ConConAI

WHITEPAPER · JULY 2026

The AI Era of Commerce.

A specialist-advisor network for the moment of decision - and $CON, the utility token that settles the commerce it creates.

conconai.com

July 2026 · v1.0

FOREWORD

The advisor is back.

Commerce, as personal as it once was. As capable as it now can be.

For most of history, buying well meant one thing: someone who knew. The salesperson at the dealership who knew every model on the lot. The agent who knew the neighborhood street by street. The clerk who knew your taste before you did. That person - a real advisor - is what made consumption *good*. Choosing was not a burden; it was a conversation with someone you trusted.

Then the internet scaled access and, in the same motion, removed the advisor. It replaced the expert with infinite search, filters and reviews, and handed the burden of judgment back to the individual. We gained everything to choose from and lost anyone to help us choose. Deciding became lonely, exhausting, a quiet guess repeated a hundred times a week.

ConConAI carries commerce into the age of AI - and uses it to bring the advisor back. A human expert could never be everywhere, for everyone, in every category. Artificial intelligence can. A category specialist that listens, understands what you actually want, and returns the one thing that honestly fits - with its reasoning attached, the way a great salesperson always did. You describe. It understands. It recommends. **You decide.** The point was never that a machine chooses for you; it is that real guidance has finally returned.

Around these advisors sits an honest marketplace - partners are surfaced only on a genuine match, never for paying to the top - and beneath it a single utility token, **$CON**, that settles the network and funds its beginning. The pages that follow make the case in the language of economics, evidence and mechanism. This one states the belief that animates all of them:

> *Buying well should feel, once again, like walking in and having someone who gets it say: "This one. Here's why."*

ABSTRACT

Abstract

Choosing what to consume is one of the most frequent and most costly activities in a modern economy, yet it remains structurally under-served. Classical information economics shows why: consumers face real search costs (Stigler, 1961), quality is obscured by asymmetric information (Akerlof, 1970), and many purchases are experience or credence goods whose value cannot be verified before - sometimes even after - the transaction (Nelson, 1970; Darby & Karni, 1973). The economic response to these frictions has always been the trusted intermediary, and value has consistently concentrated wherever such an intermediary owns the moment of decision, from price-comparison sites to online travel agencies to marketplace-plus-advertising platforms.

This paper argues that the next such intermediary is emerging in artificial intelligence, and that it is being built at the level of the specialist advisor rather than the generalist chatbot. Consumers already research and compare with AI at scale and increasingly prefer it to search (McKinsey, 2026; Adobe, 2025), but a pronounced trust gap separates research from delegated purchase: in one 2025 study, 73% had used or would use AI to research a product while only 10% had let it complete a purchase (Bain, 2025). That guided, trusted middle is largely unclaimed. ConConAI is a network of category-specialist AI advisors designed to occupy it - each advisor converting an open-ended question into a single, reasoned, honestly-placed recommendation - monetized through curated listings and settled in a native utility token, $CON.

We size the opportunity conservatively (a ~$63.2T consumption economy with a fast-growing AI-guided slice projected at $3-5T globally by 2030), examine the platform and data-network economics that make such a network defensible, and analyze the token economy without evasion - including the token-velocity problem that has undone many payment tokens, and what a fixed-supply, deposit-based design does to address it. We also set out how network fees are reinvested to expand the system. Throughout, we cite counter-evidence where it exists and mark forecasts as forecasts. Nothing herein is an offer or investment advice; $CON is a utility token with a risk of total loss (see §25, §27).

CONTENTS

Contents

PART I

The problem and its economics

Why choosing is costly, quality is hidden, and value concentrates wherever a trusted intermediary owns the moment of decision.

01 · INTRODUCTION

The decision is the product.

Every act of consumption begins as a decision. Whoever makes that decision easier, and can be trusted while doing so, owns the most valuable moment in commerce.

A modern consumer economy is, viewed from a certain height, an enormous machine for producing decisions. Households in 2024 directed roughly **$63.2 trillion** of final consumption expenditure (World Bank, 2024)[1] - and behind every dollar sat a choice: which, among an overwhelming field of alternatives, to buy, book, watch, or fund. The dominant tools built to help with those choices - the search engine and the filterable catalogue - share a defining limitation. They return lists. They transfer the burden of judgment back to the person, who must still read, compare, discount marketing claims, and decide, largely alone.

$63.2T

Household consumption directed by choice, 2024

World Bank, 2024

73% / 10%

Would research with AI vs. let it complete the purchase

Bain, 2025

$3-5T

AI-guided share of consumption, projected by 2030

Projection

This paper makes a single, cumulative argument. Deciding what to consume is an economically costly, trust-dependent problem (Part I). Historically, value has concentrated wherever a trusted intermediary reduces that cost and owns the moment of decision (Part I). The locus of that moment is now shifting to artificial intelligence - but a large gap between researching with AI and letting it act leaves the highest-value, guided middle unclaimed (Part II). That middle sits over a multi-trillion-dollar economy whose AI-mediated share is growing quickly (Part II). ConConAI is designed to occupy it with a network of honest, category-specialist advisors that already operate today (Part III), a position that compounds through two-sided-market and data-learning dynamics (Part IV) and is settled in a native utility token whose demand is tied to real network usage (Part V). We close with an unembellished account of risks and a synthesis of the case (Part VI).

On method and intellectual honesty

An investment thesis is only as strong as the evidence that survives scrutiny. We have therefore tried to hold this document to an academic standard rather than a promotional one. Where a popular claim is stale or contested - the frequently-cited “35% of Amazon sales come from recommendations,” for instance, which traces to a single 2013 McKinsey article with no published methodology (MacKenzie et al., 2013)[2] - we say so. Where the academic record is genuinely mixed, as it is for choice overload (§3), we present the correction alongside the original. Where figures are projections, we label them as such. The aim is not to soften the case but to make it credible: a thesis that names its own weak points is more trustworthy than one that does not, which is, not coincidentally, the same principle on which the product itself is built.

02 · THE ECONOMICS OF DECIDING

Why choosing is costly - and why trust has a price.

Three strands of economic theory, each a Nobel-recognized idea, explain why the market for consumer decisions is structurally imperfect - and why an intermediary that resolves those imperfections captures real value.

