SEED·VENTURES

RESEARCH BRIEFING · MARKETPLACE STRATEGY · 2026

The Omnichannel Operating *Model*

Channel count is not a strategy. How consumer brands convert marketplace proliferation into durable, margin-accretive growth — and what it costs to get the operating model wrong.

Nine exhibits. A four-layer operating model. A five-stage maturity diagnostic. Built from operating data across a live multi-marketplace portfolio and current research from McKinsey & Company, EMARKETER, ICSC and Marketplace Pulse.

— Contents

About this briefing

This paper is written for brand founders, chief commercial officers and marketplace leaders operating consumer products across Amazon, Walmart, TikTok Shop and adjacent channels. It is deliberately operational: the argument is that omnichannel underperformance is almost never a strategy failure and almost always an operating-model failure.

Third-party research is cited throughout and attributed in full on page 21. Illustrative unit economics are composite benchmarks drawn from Seed Ventures' managed portfolio and are labelled as such wherever they appear; they are directional, not category-specific guidance. Nothing in this document constitutes financial, legal or tax advice.

Seed Ventures LLC is a marketplace growth partner for consumer brands. We operate a wholesale-to-retail model: we buy inventory, take title, and earn retail margin — which means our return depends entirely on sell-through, not on management fees.

— Executive summary

> Most brands do not have an omnichannel problem. They have a coordination problem that only becomes visible once they are on more than three channels — at which point it is usually diagnosed as a demand problem and solved with advertising spend.

Between 2024 and 2026 the practical definition of “being everywhere” changed twice. First, marketplace demand concentrated further into a small number of platforms while simultaneously spawning a credible second tier — TikTok Shop, Temu and Walmart now compete in the same volume band, arriving from three entirely different directions.[1] Second, and more consequentially, the discovery layer itself began migrating into AI assistants that return a shortlist rather than a results page.[2]

Both shifts reward the same underlying capability: clean, consistent, machine-readable product data governed centrally and deployed channel-specifically. Both punish the same failure: catalogues that are complete but not managed. The gap between those two states is where margin goes to die.

39.7%	68%	$1T	1.7×
of US ecommerce sales run through Amazon in 2026 — no other retailer is close[3]	of US consumers used at least one AI tool as part of shopping in the prior three months[4]	of US B2C retail revenue may be orchestrated by agentic commerce by 2030[4]	shopping frequency of omnichannel customers versus single-channel customers[5]

Five findings

01 Channel count correlates with revenue. It does not correlate with contribution margin.
On a representative marketplace P&L, roughly 76 cents of every retail dollar is consumed before contribution — referral fees, fulfilment, advertising, returns leakage, promotions and cost of goods. Adding a channel without redesigning the operating model adds the cost line first and the revenue line second.

02 Content scales linearly. Coordination scales quadratically.
Seventy SKUs across six channels is 420 listings and roughly 9,200 managed data fields — but fifteen cross-channel reconciliation relationships covering price, inventory, promotion and claim consistency. This is the capability cliff most brands hit between channel three and channel four.

03 “Having content” is not optimisation.
Images, copy and A+ modules existing on a listing is baseline compliance. Optimisation is a repeating diagnostic loop — query performance, conversion by term, margin by SKU by channel — run against a standard. Most underperforming catalogues are compliant and unmanaged.

04 AI visibility is a data-quality discipline, not a marketing channel.
When an assistant returns three products instead of forty results, inclusion becomes binary. Structured attributes, entity consistency, review corpus depth and factual claim alignment now determine whether a brand is in the shortlist — and none of those are advertising levers.

05 Net revenue retention, not new-logo growth, is the honest measure of a channel

01 The market has consolidated and fragmented at the same time

> Demand is more concentrated than ever. Opportunity is more distributed than ever. Both statements are true, and holding them together is the entire strategic problem.

The concentration case is straightforward. **Amazon holds an estimated 39.7% of all US ecommerce sales in 2026**, with no other retailer close, and the **top 18 ecommerce retailers together account for 74.1% of total US retail ecommerce**.[3] Concentration is even sharper on the monetisation side: Amazon is estimated to absorb more than **79% of all US retail media digital ad spending** in 2026, because advertising follows transaction volume.[3]

The fragmentation case is equally real. Below Amazon, a genuine second tier has formed. On 2026 estimates, Temu, TikTok Shop and Walmart's marketplace now compete inside a tight band, with TikTok Shop having grown roughly 68% year over year to reach it.[1] Walmart's wider online business — marketplace plus first-party — reached approximately $150 billion, up more than 20% year over year.[6] Below that, specialist platforms in home, resale, handmade and live auction have each built defensible billion-dollar positions serving demand that mass-market platforms structurally under-serve.[1]

EXHIBIT 1

One anchor, then a tightly bunched second tier

Estimated 2026 US marketplace gross merchandise volume, $ billions

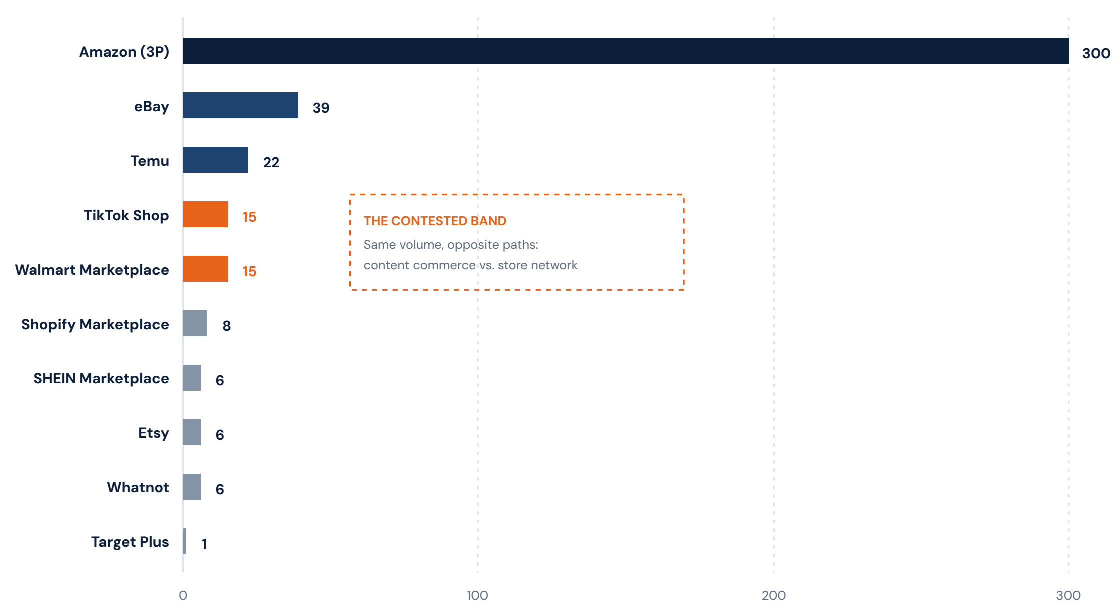

Source: Marketplace Pulse, "Top 10 E-Commerce Marketplaces in 2026" (2026 estimates). Figures are gross merchandise volume estimates and are not directly comparable across hybrid retailer-marketplace models. Analysis by Seed Ventures LLC.

