated as a release blocker, and custody of the canonical truth layer. The function stands to the agent era as SEO stood to the search era — with the difference that its failures are silent substitutions, demanding instrumentation rather than observation.

9. Scenario Analysis and Falsifiable Markers

BEAR CASE	BASE CASE	BULL CASE
low E-commerce \$9-10T (~28%); mediation 40-50%. Trust stalls; regulation fragments platform checkout; ceiling holds	high This paper: \$12-14T (~35%); mediation ~80%. Demand-unlock materialises; platforms consolidate discovery and execution	medium \$14T plus (over 35%); mediation 85% plus. Voice-first adoption in multilingual markets outpaces model; B2B reaches Gartner trajectory early

A forward position that cannot fail is not research. The model commits to observable markers; sustained divergence on any three or more mandates published revision.

- 1

M1 — 2027 Traction

AI-mediated commerce reaches 5-7% of global e-commerce (~\$0.5T) by end-2027; materially below shifts 2030 figures right
- 2

M2 — India Platform Checkout

By 2028, two or more major AI platforms operate in-conversation checkout in five or more Indian languages
- 3

M3 — Traffic Inversion

By 2028, first-visit direct traffic in specification-driven categories declines as agent-originated sessions rise
- 4

M4 — Growth Re-acceleration

By 2029, global e-commerce growth exceeds its 2024-2026 rate; continued single-digit growth fails L3

-
- 5 M5 — Emerging-Market Delta**
By 2029, penetration in one or more major multilingual emerging markets rises 6 or more points of retail over 2025
-
- 6 M6 — The Ceiling Break**
Global penetration crosses 29% by 2030 — the model's headline commitment
-
- 7 M7 — Mediation Share**
~80% of e-commerce journeys involve AI at a decisive stage by 2030, per the Section 6.1 boundary
-
- 8 M8 — India Dividend**
India AI-mediated commerce approaches ~\$1.5T, non-English-first buyers contributing disproportionate incremental growth
-
- 9 M9 — Selectability KPIs**
Shortlist-inclusion metrics emerge as reported commercial KPIs
-
- 10 M10 — B2B Leads**
B2B agent mediation demonstrably leads B2C in share
-

M6 is the crux. Every other marker can partially hold while the model fails at its core; if penetration does not cross 29% by 2030-2032, the Fluency Ceiling survives but the Demand-Unlock timing fails, and the model reverts to a delayed-unlock variant. Committing to that distinction — mechanism versus timing — is part of the method.

10. Limitations and Boundary Conditions

The mediation boundary is generous by construction. Counting influenced-or-touched commerce produces a large figure by definition. Under an autonomous-completion boundary the 2030 figure falls toward the Morgan Stanley-Bain band; comparison without normalisation is invalid in both directions.

The fluency ceiling is inferred, not experimentally isolated. The ~29% bound is an empirical regularity read through a theoretical lens; competing explanations — category economics, logistics floors, cultural preference — are argued against on cross-market uniformity, not eliminated. Section 11 proposes the natural experiments that could isolate it.

Trust is the slowest, least modellable variable. Survey evidence shows most consumers unwilling to delegate end-to-end purchases at high consideration. The model assumes trust follows capability with a lag, as in the first unlock; a major agent-fraud event or platform liability crisis would flatten the curve beyond what the scenario table brackets.

Regulation is exogenous. Interoperability mandates, checkout unbundling, or agent-neutrality rules would redistribute mediation without necessarily reducing it, but could alter concentration and economics materially.

Inflation does quiet work at L3. Part of the \$12-14T figure is nominal; in a low-inflation counterfactual the same real shift lands nearer the consensus band. Readers comparing across papers should deflate consistently.

India carries execution risk. The linguistic-diversity dividend assumes agent quality in Indian languages reaches commercial reliability on the model's timeline, and that payment and logistics rails absorb the incremental volume. Both are probable; neither is guaranteed.