Search is not free

The modern economics of information begins with George Stigler's observation that acquiring knowledge about prices and products is itself a costly activity (Stigler, 1961)[3]. A rational buyer searches only until the expected benefit of one more comparison no longer exceeds its cost, which is why identical goods persist at different prices and why most people stop searching well short of the optimum. Stigler's insight reframes discovery as an economic good with a real, if hidden, price. Any mechanism that sharply lowers the cost of reaching a good decision therefore creates genuine surplus - and can capture a share of it. This is the theoretical seat of every discovery intermediary, and it is the first pillar beneath ConConAI.

Quality is hidden - so trust becomes the scarce asset

George Akerlof's "market for lemons" (Akerlof, 1970)[4] showed that when sellers know more about quality than buyers do, the market can unravel: buyers, unable to tell good from bad, will only pay an average price, which drives the best goods out and degrades what remains. The institutional cure for such asymmetry is trusted intermediation - certification, reputation, and credible signals that let quality be priced correctly. The scarce, value-bearing asset in an information-asymmetric market is not the product; it is *trust*. An advisor that can be believed is, in Akerlof's terms, a mechanism for restoring a market that information asymmetry would otherwise impair.

> **The scarce, value-bearing asset in an information-asymmetric market is not the product - it is trust.**
>
> After Akerlof, 1970

Some goods can never be fully judged

The difficulty deepens with the nature of the goods themselves. Phillip Nelson distinguished *search goods*, whose quality can be assessed before purchase, from *experience goods*, which can only be judged after use (Nelson, 1970)[5]. Darby and Karni added a third and harder category -

credence goods, whose quality the buyer struggles to verify even after consumption (Darby & Karni, 1973)[6]. Financial, health, travel, and many high-consideration purchases are substantially experience or credence goods: the buyer often cannot tell, even in hindsight, whether the advice was optimal. It is precisely in these categories that trusted, accountable guidance is most valuable and most under-supplied - and precisely these categories that ConConAI's specialist advisors address.

THE SYNTHESIS

Search is costly (Stigler), quality is hidden (Akerlof), and much of what we buy cannot be verified in advance (Nelson; Darby & Karni). Together these describe a large, permanent market failure in consumer decision-making. The historical response is always the same: a trusted intermediary that lowers search cost and credibly signals quality. The question this paper pursues is who provides that intermediation in the age of AI - and how the value it creates is captured and shared.

03 · THE BEHAVIORAL REALITY OF CHOICE

More options, fewer - and worse - decisions.

The classic evidence that abundance can paralyze is vivid but, on honest inspection, conditional. Stated precisely, it becomes a sharper and more defensible argument.

The famous result

In a now-canonical field experiment, Iyengar and Lepper set up a supermarket tasting booth offering either 24 jams or 6 (Iyengar & Lepper, 2000)[7]. The larger display drew more attention - roughly 60% of passers-by stopped, versus 40% for the smaller one - but converted far worse: about **3%** of those who stopped at the 24-jam display bought, against roughly **30%** at the 6-jam display, an order-of-magnitude difference. Barry Schwartz later popularized the broader claim in *The Paradox of Choice* (Schwartz, 2004)[8]: beyond a point, added options raise anxiety, regret, and paralysis rather than welfare.

The honest correction

Intellectual honesty requires the next sentence. When Scheibehenne, Greifeneder and Todd meta-analyzed 50 experiments (63 conditions, N = 5,036), they found a **mean effect size of virtually zero**, with large variance: some studies reproduced overload, others found none, and some found that more choice helped (Scheibehenne et al., 2010)[9]. Chernev, Böckenholt and Goodman then reconciled the field across 99 observations (N = 7,202): the average main effect is small, but choice overload becomes reliable under four moderators - high *choice-set complexity*, high *decision-task difficulty*, high *preference uncertainty*, and an effort-minimizing *decision goal* (Chernev et al., 2015)[10]. Choice overload, in other words, is not a universal law; it is a conditional effect that appears exactly where consumer decisions are complex, unfamiliar, and effortful.

This is a stronger foundation than the popular myth, not a weaker one. The high-consideration categories ConConAI serves - investing, travel, health, automotive, real estate - are the textbook conditions under which the effect is real: many alternatives, difficult trade-offs, and buyers who often lack a formed preference. We therefore do not claim that choice always overwhelms; we claim, with the meta-analytic literature, that it reliably does so in precisely the decisions an expert advisor is built to resolve. (We treat "decision fatigue" as a supporting but genuinely contested mechanism - the foundational ego-depletion effect failed a 23-lab preregistered replication (Hagger et al., 2016)[11] - and do not rest the argument on it.)

The commercial residue of decision difficulty is measurable. Across aggregated studies, roughly **70%** of online shopping carts are abandoned (Baymard Institute)[12] - a figure that reflects, among other frictions, buyers leaving to keep researching and comparing rather than resolving the choice. Unresolved decisions are unrealized transactions.

04 · THE INTERMEDIARY AT THE DECISION MOMENT

Value concentrates where the decision is made.

Across two decades of digital commerce, the same pattern recurs: the more completely an intermediary owns the decision and the transaction, the larger the share of value it captures.

If theory says a trusted intermediary that lowers search cost should capture value, the history of digital commerce says how much. The evidence forms a ladder. At the lightest touch sit affiliate and comparison models, which merely refer a decided buyer onward and earn commissions typically in the low single to mid-double digits - physical-goods affiliate rates commonly run 5-20% (industry benchmarks)[13]. A step up, online travel agencies intermediate the booking itself: Booking Holdings realized roughly a **14%** blended take on $150.6B of gross bookings in 2023, and Expedia around **12%** (company filings, 2023)[14], with headline hotel commissions of 15-25%. Higher still, a platform that owns discovery, transaction, fulfilment and advertising captures the most: independent reconstruction of Amazon's marketplace economics estimates total seller-side take - referral fees plus fulfilment plus advertising - had climbed past **50%** of third-party seller revenue by 2023 (Marketplace Pulse, 2023)[15], and Amazon's third-party seller-services revenue grew from $53.8B in 2019 to about $172B in 2025 (from company filings)[16].