Demand concentration and monetisation concentration are different shapes

The practical consequence of Exhibit 2 is that a brand's revenue distribution and its cost distribution do not move together. A brand may earn 60% of revenue on Amazon while spending 85% of its channel budget there, because the advertising market on the anchor platform is priced against everyone's demand, not just its own. Second-tier channels are cheaper to acquire on and more expensive to operate on. That trade is the actual channel-selection decision, and it is rarely modelled.

EXHIBIT 2

Where demand sits — and where the toll is collected

US retail ecommerce sales and US retail media digital ad spending, share of total, 2026 estimates

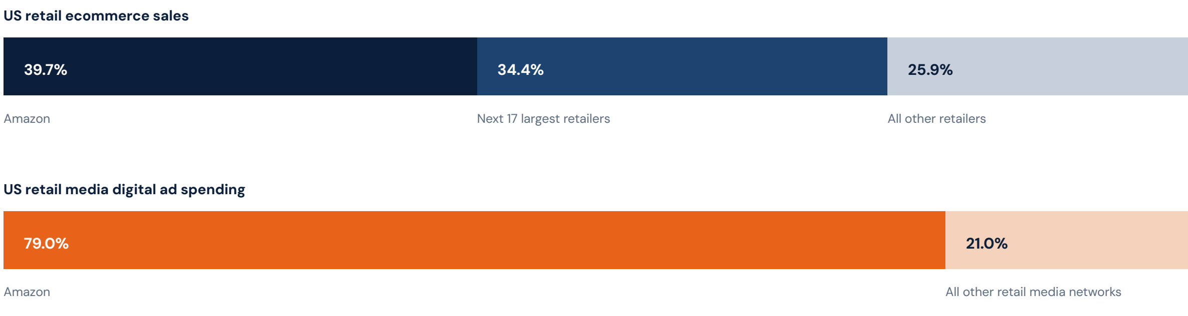

Source: EMARKETER forecasts as reported in "FAQ on ecommerce marketplaces" (2026). Shares rounded; "next 17" derived from the top-18 aggregate. Analysis by Seed Ventures LLC.

> *A channel portfolio is an allocation decision, not a coverage decision. The question is never "should we be on Walmart?" — it is "what is the contribution margin per unit of operating attention on Walmart versus the alternative use of that attention?"*
>
> SEED VENTURES LLC

What the concentration data does not say

It does not say the second tier is optional. TikTok Shop reached its position by embedding commerce inside entertainment, and its buyer base has broadened well past the launch demographic — reporting on first-quarter 2026 data indicates the fastest-growing spender cohort is now 45 and over.[7] Nor is Amazon's internal mix static: third-party sellers accounted for roughly 60% of Amazon paid units in Q1 2026, down from 62% a year earlier, as first-party regains share.[8]

The operating implication

Treat the anchor as the reach engine and the second tier as the margin and discovery engine — then hold both to a single contribution-margin standard. A channel that cannot clear a defined margin floor within a defined ramp window is a portfolio position to exit, not a capability to keep funding. Very few brands have written that floor down, which is why very few brands ever exit a channel.

02 Discovery has moved upstream

> When an assistant returns three products instead of forty results, inclusion stops being a ranking problem and becomes a binary one. That single change rewrites the economics of product data.

The most significant structural shift in consumer commerce since marketplace consolidation is not a new channel. It is the insertion of a synthesis layer between the shopper and every channel. Research conducted with ICSC across more than 3,000 US consumers in January 2026 found that **68% had used at least one AI tool as part of a shopping experience in the prior three months**, and **62% had used AI to compare brands, models, prices or reviews**.[4] McKinsey estimates agentic commerce could account for up to **$1 trillion of orchestrated US business-to-consumer retail revenue by 2030**, and $3–5 trillion globally.[4, 9]

The mechanism matters more than the forecast. Traditional search returns a page; a brand that ranks eleventh is diminished but present. An assistant returns a shortlist; a brand that is not in it is absent. The determinants of inclusion are not advertising levers. They are attribute completeness, entity consistency across surfaces, review corpus depth and recency, factual alignment between claims made on the listing and claims made everywhere else, and machine-legible structure.

EXHIBIT 3

Adoption is already mainstream; the market sizing is not settled

Left: reported US consumer AI use in shopping. Right: 2030 US market estimates, $ billions