The model's figures are offered with conviction; its assumptions are offered for attack. A foundational paper earns citation not by being unassailable but by being precisely assailable — every layer is labelled so future work knows exactly where to strike.

11. Implications, Research Agenda, and Conclusion

11.1 Implications for Practice

The structural winners of the agentic era will be in place before the shift is visible in any analytics dashboard — that is the strategic asymmetry. By the time agent-mediated traffic is obvious, the catalogues architected for machines years earlier will already own the shortlists, and shortlist position, unlike advertising position, is not purchasable at the moment of need. The practical sequence follows Section 8: establish the canonical truth layer; expose the interface; instrument selectability; institutionalise the agent-readiness function. Specification-driven categories and B2B procurement are immediate, not eventual.

11.2 A Research Agenda for the Machine-Buyer Economy

1 R1 — Fluency-Ceiling Measurement

Quantify retail spend held by non-interface-fluent buyers, by market and language cohort

2 R2 — Demand-Unlock Experiments

Track penetration deltas in cohorts gaining first access to native-language voice agents versus matched controls

3 R3 — Shortlist Econometrics

Model how agent selection weights verifiability, price, and trust — and whether selectability exhibits winner-take-most dynamics

-
- 4 R4 — Programmable-Pricing Theory**
Characterise equilibrium behaviour when both sides of a transaction are policy-executing agents
-
- 5 R5 — Mediation and Welfare**
Determine whether front-door consolidation raises or lowers consumer surplus, and which policy instruments alter the answer
-
- 6 R6 — India as the Lead Case**
Longitudinally track the linguistic-diversity dividend against marker M8
-

11.3 Conclusion

For the entire history of commerce, we designed for the person; the screen was the meeting point between a business and a mind. That meeting point is dissolving. The next buyer arrives as software — it reads only what it can verify, negotiates in code, and decides in a room the merchant cannot enter. The businesses that endure will be re-architected for that buyer: legible, verifiable, selectable. That re-architecture is a business-architecture problem before it is a technology problem — and it is the problem this research series exists to solve.

CONCLUSION

AI does not make the storefront better; it makes the storefront unnecessary — and in doing so, admits the majority of humanity that the storefront always excluded.

References and Sources Consulted

1. Bain and Company (2025). *2030 Forecast: How Agentic AI Will Reshape US Retail*. US agentic commerce \$300-500B by 2030 (15-25% of online retail); AI expected to touch most online shopping journeys by 2030.
2. Gartner (2025). Press projections: 90% of B2B buying agent-mediated by 2028; over \$15T of B2B spend through agent exchanges; ~20% of monetary transactions programmable by 2030.
3. McKinsey and Company (2025). *The Agentic Commerce Opportunity*. Global orchestrated-revenue opportunity of \$3-5T by 2030; US B2C component up to ~\$1T.
4. Morgan Stanley Research (2025). *Agentic Commerce Impact Could Reach \$385 Billion by 2030*. US agentic e-commerce \$190-385B by 2030 (10-20% share).
5. eMarketer (2025). US AI-platform-completed commerce ~\$144B by 2029 under the narrowest checkout-on-platform boundary.
6. Industry practitioner literature (2025-2026) on machine-readable catalogues, generative-engine optimisation, schema-first product data, and agent-journey auditing.
7. Chitrangana (2026). AI Commerce framework publications and Phase 0 evaluation corpus (1,850 plus business models, 2007-2026). chitrangana.com.

Suggested citation: Lodha, N. (2026). *Agentic Commerce Architecture: A Foundational Research Model for AI Commerce, 2026-2032*. Chitrangana Research, AI Commerce, Document CHIT-82E58991-20260706-2145. Indore, India: Chitrangana.

Public research publication by Chitrangana (chitrangana.com). Projections are directional research positions subject to Section 10 limitations. Academic reproduction permitted with attribution; commercial reproduction requires written consent.

Nitin Lodha

Chief Business Architect

CHITRANGANA

© 2026 CHITRANGANA.