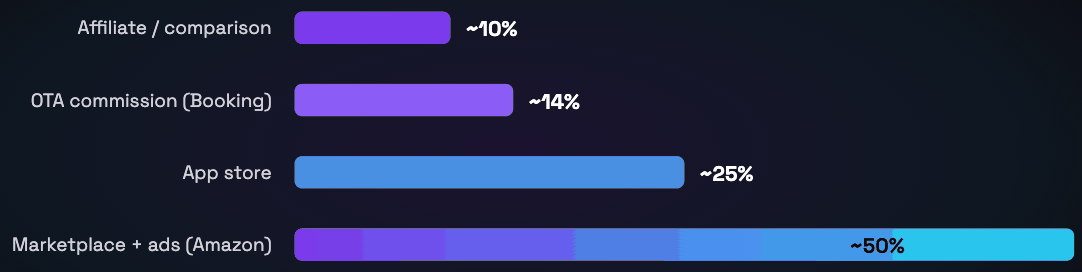

Illustrative take-rate ladder. Value captured rises with how completely the intermediary owns the decision and the transaction.
Sources: company filings; Marketplace Pulse (2023); industry benchmarks.

Two further data points underline the prize. First, ownership of the discovery moment is itself contested at the highest level: roughly half or more of US product searches now begin on Amazon rather than a general search engine (industry surveys, 2023)[17], and the European Commission fined Google **€2.42 billion** in 2017 for preferencing its own Shopping service in

search results - a measure of how valuable, and how contested, control of the decision entry-point has become[18]. Second, the pattern is directional and robust: whoever converts an undecided consumer into a decided buyer sits closest to the money.

WHY THIS MATTERS FOR WHAT FOLLOWS

ConConAI is positioned at the top of this ladder by design. It does not merely refer traffic; it produces the decision - the reasoned recommendation itself - and settles the resulting commerce. The economic case is not that AI advice is novel, but that it occupies the exact point in the funnel where, historically, the largest share of value has always accrued.

PART II

The shift and the opportunity

Decisions are moving to AI - but the guided, trusted middle, sitting over a multi-trillion-dollar economy, is still unclaimed.

05 · THE SHIFT TO AI-MEDIATED DECISIONS

The decision moment is moving - but not all the way.

Consumers already research and decide with AI at scale and trust it more each quarter. Yet a wide gap between researching and delegating leaves the most valuable, guided middle unoccupied.

The behaviors that precede a purchase are migrating from the search box to the AI assistant. In McKinsey's 2026 ConsumerWise research, **68%** of US consumers reported using an AI tool in the previous three months, and **19%** had used AI specifically to discover or decide on brands and products (McKinsey, 2026)[19]. Among AI-enabled shoppers, 62% use it to compare options, and 44% of AI-search users now name it their most-preferred information source, ahead of search engines at 31% (McKinsey, 2026). Adobe reports that **73%** of AI shoppers consider AI their primary product-research source, with generative-AI referral traffic to US retail sites rising sharply off a small base (Adobe, 2025)[20]. Capgemini, surveying 12,000 consumers across twelve countries, found **71%** want generative AI integrated into how they shop (Capgemini, 2025)[21].

Adoption of AI in the consumer journey. Sources: McKinsey ConsumerWise (2026); Adobe Analytics (2025); Capgemini Research Institute (2025).

The trust gap - the unclaimed middle

Adoption of AI for *research* is not the same as delegation of the *purchase*, and the distance between them is the whole opportunity. Bain's 2025 consumer survey (n = 2,016 US) captured it precisely: **73%** had used or would use AI to research products or compare prices, **64%** were open to AI completing a purchase - but only **10%** had actually let it do so (Bain, 2025)[22]. Between broad willingness to be advised and narrow willingness to hand over the transaction lies the guided decision: the point where a consumer wants a trusted recommendation but is not ready to cede control entirely.

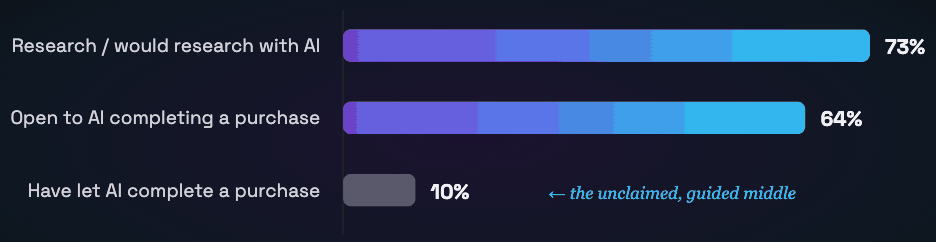

The delegation gap: research with AI is common; letting AI complete a purchase is rare. Source: Bain & Company AI Payments Consumer Survey (2025), n = 2,016 US.

Why the gap? The academic literature on trust in algorithms explains both its existence and its likely closure. Dietvorst and colleagues documented *algorithm aversion*: people abandon an algorithm disproportionately after seeing it err, even when it outperforms humans (Dietvorst et al., 2015)[23]. But Logg and colleagues documented the countervailing *algorithm appreciation*: for many judgment tasks, people actually weight algorithmic advice more heavily than human advice (Logg et al., 2019)[24]. Which force dominates is contextual - it depends on transparency, on whether errors are visible, and on the user's expertise. The design implication is direct: trust is won by advice that explains its reasoning, admits uncertainty, and lets the user stay in control - exactly the posture of a specialist advisor that recommends but does not unilaterally transact. The generalist chatbot answers everywhere and owns nothing; the trusted, category-specialist middle remains open.

06 · SIZING THE OPPORTUNITY

A decision layer over a $63 trillion economy.

The total addressable market is the consumption economy itself; the serviceable slice is the share of it that flows through AI - a slice growing from near zero toward the trillions.

The upper bound is the entire consumption economy: households directed roughly **$63.2 trillion** of final consumption expenditure in 2024, up from $60.4T in 2023 (World Bank, 2024)[1]. Every dollar of it is preceded by a decision, and is therefore a potential ConConAI touchpoint. That figure is the theoretical ceiling (TAM), not a revenue claim; the relevant serviceable market (SAM) is the portion of consumption mediated by AI.

That serviceable slice is where the growth is. McKinsey projects the global agentic-commerce opportunity at **$3-5 trillion** by 2030, with up to $1 trillion of US B2C retail revenue potentially "orchestrated" by AI - explicitly a goods-only, forecast figure (McKinsey, 2025)[25]. Bain estimates US agentic commerce at **$300-500 billion** by 2030, or 15-25% of US e-commerce (Bain, 2025)[26]; Morgan Stanley models a US range of **$190-385 billion** (Morgan Stanley, 2025)[27]. Gartner, in a distinct framing, predicts that by 2030 20% of monetary transactions will be *programmable* to give AI agents economic agency, and that by 2028 AI agents will intermediate the majority of B2B buying, routing over **$15 trillion** of spend (Gartner, 2025)[28]. And the leading edge is already visible: Salesforce reported that AI influenced roughly **19%** of global online orders over the 2024 holiday season, serving tens of billions of recommendations (Salesforce, 2024)[29].