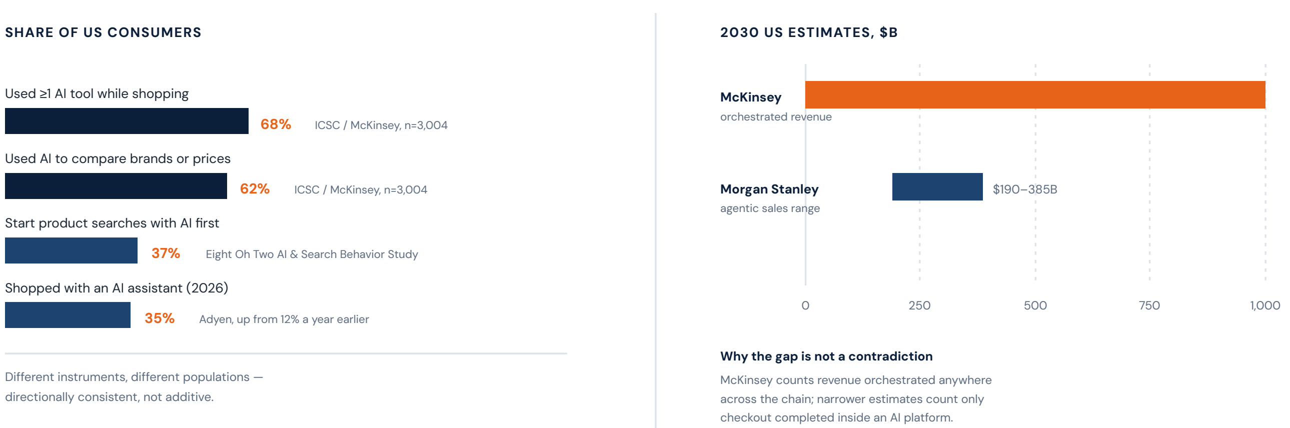

Sources: ICSC and McKinsey & Company, “Shopping in the age of AI” (April 2026), survey of 3,004 US consumers fielded 7–11 January 2026; Eight Oh Two AI & Search Behavior Study; Adyen consumer research; Morgan Stanley estimates as compiled in published agentic-commerce benchmark reviews (2026). Estimates measure different definitions of the market and should not be read as a single figure.

The retailer assistants are the near-term surface, not the frontier models

For brands selling on marketplaces, the immediate exposure is not a general-purpose chatbot. It is the assistant embedded in the retailer. Amazon disclosed that monthly active users of its shopping assistant grew 115% year over year before it was merged into a broader shopping experience in May 2026.[8] These systems read the same catalogue a brand already controls: title, bullets, structured attributes, A+ content, Q&A and reviews. A listing that is thin, inconsistent or unstructured is not merely ranked lower — it is harder to summarise, and therefore less likely to be summarised.

This is why we treat AI visibility as a data-quality discipline rather than a marketing channel. The work is unglamorous: normalising attribute taxonomies across channels, eliminating claim drift between the DTC site and the marketplace listing, building review volume against the specific use cases the assistant is being asked about, and keeping availability and pricing signals accurate enough that the brand is not filtered out before it is considered.

> *The brands most exposed to AI-mediated discovery are not the ones with the smallest ad budgets. They are the ones whose product data was assembled channel by channel, by different people, at different times, against no standard.*
>
> SEED VENTURES LLC

Where the physical channel now fits

The same research is a useful corrective to the assumption that AI adoption simply accelerates the shift online. It found that shopping mostly online was not the leading method for any generation or category studied, and that **37% of respondents ranked in-stock reliability among their top three reasons for choosing a retailer**.[4] Gen Z and millennial respondents were significantly more likely than older cohorts to shop both online and in store.[4] The store is not disappearing; its mission is narrowing to fulfilment, validation, immediacy or experience.

For a consumer brand without owned retail, the translation is direct. If availability signals and validation content are the two things shoppers and agents now look for, then inventory accuracy and content quality stop being operational hygiene and become demand drivers. A stockout is no longer a lost sale; it is a delisting from the consideration set that persists after the inventory returns.

A note on enterprise readiness for AI itself

The adoption gap on the operator side mirrors the one on the consumer side. McKinsey's 2025 State of AI survey found that while 88% of organisations report using AI in at least one business function, only around a third are scaling it across the enterprise — 62% were experimenting with AI agents, and just 23% had begun scaling agentic AI in any function.[10]

The implication for channel teams is that the constraint is rarely tooling. It is that no one owns the standard the tooling is supposed to enforce.

03 Why adding channels destroys margin

Roughly three-quarters of every retail dollar is consumed before contribution. A new channel adds the cost structure on day one and the demand curve on month nine.

Omnichannel customers are demonstrably more valuable. McKinsey research indicates omnichannel customers shop about **1.7 times more often** than single-channel customers, and independent analysis puts the omnichannel order-value premium at around **16%** with materially higher engagement for retailers operating three or more channels.[5, 11] None of that value is automatic. It is realised only if the incremental channel clears its own cost structure.

Exhibit 4 decomposes a representative marketplace P&L on an indexed basis. The specific percentages vary enormously by category, weight, price point and fulfilment method — but the shape is remarkably stable, and the shape is the point.

EXHIBIT 4

The channel margin waterfall

Indexed decomposition of $100 in marketplace retail revenue to net contribution; illustrative composite

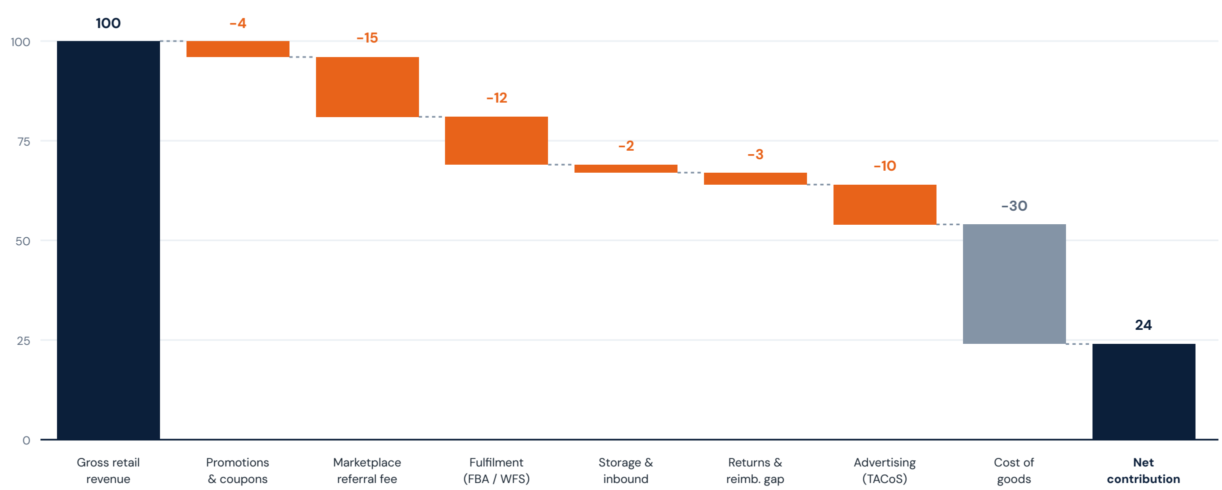

Six of the eight deductions above are operating-model outputs, not market prices. Only the referral fee is genuinely fixed.

Illustrative composite benchmark based on Seed Ventures' managed portfolio across general merchandise and consumer packaged categories, indexed to $100 of gross retail revenue. Directional only; actual economics vary substantially by category, unit weight, price point and fulfilment method.