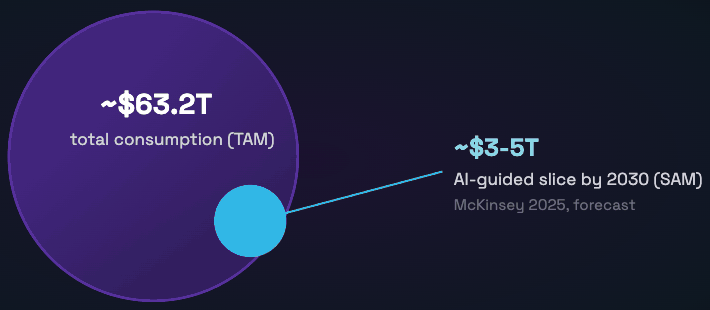

TAM vs. serviceable slice: the consumption economy (~$63.2T, World Bank 2024) and the AI-guided commerce slice projected by 2030 (~$3-5T; McKinsey 2025). Forecast figures; areas proportional.

Two features of these numbers matter for the thesis. First, the estimates diverge in scope and definition - goods vs. services, US vs. global, agent-completed vs. AI-influenced - and should be read as a *direction and order of magnitude*, not a single point. We treat them accordingly. Second, they converge on the same conclusion: the AI-mediated share of consumption is moving from negligible to material within this decade. A network positioned at the decision layer does not need to capture the whole slice; a modest share of a multi-trillion-dollar, fast-growing SAM is a large business.

TAM → SAM → SOM

TAM - the ~$63.2T consumption economy (every decision). **SAM** - the AI-guided slice, ~$3-5T globally by 2030. **SOM** - the share ConConAI can realistically serve, a function of category coverage, advisor quality, and partner supply. Revenue is a take-rate on the transactions the network guides (see §11); token demand is a function of that fee volume (see §18).

07 · THE SEVEN VERTICALS

Seven live markets, each measured in the billions or trillions.

The thesis is not abstract. Seven ConConAI advisors already operate across categories whose individual sizes span from tens of billions to over a hundred trillion dollars.

ConConAI's first seven advisors were chosen to demonstrate the model across structurally different, individually enormous categories. Each is a market in its own right, with a verifiable size drawn from an authoritative source.

ADVISOR	CATEGORY	CATEGORY SIZE (LATEST)	SOURCE
AInvest	Investing	**$139T** global AUM → ~$200T by 2030	PwC (2025)
FlyAwAI	Travel & tourism	**$11.1T** total contribution to GDP (~10% of world GDP)	WTTC (2024)
33Cars	Automotive	**74.6M** new cars sold in 2024 (multi-trillion $ market)	ACEA / OICA (2024)
Dr. Intelligence	Digital health	**$285B** → **$946B** by 2030 (22.2% CAGR)	Grand View Research
Mo' Vie	Film & streaming	**$95.5B** SVOD in 2024 → ~$165B by 2030	Grand View Research
WAV247	Music	**$29.6B** recorded music (+4.8%); 752M paid subscriptions	IFPI (2025)
FindYourCrib	Real estate	Global real estate assets valued at **~$393T** (a stock, not annual revenue)	Savills (2025)

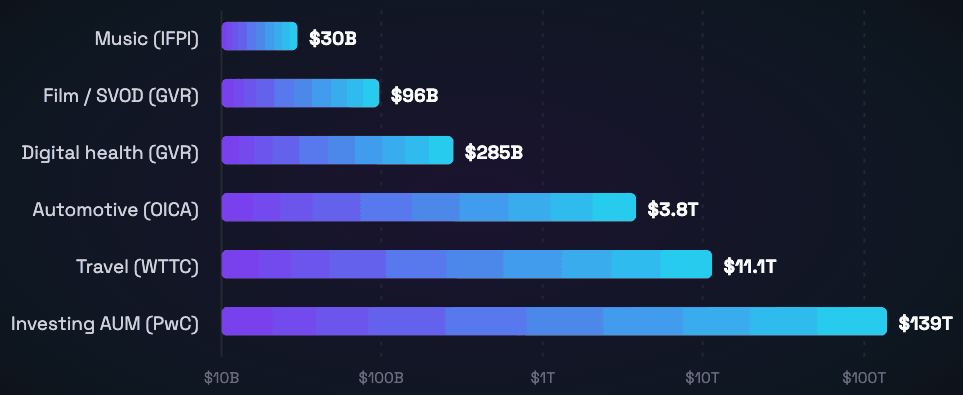

Category sizes on a logarithmic scale (US$). Definitions differ by category - AUM and real-estate figures are asset stocks, others are annual market values - and are not directly additive. Sources as tabled.

Two cautions keep the claim honest. These figures use different definitions - some are annual market values (music, streaming, digital health), one is an activity measure (new-car units), and two are asset stocks rather than flows (AUM; the ~$393T total value of global real estate, which is emphatically not annual revenue). They are therefore illustrative of scale, not a single summable TAM. What they establish is narrower but sufficient: each category ConConAI serves is independently large, structurally different, and a real market in which its advisor already operates. The model is not confined to one niche; it is a repeatable pattern, and seven is the beginning rather than the boundary (see §13).

PART III

The solution: ConConAI

A network of honest, category-specialist advisors that produce the decision - and already operate today.

08 · THE CONCEPT

One brand, a specialist for every decision.

ConConAI is the trusted decision layer that theory predicts should exist - built as a network of category specialists rather than a single generalist.

The preceding parts define a shaped hole in the market: a costly, trust-dependent decision problem (Part I), whose resolution is migrating to AI but stalling at the guided middle (Part II). ConConAI is designed to the shape of that hole. It is a single consumer brand over a network of *category-specialist* AI advisors, each of which converts an open-ended question - "a reliable family SUV under $30k, mostly city" - into a single, reasoned recommendation, with the reasoning attached.

The choice of a specialist architecture is deliberate and maps directly to the theory. A specialist advisor confronts precisely the conditions under which choice overload is real - complex sets, hard trade-offs, uncertain preferences (Chernev et al., 2015) - and resolves them with domain depth a generalist cannot match. It lowers Stiglerian search cost to near zero by replacing a list with an answer. And because it is accountable within a domain and shows its reasoning, it supplies the credible signal that an information-asymmetric, credence-good market requires (Akerlof, 1970; Darby & Karni, 1973). The brand-level network then does what no single-vertical tool can: it accumulates trust and behavioral data *across* categories under one name, so a consumer who trusts the travel advisor arrives pre-disposed to trust the investing one.

MAPPING THEORY TO DESIGN

Search cost (Stigler) → a reasoned answer replaces a list. Asymmetric information (Akerlof) → transparent, accountable recommendations restore a credible quality signal. Credence goods (Darby & Karni) → a domain specialist supplies verifiable rationale. Choice overload under complexity (Chernev) → the advisor operates exactly where the effect is strongest. Each pillar of the product answers a specific, documented market failure.

09 · PRODUCT, EXPERIENCE & HONEST PLACEMENT

Trust is not a slogan; it is the mechanism.

The product is engineered so that being honest is also being profitable - the only configuration in which advice remains trusted, and therefore valuable, over time.

The consumer experience is deliberately narrow: describe the need in plain language; answer a few clarifying questions; receive one recommendation with its justification. This is not a feature list but the enactment of the theory - the advisor absorbs the search and comparison work and returns judgment, which is the scarce good.

The advisor experience: an open question becomes a few smart clarifications, then a single recommendation with its reasoning attached - judgment, not a list.

The commercial design is where trust is either protected or destroyed. Akerlof's lemons logic implies that if recommendations could be bought, buyers would rationally discount all of them, and the advice would lose the very credibility that makes it valuable - the market would unravel. ConConAI therefore separates payment from ranking: partners can list and pay to participate, but *placement is earned by fit, never by bid*. Relevance is gated by genuine match to the stated need. This is not an ethical flourish; it is the equilibrium condition for the asset to retain value. An advisor that can be trusted is, economically, the entire product, and a pricing model that would corrode that trust would corrode the business with it.

This principle also shapes how the token is used. As detailed in Part V, partners settle listing, placement, and performance fees in $CON, but paying in $CON never purchases better ranking. The token is the medium of payment for the network's services; it is never a lever on the advice. Keeping those two things strictly separate is what allows the network to be both monetized and believed.

10 · THE ECOSYSTEM AS PROOF OF CONCEPT

The model already runs.

Seven advisors are live on the web today. The thesis is not a promise to build; it is an operating system to scale.

A whitepaper's weakest form asks readers to fund an idea. ConConAI's advisors - across music, travel, film, health, automotive, real estate, and investing - are already live products on the web, each rendered in its own identity and answering real consumer questions in its domain (see §7 for the categories and their sizes). This materially de-risks the central claim: the pattern of specialist advice, curated listings, and honest placement is demonstrated, not hypothesized. What remains is expansion - more categories, more depth, more reach, and native applications - rather than invention.

The strategic significance is that the seven prove *repeatability*. Because each advisor instantiates the same underlying pattern - ingest a need, ask, reason, recommend, list, settle - adding an eighth or an eightieth does not require a new business model, a new brand, or a new settlement layer. It requires the marginal work of a new vertical on an existing machine (see §13). The proof of concept is therefore also the proof of scalability.

11 · THE BUSINESS MODEL

Paid on the decision, settled in the token.

ConConAI earns where the value ladder of §4 is highest: at the produced decision itself. Revenue is a take-rate on guided commerce, denominated in $CON.

The revenue mechanism is the funnel made legible. Traffic arrives at an advisor with intent; the advisor narrows thousands of options to one reasoned pick; a partner's product surfaces only when it genuinely fits; and ConConAI earns on the resulting click, lead, or sale. Because the network produces the decision rather than merely referring traffic, it sits at the top of the take-rate ladder established in §4 - closer to the marketplace-plus-advertising end than to thin affiliate referral. Performance-based pricing (per click, lead, or action) aligns cost with value delivered and removes upfront risk for partners. In its current phase ConConAI is an advertising and listings platform - it is paid to surface the right option while the purchase completes on the partner's own site; a later, optional phase can process the transaction directly, moving the network up the take-rate ladder into a structurally larger revenue pool. Section 25 quantifies both phases.

Crucially, all partner spend across the network - listing, placement, and performance fees - is denominated and settled in **$CON**. This is the connective tissue between the commercial engine and the token economy: as guided commerce grows, so does the volume of fees that must be transacted in the token. The monetization model plugs into an established, growing market - affiliate and performance-marketing platforms alone are projected to grow from $22.6B (2025) to $35.7B (2033) (Grand View Research)[30] - but captures value at a structurally higher point in the funnel than classical affiliates do.

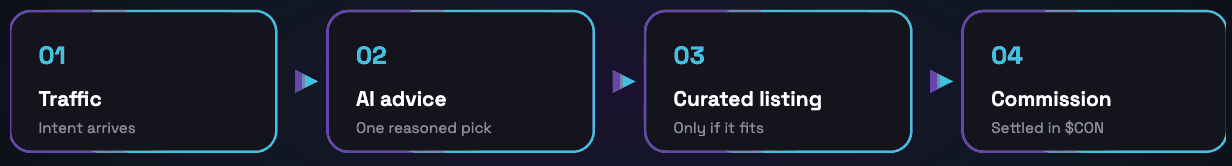

The value flow: intent becomes a reasoned decision, a curated listing, and a commission settled in $CON - each step closer to the money than a classical affiliate referral.

UNIT LOGIC

Network revenue ≈ guided GMV × take-rate. Token demand ≈ (fee volume denominated in $CON) ÷ velocity (see §16-18). The business and the token are two views of one flow: the commerce the network guides.

PART IV

Why it compounds

Two-sided-market dynamics and across-user data learning turn a market position into a defensible, self-reinforcing advantage.

12 · PLATFORM & NETWORK ECONOMICS

Why the position compounds - and where it does not.

ConConAI is a two-sided market with a data-learning core. Both dynamics are real sources of defensibility; both are frequently overstated, so we bound them honestly.

ConConAI serves two interdependent groups - consumers seeking decisions and partners seeking qualified buyers - which makes it a two-sided market in the sense of Rochet and Tirole (2003)[31]. Their central result is that such platforms succeed by getting both sides on board, and that the *structure* of pricing (which side is subsidized) matters as much as its level. ConConAI follows the canonical pattern: the consumer side is free and frictionless to maximize the decisions flowing through the network, while the partner side is monetized. This is not incidental; it is the price architecture the theory prescribes.