Only one line in the waterfall is a market price

The referral fee is set by the platform. Every other deduction is an output of how the brand operates. Advertising load is a function of organic conversion, which is a function of content and review quality. Returns leakage is a function of listing accuracy and packaging. Storage cost is a function of forecast quality. Promotional depth is a function of pricing architecture and MAP enforcement. Cost of goods is a function of supplier terms and order cadence, which are functions of demand visibility.

This is why "we need more traffic" is usually the wrong diagnosis. A brand converting at 8% with a 10% advertising load and a brand converting at 14% with the same load are not running the same business, and no amount of incremental spend closes that gap. It closes by fixing the inputs.

The symptom

Revenue is flat or growing slowly, advertising spend is rising to hold position, and gross margin is compressing quarter over quarter. Channel teams request budget; finance asks why last quarter's budget did not compound.

The cause

Each channel has been added as an incremental workstream on top of an operating model designed for one channel. Nothing was retired, no standard was written, and no single owner holds contribution margin per SKU per channel.

Content scales linearly. Coordination scales quadratically.

The second cost of channel expansion is not visible in any P&L line. It is the coordination burden created by every additional pair of channels that must stay consistent with each other on price, availability, promotion, claim and image. Content volume grows in proportion to channels. Reconciliation relationships grow as $n(n-1)/2$. This is the mechanism behind the capability cliff that most brands hit somewhere between the third and fourth channel.

EXHIBIT 5

The coordination penalty

Managed data fields and cross-channel reconciliation relationships, 70-SKU catalogue

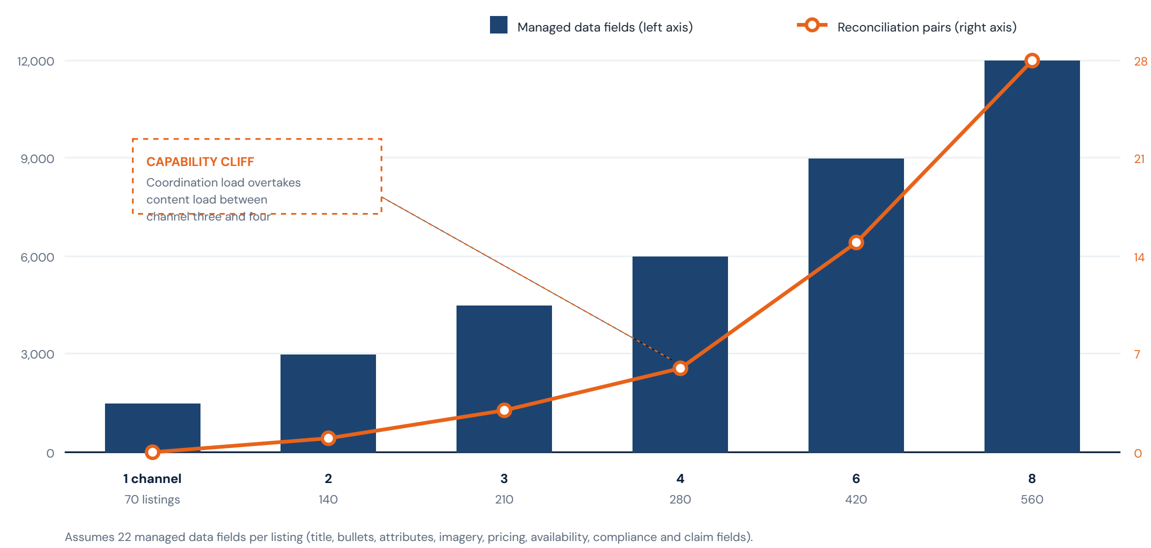

Seed Ventures LLC analysis. Reconciliation relationships calculated as $n(n-1)/2$ across active selling channels. Illustrative model based on a 70-SKU catalogue.

"Having content" is not optimisation

The single most common finding in a channel audit is that the listing is complete and unmanaged. Seven images exist. Bullets exist. A+ content exists. Someone wrote them, once, at launch. Nothing since has tested whether the title reflects the terms that actually convert, whether the hero image survives a mobile thumbnail crop, whether the top three search terms driving impressions are the same three driving purchases, or whether the SKU still clears its margin floor after eighteen months of fee changes and promotional creep.

Optimisation is a loop, not an asset. It has a diagnostic input (query-level performance data), a decision rule (a written standard for what good looks like), an intervention, and a measurement window. Brands that describe their listings as "optimised" are usually describing the asset. The distinction is not semantic — it is the difference between a fixed cost incurred once and a compounding capability.

> *Compliance asks: does the listing have what the channel requires? Optimisation asks: is this listing winning the specific queries that carry margin — and how would we know within fourteen days if it stopped?*
>
> SEED VENTURES LLC

The three failure modes

FAILURE MODE	WHAT IT LOOKS LIKE	WHAT IT ACTUALLY COSTS
Coverage without control	Live on six channels; pricing, promotions and claims are inconsistent across at least three of them.	MAP erosion and buy-box instability on the anchor channel. Wholesale partners lose confidence. Margin leaks fastest where the brand is least watching.
Content without diagnosis	Complete listings, no query-level review cadence, no written content standard, no owner for conversion rate.	Advertising load rises to compensate for organic conversion decay. The cost is invisible because it appears as a media line, not a content line.
Forecasting inside operations	Reorder decisions made by whoever manages the warehouse, using trailing sell-through and instinct.	Simultaneous stockouts and overstock in the same catalogue. Stockouts now cost ranking and assistant-shortlist inclusion, not just the lost unit.

A structural point about team constraints

In our experience, omnichannel underperformance is almost never explained by the capability of the people doing the work. It is explained by role structure and mandate: no single owner for contribution margin, forecasting sitting with execution rather than with data, brand standards distributed across every pod instead of held centrally, and no named owner for net revenue retention.

That distinction matters because the two diagnoses lead to opposite interventions. One leads to replacing people. The other leads to redrawing the operating model — which is almost always the correct move, and almost always the harder conversation.

04 The four-layer operating model

> Four layers, four owners, four measures. The test of the model is not whether the work gets done — it is whether a single named person can be asked why a number moved.

Most channel organisations are structured by platform: an Amazon person, a Walmart person, a TikTok person. That structure guarantees the coordination penalty described in Exhibit 5, because every cross-channel standard becomes a negotiation. The alternative is to organise by discipline and deploy across channels — which concentrates standards centrally and pushes only execution to the channel edge.

EXHIBIT 6

Four layers, four accountable measures

Disciplines held centrally; execution deployed channel-specifically

LAYER 01 **Strategy**	LAYER 02 **Execution**	LAYER 03 **Distribution**	LAYER 04 **AI Visibility**
SCOPE Channel portfolio design. Pricing architecture and MAP policy. Assortment, bundling and pack strategy. Margin floors and exit criteria.	SCOPE Catalogue and content standards. Query-level diagnostics. Advertising structure and budget governance. Compliance, cases and customer experience.	SCOPE Demand forecasting and reorder logic. Inventory placement across FBA, WFS, AWD and 3PL. Cost-to-serve and supplier terms.	SCOPE Structured attribute taxonomy. Entity and claim consistency across surfaces. Review corpus depth. Availability and pricing signal accuracy.
ACCOUNTABLE MEASURE **Contribution margin per SKU per channel**	ACCOUNTABLE MEASURE **Organic conversion rate and advertising load**	ACCOUNTABLE MEASURE **In-stock rate and inventory turns**	ACCOUNTABLE MEASURE **Shortlist inclusion rate and attribute completeness**
COMPANY-LEVEL NORTH STAR			**Net revenue retention**

Seed Ventures LLC operating framework. Layers are disciplines, not departments; in organisations below roughly $25M in channel revenue, one person may own two adjacent layers — but never Strategy and Execution together, and never Execution and Distribution together.

Three separations that do most of the work

Separate forecasting from execution. Reorder logic is a data-science problem: demand signal, seasonality, lead-time variance, promotional lift, channel cannibalisation. Execution — receiving, labelling, shipping, replenishing — is an operations problem. When forecasting sits inside the warehouse, the forecast inherits the warehouse's incentives, which are to avoid stockouts and to avoid inbound complexity. Those two incentives produce overstock in slow SKUs and stockouts in fast ones, simultaneously.

Separate brand standards from channel execution. Catalogue brand standards, imagery rules, claim libraries and MAP violation resolution should be held once, centrally, and consumed by every channel team. When brand standards are distributed to channel owners, drift is not a risk — it is a certainty, because each owner optimises locally against a different platform's conventions.

Separate margin ownership from revenue ownership. The person accountable for growth should not also be the person who signs off on the margin maths that justifies the growth. Accounting owns margin, bookkeeping and P&Ls; growth owns retention and expansion. The friction between those two roles is the control.

> *If you cannot name the single person who owns net revenue retention, the honest answer is that nobody does — and that revenue arriving through new channels is quietly offsetting revenue leaving through old ones.*
>
> SEED VENTURES LLC

What each layer produces, weekly

LAYER	WEEKLY ARTEFACT	REVIEW CADENCE	ESCALATION TRIGGER
Strategy	Contribution margin by SKU by channel, with variance to margin floor	Weekly, reviewed monthly with finance	Any SKU below floor for two consecutive weeks
Execution	Query performance delta: impression share, click share, purchase share by term	Weekly	Purchase share decline >15% on a top-10 term
Distribution	Weeks of cover by SKU by fulfilment node, plus inbound exceptions	Weekly	Cover below 4 weeks or above 16 weeks
AI Visibility	Attribute completeness score and shortlist inclusion sample	Fortnightly	Completeness below 90% on a revenue-decile SKU

The cadence is the capability

None of the artefacts above are difficult to produce. Every one of them is difficult to produce *every week, without being asked*. That is the actual gap between a channel programme that compounds and one that plateaus — and it is why tooling investments so often fail to change outcomes. Software enforces a standard; it does not write one.

05 The omnichannel maturity curve

Five stages. Most brands sit at stage two and describe themselves as stage four. Essentially all of the economic value is created in the transition from three to five.

EXHIBIT 7

Five stages of omnichannel maturity

Capability, ownership and indicative economic outcome by stage

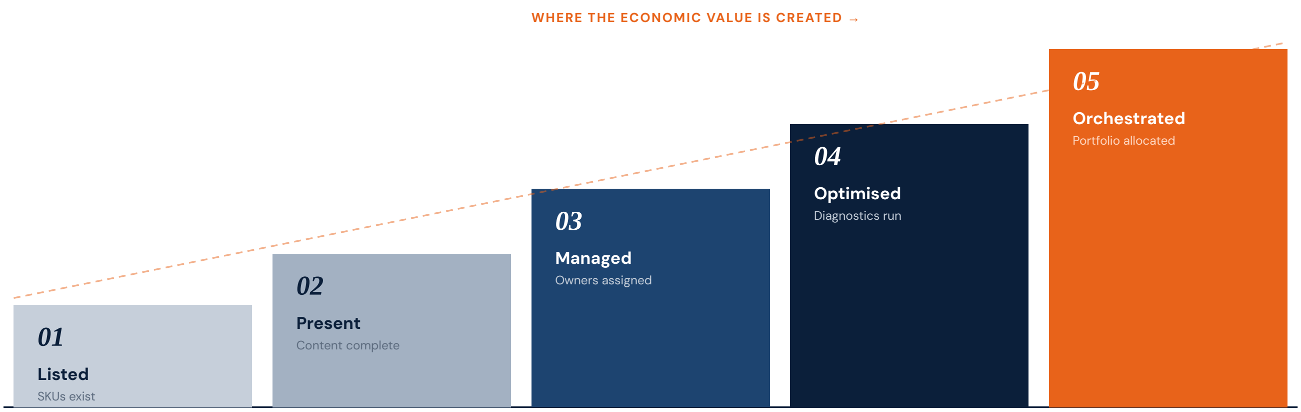

Bar height indicates relative operating capability, not revenue.