A two-sided market: the consumer side is free to maximize decisions and data; the partner side is monetized in $CON. The price structure follows Rochet & Tirole (2003).

The data-learning moat - stated precisely

The more durable advantage is informational. Every decision the network guides teaches it what fits, which sharpens future questions and recommendations. The relevant distinction, from Hagiu and Wright (2023)[32], is between *within-user* learning, which only creates switching costs, and *across-user* learning, in which one user's data improves the product for all - the latter being a genuine, compounding competitive advantage. ConConAI's advisors are of the across-user kind: aggregate decision data lifts recommendation quality network-wide, which attracts more consumers, which yields more data. Better advice lowers customer-acquisition cost as more traffic arrives organically - the measurable signature of network effects (a16z, 2018)[33].

The honest boundary

We resist the usual overclaims. Metcalfe's "value scales with the square of users" is a heuristic, not a law; for large networks an n·log(n) relationship is more defensible (Briscoe, Odlyzko & Tilly, 2006)[34]. Nor do network effects guarantee a winner-take-all outcome: multi-homing and differentiation frequently prevent tipping, and market power is an empirical question, not an automatic consequence. And Hagiu and Wright themselves caution that data network effects are often weaker and more quickly saturating than assumed. We therefore rest the defensibility case on the better-founded mechanism - across-user data learning under one trusted brand - and treat winner-take-all dynamics as a possibility, not a premise.

13 · SCALING & THE ROAD AHEAD

How the concept compounds.

The seven advisors are a starting position. The model is built to widen across categories and geographies, deepen from advice toward the transaction, and open into a platform - each axis reinforcing the token.

1 · More categories, the same machine. Because the pattern - specialist advisor, honest pick, curated listing, fee settled in $CON - is category-agnostic, expansion does not reinvent the product. Insurance, education, home services, consumer electronics, fashion, and B2B procurement are each multi-billion-to-trillion markets that plug into the same backend and the same token. Adjacency is native to network businesses: a network built in one vertical can extend into neighboring ones on the base it has already assembled (a16z, 2018).

2 · Geographic expansion. Consumption is global; the ~$63.2T TAM is worldwide. Advisors localize by language, catalogue, and partner base while sharing one brand, one data-learning core, and one settlement layer.

3 · The data flywheel. More usage yields better questions and sharper matches (across-user learning; §12), which raises conversion and trust, which draws more consumers and more partners, which funds more advisors. This is the network variant of a reinvestment flywheel: build → usage → fees → reinvest → build (see §23).

4 · From advice toward the transaction. Today the trusted middle is guidance; only ~10% of consumers have let AI complete a purchase (Bain, 2025). As trust matures, the advisor can move - with the user's consent - from "here is the pick" toward "shall I complete it?", capturing more of the $3-5T agentic-commerce slice (McKinsey, 2025) and, mechanically, routing more fee volume through $CON.

5 · A platform, not just apps. The listings backend and matching layer can power third-party surfaces - an advice layer other products embed - turning ConConAI from a set of consumer apps into infrastructure for AI-guided commerce, with $CON as its unit of account. Each of these axes increases the volume of commerce that must be paid for in the token; utility demand scales with the network rather than with sentiment.

14 · COMPETITIVE LANDSCAPE & DEFENSIBILITY

Why the middle holds against giants.

The obvious question - what if OpenAI, Perplexity, Google, or Amazon do this? - deserves a direct answer rather than avoidance.

Large incumbents are moving into AI shopping, and it would be dishonest to pretend otherwise. Perplexity and OpenAI have added shopping and checkout features; Amazon has embedded an assistant into its store; Google is reshaping product discovery around AI. Their advantages in distribution and capital are real. But the position ConConAI occupies is defensible for several structural reasons, each grounded in the theory already established.

Generalists own breadth; specialists own the decision. A horizontal assistant answers everything shallowly; the documented value, and the documented conditions for choice overload, live in deep, high-consideration categories where domain specialization wins (§3, §8). **Incentive alignment.** A retailer's assistant recommends within that retailer's inventory and is structurally conflicted; a marketplace's assistant is optimized to sell what the marketplace profits from. ConConAI's honest, relevance-gated placement (§9) is a credible-signal advantage precisely because it is not a store. **Cross-domain trust and data.** A single brand accumulating across-user learning over many verticals builds a compounding asset that a collection of siloed features does not (§12). **A neutral settlement layer.** An open, token-settled network aligns independent partners and advisors in a way a closed corporate platform cannot, and does not require partners to feed a competitor's marketplace.

We do not claim these advantages are insurmountable moats - incumbents may replicate parts of the model, and execution risk is real (§25). The claim is narrower and defensible: the specific combination of category specialization, structurally honest placement, cross-domain data learning, and a neutral token-settled network is difficult for a conflicted generalist to copy without undermining its own model. That combination is where the middle holds.

PART V

The token economy

$CON, the settlement layer of the network - its value accrual, mechanism, and precedents, analyzed without evasion.

15 · $CON - THE LISTINGS-TOKEN PRINCIPLE

A token the network actually needs.

$CON is not an accessory to the product; it is the medium in which the network's core commercial activity - paying to be surfaced at the moment of decision - is transacted.

Across the ConConAI network, partners pay to list products, to be placed, and for performance, and those payments are denominated and settled in $CON. In economic terms $CON is a *utility token*: an instrument whose function is access to and payment for a service, not a claim on an enterprise. This is the same category the EU's MiCA regulation defines as a crypto-asset "intended to provide access to a good or a service supplied by its issuer" (see §22). It confers no equity, dividend, revenue share, or governance right, and it carries a risk of losing all value (§25).

Why denominate network payments in a native token at all, rather than simply invoicing in fiat? Three reasons, in ascending order of importance. First, *alignment*: a shared instrument gives partners, and eventually the wider ecosystem, a common unit tied to the network's activity. Second, *programmability*: on-chain settlement makes fees, deposits, and reconciliations auditable and automatable across many independent advisors and partners. Third, and most fundamentally, a *single neutral settlement layer* across every vertical and geography - one unit of account for AI-guided commerce that no single incumbent controls. The token's reason to exist is therefore endogenous to the business: the network genuinely transacts in it. As the following chapters show, that is a necessary but not sufficient condition for the token to hold value - and we examine the sufficiency question without evasion.