STAGE	WHAT IS TRUE	WHAT IS OWNED	TYPICAL ECONOMIC SIGNATURE
01 Listed	Products are live. Content was written once, at launch.	Nobody, explicitly. Channel work is absorbed by whoever has capacity.	Revenue exists; contribution margin is unmeasured.
02 Present	Content is complete and compliant. Brand registry in place.	A channel manager per platform.	Margin visible at channel level, not SKU level. NRR flat to negative.
03 Managed	Pricing, MAP, inventory and advertising have named owners and written rules.	Disciplines, not platforms. Standards held centrally.	Advertising load stabilises. Stockouts decline. Margin becomes forecastable.
04 Optimised	Query-level diagnostics, bundling and content iteration run on a fixed cadence.	Contribution margin per SKU per channel has a single owner.	**Organic conversion rises; advertising load falls. NRR turns positive.**
05 Orchestrated	Inventory and advertising are allocated across channels against margin, not habit. AI visibility governed.	Portfolio-level allocation, reviewed monthly against exit criteria.	**Channel additions become accretive by design rather than by hope.**

Seed Ventures LLC maturity framework, derived from channel audits across managed and advisory engagements. Economic signatures are qualitative patterns, not guaranteed outcomes.

Why stage two is so stable — and so expensive

Stage two is a local optimum. Everything looks finished. Listings are complete, the brand is on the channels competitors are on, and the team is visibly busy. Revenue may even be growing, because channel additions produce their own short-lived lift. The cost of stage two is invisible precisely because it is a cost of omission: the advertising that would not have been needed, the returns that would not have occurred, the inventory that would not have been written down.

Brands leave stage two for one of two reasons. Either a shock forces it — a large partner leaves, a fee change compresses margin, a stockout costs a hero SKU its ranking — or someone writes down a margin floor and starts enforcing it. The first route is more common. The second is cheaper.

The retention lens

Net revenue retention is the honest measure because it nets expansion against churn inside the existing book. A brand programme that adds two channels and loses a major wholesale partner has not grown; it has substituted. Reporting that leads with gross new revenue will show the substitution as growth for at least two quarters.

Where revenue declines are driven by partner or listing churn rather than by execution quality, the correct read is structural, not performance-based — and the correct fix is retention ownership, not more activity.

The compounding lens

Stages four and five compound because their outputs are inputs to each other. Better content raises organic conversion, which lowers advertising load, which frees budget for testing, which improves content. Better forecasting raises in-stock rate, which protects ranking and shortlist inclusion, which raises velocity, which improves the forecast.

Stages one to three do not compound. They are maintenance. This is the single strongest argument for concentrating effort rather than distributing it thinly across every channel available.

16%	250%	10–30%	37%
higher average order value among omnichannel customers versus single-channel[11]	higher engagement reported by retailers operating three or more channels[11]	revenue uplift associated with leading personalisation capability[12]	of consumers rank in-stock reliability in their top three retailer criteria[4]

These figures describe the size of the prize, not the probability of capturing it. Each is contingent on the operating conditions described in chapters three and four. A brand at stage two should read them as an estimate of what is currently being left on the table.

06 Build, hire, or partner

> The real question is not who does the work. It is who carries the inventory risk, and whether the party advising on growth is paid when growth is profitable or merely when it is busy.

Three structures dominate. Each is defensible; each fails differently. What distinguishes them is not capability but the alignment between how the counterparty is paid and what the brand actually needs.

EXHIBIT 8

Three structures, three risk positions

Comparison of cost behaviour, risk transfer and incentive alignment

	BUILD IN-HOUSE	RETAIN AN AGENCY	WHOLESALE-TO-RETAIL PARTNER
Cost behaviour	Fixed. Salaries, tooling and overhead accrue regardless of sell-through.	Semi-fixed. Retainer plus percentage of spend or sales; rises with activity.	Variable by construction. Partner earns retail margin on units actually sold.
Time to capability	12–18 months to reach stage three, assuming the right first hire.	30–90 days to activity. Capability remains outside the business.	30–60 days to activity, using an existing operating standard.
Who holds inventory risk	**The brand.**	**The brand.**	**The partner, on purchased units.**
Who holds working capital	The brand, across production and channel inventory.	The brand, plus fees payable before revenue.	Released on purchase order; the partner funds channel inventory.
Incentive alignment	Aligned in principle; in practice tied to headcount and budget defence.	Paid on spend or activity. Profitable growth and expensive growth pay the same.	Paid on sell-through margin. Unsold inventory is the partner's loss.
Primary failure mode	The first hire is a platform specialist, not an operator; the model never reaches stage three.	Activity substitutes for outcome. Reporting measures work done, not margin earned.	Partner selection is wrong: too little category fit, or unwillingness to hold brand standards.
Best fit	Brands above roughly $25M in channel revenue with an existing data function.	Brands needing a specific capability gap filled for a defined period.	Brands wanting channel revenue without fixed cost, working capital drag or inventory risk.

Seed Ventures LLC analysis. Ranges are directional and vary with category complexity, catalogue size and existing internal capability. Seed Ventures operates the third model.

The question that separates the three

Ask any prospective partner a single question: *what happens to you if these units do not sell?* Under a retainer, nothing happens. Under a percentage-of-spend arrangement, the partner is paid more for spending more. Under a wholesale-to-retail structure, the partner has already paid for the inventory and carries the loss. That is not a moral distinction — it is a structural one, and it determines which conversations are easy to have in month nine.

The corollary matters too. A partner carrying inventory risk will be more selective about assortment, more insistent on pricing discipline, and more willing to say that a SKU should not be launched. Brands sometimes read that as a lack of ambition. It is the alignment working.

> ***Zero management fees is not a pricing gimmick. It is a statement about where the risk sits: if the partner earns nothing until units sell through at margin, the brand is buying an outcome rather than an activity.***
>
> SEED VENTURES LLC

What to test before you sign anything

- **Ask for the margin floor.** If a partner cannot tell you the contribution margin below which they will stop selling a SKU, they do not have one, and neither will you.
- **Ask who writes the content standard.** If the answer is "we follow your brand guidelines," ask to see the channel-specific translation of those guidelines. Brand guidelines are not a listing standard.
- **Ask how forecasting is done and by whom.** If forecasting sits with the same function that executes fulfilment, expect simultaneous overstock and stockout.
- **Ask what they will refuse to do.** A partner with no exclusions has no standard. The specific exclusions tell you what they have been burned by, which is the most useful information available.
- **Ask for the reporting cadence, then ask for last week's.** The artefact that exists without being requested is the one that reflects how the business is actually run.

A note on simplicity in commercial documents

One of the more durable lessons from our own partner development: complexity in the agreement is a leading indicator of a stalled deal. A long exclusive distribution agreement produced very few responses from suppliers. A single-page visual partnership summary — scope, economics, obligations, term — produced conversations, and the detailed paper followed once both sides wanted the outcome.