16 · HOW THE TOKEN ACCRUES VALUE

The honest model - including the problem most tokens ignore.

A payment token's value follows from the volume it settles and the rate at which it changes hands. Ignoring the second term is how many token economies quietly failed.

The standard framework for valuing a utility token adapts the equation of exchange, MV = PQ (Burniske, 2017)[35]. Rearranged, the aggregate token value required to support a network economy is **M = PQ / V**, where PQ is the annual on-network economic throughput - here, the fee volume settled in $CON - and V is velocity, the number of times the token turns over per year. In words: the network value the token can support equals annual fee volume divided by velocity.

The uncomfortable term is V, and honesty about it is what separates a credible token thesis from a promotional one. Samani's analysis of token velocity (Samani, 2017)[36] makes the danger explicit: if a token is bought only at the instant of payment and immediately sold by the recipient, velocity is very high, M is correspondingly small, and - his exact warning - "transaction volume could grow a million-fold and network value could remain constant." Buterin reaches the same conclusion by a different route, expressing market capitalization as transaction value multiplied by average holding time; as holding time collapses toward zero, so does value (Buterin, 2017)[37]. A pure pay-and-dump token can process enormous volume and still be nearly worthless.

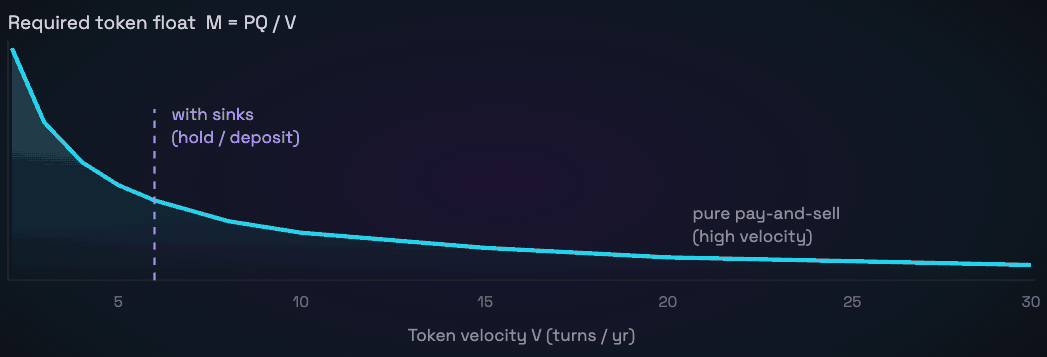

The velocity problem: for a fixed fee volume PQ, the token value the network can support (M = PQ/V) falls as velocity rises. Sinks that induce holding lower velocity and raise supportable value. Illustrative.

The recognized remedies are *velocity sinks* - mechanisms that give holders a reason to hold rather than immediately sell: staking or bonding, required deposits, cash-flow or access rights, and burn-and-mint designs. The lesson from precedent is sobering and instructive. Basic Attention Token reached roughly 100M monthly active users by 2025 while its price fell more than 90% from its peak, because advertisers were never required to hold it; Chainlink introduced “payment abstraction” in 2025 so users need not hold LINK to pay; and Helium’s early network burned only a few thousand dollars of usage per month against a large token supply (see §19). The through-line: making a token easy to avoid holding does not, by itself, create demand. A token economy must be designed so that value scales with usage - which is the subject of the next chapter.

17 · $CON'S MECHANISM VS. THE ALTERNATIVES

Fixed supply, held balances, no dilution.

$CON addresses the velocity problem through a fixed supply and a prepaid-deposit model that requires partners to hold token balances to transact - a design chosen for trust-minimization over mint-based cleverness.

The cleanest textbook answer to "pay for a service in a native token" is a *burn-and-mint equilibrium* (BME), used by Helium and Factom: users burn the native token to mint non-transferable, fiat-priced usage credits, while new tokens are minted on a schedule to reward the supply side, and price settles where value burned balances value minted. BME elegantly ties demand to consumption and shields buyers from volatility. It has one requirement ConConAI deliberately does not meet: it depends on *minting*.

$CON is instead a **fixed-supply token - 100,000,000 units, minted once at genesis, with no mint function, no pause, and no blacklist** (§21). This is a trust-minimization choice: a token that cannot be inflated or frozen is a more credible neutral settlement asset, and avoids the governance surface and securities-adjacent optics of discretionary issuance. Because there is no mint, a pure BME is not available. The velocity sink is provided differently: partners fund network usage by *depositing* $CON into a listings contract and drawing it down as they consume services, which requires them to acquire and hold token balances proportional to their expected activity. Held, prepaid balances reduce effective velocity in the same direction a stake would, while the fixed supply guarantees no dilution of those balances over time. Liquidity-provider tokens are locked at listing (§20), further constraining freely-circulating supply.

DESIGN TRADE-OFF, STATED PLAINLY

A fixed-supply, deposit-based sink is more trust-minimized but a softer velocity brake than a well-run BME, whose deflationary link to usage is tighter. We consider a burn-and-mint or fee-burn layer a legitimate *v2* option should network fee volume warrant it. The present design prioritizes credibility and simplicity: no inflation, no admin power over balances, demand tied to real prepaid usage. We state this trade-off rather than obscure it.

18 · A DEMAND & SENSITIVITY FRAMEWORK

How token demand follows network usage.

We give the structure of the demand model and its levers, not invented figures. Token demand is a transparent function of fee volume, take-rate, adoption, and velocity.

The required aggregate $CON float follows from §16: **required float ≈ (fee volume settled in $CON) ÷ velocity**, where fee volume ≈ guided GMV × take-rate × the share of fees denominated in $CON. A defensible model therefore has four levers, each anchored to observable benchmarks rather than assertion:

LEVER	WHAT IT IS	EMPIRICAL ANCHOR
Guided GMV	Commerce the network guides, as a share of the SAM	AI-guided slice $3-5T by 2030 (McKinsey)
Take-rate	Fee as % of guided GMV	~5% (ad/retail-media) to 10-30% (marketplace); Etsy ~21% effective
$CON share	Portion of fees settled in the token	Network policy (targeting all partner spend)
Velocity	Token turnover per year	~5-20/yr modeling band; sinks lower it (§17)

The point of stating the model this way is discipline, not spectacle. We deliberately do not populate it with a headline "token will be worth X" figure - Burniske himself cautions that "the granularity of our understanding of cryptoeconomics is not there yet" and that such outputs are frameworks, not price targets (Burniske, 2017). What the framework does establish is the *shape* of the relationship: token demand rises with guided GMV, with take-rate, and with the $CON-settled share, and falls with velocity. The design choices in §17 push three of those four levers in the value-supporting direction (fixed supply, held deposits, all-fees-in-$CON), and the strategy in §13 pushes the fourth (GMV) by widening the network. The sensitivity that matters most - velocity - is precisely the one the deposit model is built to contain.