The same holds for channel proposals. If the value cannot be stated on one page, the thinking behind it is usually not finished.

07 A 90-day sequencing plan

> Sequence matters more than speed. Remediating content before writing the standard means doing the work twice; restructuring advertising before fixing conversion means paying for the old conversion rate more efficiently.

EXHIBIT 9

Ninety days from audit to allocation

Indicative workstream sequencing for a brand at maturity stage two

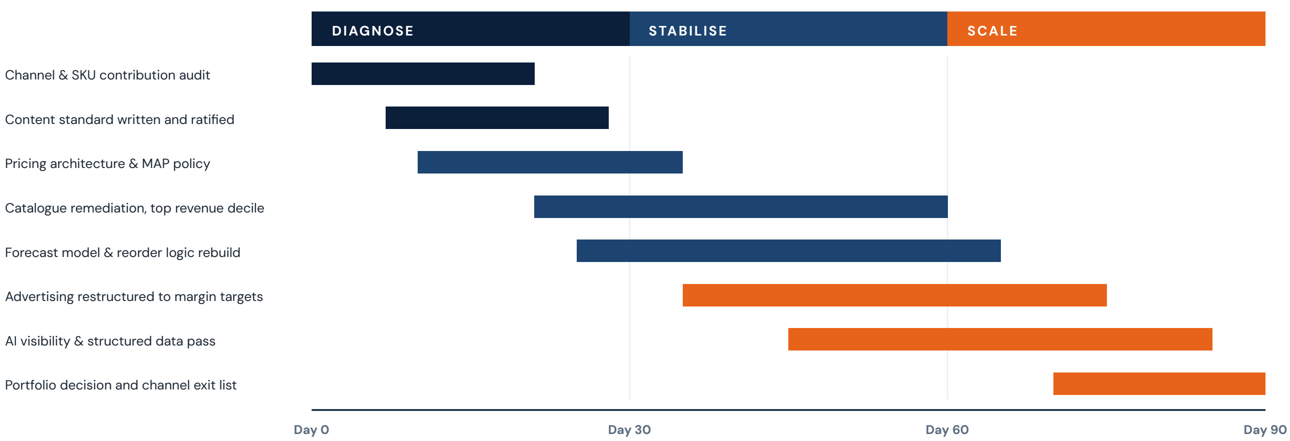

Dependencies are strict: no remediation before the standard exists; no advertising restructure before conversion inputs are fixed.

Seed Ventures LLC. Indicative only; duration scales with catalogue size and the number of active channels. Assumes an existing team of two to four people on channel operations.

PHASE	OBJECTIVE	EXIT CRITERION
Days 0–30 Diagnose	Establish contribution margin per SKU per channel. Write the content and pricing standards. Identify the revenue decile that carries the business.	A single sheet showing every SKU, every channel, and margin against a floor you have committed to.
Days 31–60 Stabilise	Remediate the top decile against the standard. Rebuild forecasting outside operations. Enforce MAP and close pricing inconsistencies across channels.	In-stock rate above target on hero SKUs; zero unresolved cross-channel price conflicts.
Days 61–90 Scale	Restructure advertising against margin rather than revenue. Complete the structured-data pass. Decide which channels to fund, hold or exit.	**A written allocation decision with named owners, and at least one channel explicitly deprioritised.**

The three things that go wrong in the first ninety days

The audit becomes a project. Contribution margin per SKU per channel is not a data-warehouse initiative. It is a spreadsheet built from settlement reports, fee schedules, landed cost and advertising spend, and it should exist in under three weeks in a rough form. Precision can improve later; the decisions it unlocks cannot wait for it.

Remediation starts at SKU one. Catalogues are not uniform. In most portfolios, the top decile of SKUs carries the large majority of contribution. Remediating alphabetically, or by whatever is easiest, spends the scarcest resource — attention — on the least consequential inventory.

Nothing gets stopped. A ninety-day plan that only adds work does not survive contact with a real team. The exit criterion for day ninety deliberately requires that at least one channel is deprioritised. If nothing was stopped, the operating model has not changed; it has only been documented.

Week 1 artefact

Fee-and-cost schedule per channel, agreed with finance. Nothing else is calculable without it.

Week 4 artefact

The content standard: title convention, image sequence, attribute requirements, claim library, per channel.

Week 12 artefact

The allocation memo: what gets funded, what gets held flat, what gets exited, and who owns each.

> ***A ninety-day plan is not a transformation programme. It is the minimum work required before anyone in the business can answer, with evidence, whether the next channel should be added at all.***
>
> SEED VENTURES LLC

On tooling

Automation earns its place after the standard exists, not before. Once title conventions, attribute requirements and image rules are written down, generating and maintaining them across hundreds of listings is a tractable software problem — and one where the time savings are real. Applied before the standard exists, the same tooling simply produces inconsistency faster.

08 Omnichannel readiness diagnostic

> Twelve questions. Answer each strictly — “we could produce that” is a no; only “that exists and someone owns it” is a yes.

	LAYER	QUESTION	YES / NO
01	Strategy	Can you produce contribution margin for any SKU, on any channel, for last month, within one hour?	☐ ☐
02	Strategy	Is there a written margin floor below which a SKU is delisted or repriced, and has it been enforced in the last quarter?	☐ ☐
03	Strategy	Is there a written exit criterion for a channel, and has any channel ever been exited against it?	☐ ☐
04	Execution	Does a documented content standard exist per channel — title convention, image sequence, attribute set, claim library?	☐ ☐
05	Execution	Is query-level performance reviewed on a fixed cadence, with a named owner for organic conversion rate?	☐ ☐
06	Execution	Is advertising budgeted against contribution margin targets rather than revenue or ACoS alone?	☐ ☐
07	Distribution	Does demand forecasting sit outside the function that executes fulfilment?	☐ ☐
08	Distribution	Do you know weeks of cover by SKU by fulfilment node today, without asking anyone to build it?	☐ ☐
09	Distribution	Are supplier terms and order cadence reviewed against demand visibility rather than against convenience?	☐ ☐
10	AI Visibility	Is there a single attribute taxonomy applied across every channel and your own site?	☐ ☐
11	AI Visibility	Have you tested how AI assistants describe and recommend your category, and where you appear?	☐ ☐
12	Company	Can you name the single person accountable for net revenue retention?	☐ ☐

Scoring

0–3 yes — Stage one or two. The immediate priority is the audit and the margin floor. Do not add a channel this year. Do not restructure advertising yet; you do not know what you are optimising toward.

4–7 yes — Stage two to three. Standards exist in parts. The gap is almost always ownership rather than documentation. Assign the four accountable measures in Exhibit 6 to named people before doing anything else.

8–10 yes — Stage three to four. You have a managed programme. The remaining value is in cadence and in the AI visibility layer, which is usually the last to be assigned an owner and the fastest-moving.

11–12 yes — Stage four to five. The constraint is now allocation, not execution. The question shifts from “how do we run this channel better” to “where should the next dollar of inventory and attention go.”

What this adds up to

Omnichannel is no longer a growth strategy. It is the ambient condition of consumer commerce, and the differentiator has moved from presence to discipline.

The market data in chapters one and two describes an environment that rewards a specific and slightly boring capability: clean, consistent, machine-readable product data, governed centrally, deployed channel-specifically, and reviewed on a cadence nobody has to ask for. Demand concentration means the anchor channel is unavoidable. Second-tier growth means it is insufficient. AI-mediated discovery means the quality of the underlying data now determines inclusion in a way that advertising cannot compensate for.

The economics in chapter three explain why most brands do not build that capability: the cost of not having it is a cost of omission, and costs of omission do not appear in management reporting. They appear as advertising load, as returns, as write-downs, and as the quiet substitution of new-channel revenue for churned revenue.

The operating model in chapter four is deliberately unexciting. Four disciplines, four measures, one north star. Its only real claim is that it makes accountability legible — which is the prerequisite for everything else, and the thing most channel organisations lack.

> *If you take one thing from this paper: write down the margin floor, name the owner of net revenue retention, and move forecasting out of operations. Those three moves cost nothing and change what is possible in the following quarter.*
>
> SEED VENTURES LLC

Three questions for your next leadership review

- What is our contribution margin on our tenth-largest SKU, on our second-largest channel, last month — and who produced that number without being asked?
- Which channel would we exit if we had to exit one, and what would that cost us in the following two quarters?
- When an AI assistant is asked to recommend a product in our category, are we in the answer — and who in this business is responsible for that?

A closing caveat

Every benchmark in this paper is directional. Category, price point, unit weight, seasonality and brand equity move these numbers substantially. The frameworks are intended to structure a decision, not to substitute for the arithmetic of your own P&L — which is, in the end, the only analysis that matters.

Notes and sources

1. Marketplace Pulse, “Top 10 E-Commerce Marketplaces in 2026,” March 2026. US gross merchandise volume estimates.

2. Agentic Commerce Review, “AI in Ecommerce in 2026: The New Shopping Funnel,” April 2026, citing the Eight Oh Two AI & Search Behavior Study.

3. EMARKETER, “FAQ on ecommerce marketplaces: The shift from Amazon-only to multi-marketplace,” 2026. Forecast shares of US retail ecommerce sales and US retail media digital ad spending.

4. ICSC and McKinsey & Company, “Shopping in the age of AI: Redefining stores for a new era,” April 2026. Based on the ICSC Consumer Survey of 3,004 US consumers fielded 7–11 January 2026, executive interviews and McKinsey analysis. Agentic commerce sizing of up to $1 trillion in US business-to-consumer retail revenue by 2030 is McKinsey’s estimate as reported in that publication and in contemporaneous coverage.

5. McKinsey & Company research on omnichannel customer behaviour, as reported in retail industry analyses (2025). Omnichannel customers shop approximately 1.7 times more frequently than single-channel customers.

6. Industry analysis of Walmart’s reported fiscal 2026 results; global ecommerce sales of approximately $150 billion, growing more than 20% year over year.

7. Nova Analytics, “TikTok Shop Q1 2026,” May 2026, summarising Charm.io and Wall Street Journal reporting on first-quarter gross merchandise volume and buyer demographics.

8. Nova Analytics, “Ecommerce Statistics 2026,” June 2026, citing Amazon disclosures on shopping-assistant monthly active user growth and third-party share of paid units.

9. McKinsey & Company, “The automation curve in agentic commerce” and “The agentic commerce opportunity,” January 2026. Global estimate of $3–5 trillion of consumer commerce mediated by AI agents by 2030 under moderate scenarios.

10. McKinsey & Company, State of AI survey (2025), as summarised in commercetools, “7 AI Trends Shaping Agentic Commerce in 2026,” April 2026.

11. Capital One Shopping Research, omnichannel statistics compilation, 2026. Order-value premium and multi-channel engagement figures.

12. McKinsey & Company research on personalisation, as reported in retail marketing analyses (2025). Revenue uplift of 10–30% associated with leading personalisation capability.

13. Adyen and Morgan Stanley figures referenced in Exhibit 3 are drawn from published 2026 agentic-commerce benchmark compilations. Market-size estimates from different institutions measure different definitions of the market and are not additive.

Method and limitations

Third-party figures are reproduced as published and have not been independently verified by Seed Ventures. Where a statistic is attributed to a primary research organisation via a secondary compilation, that chain is stated in the note. Marketplace gross merchandise volume estimates in particular vary materially between published sources because platforms define first-party, third-party and hybrid volume differently.

All unit economics, maturity stages, coordination models and sequencing plans presented as Seed Ventures analysis are composite benchmarks and frameworks derived from managed portfolio experience. They are directional and are not category-specific guidance. This document is published for general information and does not constitute financial, legal, tax or investment advice.

ABOUT SEED VENTURES

We are a marketplace growth partner that *carries the risk* alongside the brand.

Seed Ventures LLC helps consumer brands sell across Amazon, Walmart, TikTok Shop and more than sixty additional channels. We operate a wholesale-to-retail model: we purchase inventory, take title, and earn retail margin. Brand partners pay **zero management fees**. If units do not sell through at margin, that is our loss, not theirs.

We also own and operate our own private-label brands, which is where the frameworks in this paper were built and stress-tested — on our own inventory, against our own P&L, before they were ever applied to a partner's catalogue.

01	02	03	04
Strategy	**Execution**	**Distribution**	**AI Visibility**

NEXT STEP

Take the twelve-question diagnostic on page 19 to your next leadership review. If you would like us to run the channel and SKU contribution audit with you, we will do the first pass at no cost.

Seed Ventures LLC · seedventuresllc.com · Innovation Hub product validation tool at seedventuresllc.com/pages/innovation-hub · Marketplace automation platform at assistra.ai