19 · PRECEDENTS - WHAT WORKS, WHAT FAILS

Fee-linked tokens can win; many tokens have lost.

Real designs that tie token demand to real revenue have created durable value. Designs that decoupled the token from usage - or misallocated proceeds - have not. We present both.

Where the design works

When a token is bound to genuine, growing usage, value can accrue. Binance's BNB permanently reduces supply toward a 100M floor through algorithmic and fee-based burns - individual quarterly burns have exceeded $1 billion in value[38]. Ethereum's EIP-1559 has burned roughly 4.6M ETH by tying destruction to blockspace demand - though, in the interest of honesty, we note the "ultrasound money" narrative has weakened: post-2024 upgrades reduced burn and ETH turned mildly net-inflationary in 2026[39]. In late 2025 Uniswap governance approved routing protocol fees to a burn, linking usage directly to supply reduction[40], and Aave, dYdX, and Maker run revenue-funded buybacks. The common denominator of the durable cases is a real, sizable revenue or usage stream feeding the token - not emissions or narrative.

That revenue mechanic is exactly the take-rate-on-GMV engine of §11, and its scaling is well-precedented outside crypto: Apple's App Store facilitated roughly $1.3T in billings in 2024, and Amazon's third-party seller-services revenue compounded from $53.8B (2019) to ~$172B (2025) (§4). Fee revenue that rides transaction volume is a proven way to grow with a network.

Where tokens fail - and how $CON differs

Credibility requires naming the failures. Terra/LUNA erased roughly $45 billion in a week when a yield subsidized beyond organic revenue collapsed (Richmond Fed, 2022)[41]. EOS raised on the order of $4 billion and its own foundation head later called it "a failure." Academic analysis found only about 44% of 2017-era ICOs were still active four months after their sale (Benedetti & Kostovetsky)[42]. And the pay-with-token precedents of §16 - BAT, Chainlink, early Helium - show usage and price decoupling when the token is abstracted away from real, held demand. $CON is designed against these failure modes: no inflationary emissions to prop up yield (fixed supply), demand tied to real prepaid listing fees rather than subsidy, no discretionary admin control, and proceeds reinvested into building the network that generates the fees (§23) rather than into unsustainable rewards. None of this guarantees success - it removes the specific mechanisms by which comparable tokens most often failed.

20 · TOKENOMICS

Fixed supply, transparent allocation.

One hundred million $CON, minted once. Half funds the public presale; the remainder covers the seed round, exchange liquidity, and a modest team allocation.

100M
Fixed total supply, minted at genesis; no mint-on-sale

$0.005
Phase 1 presale price, rising to $0.009 across five phases

$0.01
Targeted DEX listing price

ALLOCATION	SHARE	TOKENS	NOTE
Presale (Phases 1-5)	50%	50,000,000	$0.005 → $0.009
Seed	20%	20,000,000	$0.004, outside the ICO
DEX liquidity	15%	15,000,000	Paired at listing; LP tokens locked
CEX liquidity	10%	10,000,000	Exchange liquidity
Team	5%	5,000,000	Modest allocation

The presale follows a seven-step price ladder from the $0.004 seed through five public phases ($0.005 to $0.009) to a targeted $0.01 listing. A stacking purchase bonus rewards larger cumulative contributions - totalling an additional 50,000 / 200,000 / 600,000 $CON at the $10k / $25k / $50k thresholds respectively, drawn from the presale allocation and fixed in the contract. Purchased $CON is booked at the time of purchase and becomes *claimable at the DEX listing*; it is not transferred at the moment of sale. Liquidity provided at listing is paired at the target price and the resulting LP tokens are locked, constraining freely-circulating supply. Every parameter here is enforced by the deployed, verified contract rather than by promise (§21).

21 · SECURITY & AUDIT

Trust-minimized, and verifiable.

The token and sale are engineered so that no party - including the team - can quietly change the rules, and the code is open to inspection.

$CON's contract was independently audited by SolidProof with **no critical, high, or medium-severity findings** - only minor optimization and informational notes - and is verified on Etherscan, so anyone can read the exact code that runs. The token itself is deliberately powerless in the ways that matter: a fixed supply minted once at genesis, *no mint function*, no pause, no blacklist, and no owner authority over balances. It implements standard ERC-20 behavior plus gasless approvals (EIP-2612).

The sale contract is equally conservative. Purchases are *booked, not transferred*, with claiming enabled only after the presale; proceeds are held for a treasury and the sale can never touch buyers' claimed allocations; and the code uses reentrancy protection, safe token transfers, two-step ownership, and one-way state flags, with full test and invariant coverage. An audit reduces but does not eliminate risk, and no audit is a guarantee (§25) - but the combination of a trust-minimized token, a conservative sale, third-party review, and on-chain verifiability places $CON well outside the design patterns that produced the failures catalogued in §19.

22 · REGULATORY CONTEXT

Where a utility token sits - factually.

This chapter states the regulatory landscape factually. It is not legal advice, and independent counsel review precedes any token sale.

In the European Union, the Markets in Crypto-Assets Regulation (MiCA, Regulation (EU) 2023/1114) places utility tokens in its residual "other crypto-assets" category, defining a utility token as one "intended to provide access to a good or a service supplied by its issuer" (Art. 3). The principal obligation is to draw up, notify, and publish a crypto-asset whitepaper; notably, this is a *notification* regime, not a prior-approval one, which distinguishes it from a securities prospectus. MiCA entered into force in June 2023, with the provisions relevant to utility tokens applying from 30 December 2024[43].

In the United States, whether a token is a security turns on the *Howey* test (SEC v. W.J. Howey Co., 1946): an investment of money in a common enterprise with a reasonable expectation of profit derived from the efforts of others[44]. The SEC's 2019 "Framework for Investment Contract Analysis of Digital Assets" applied Howey to tokens, weighing factors such as whether a network is functional and a token consumptively usable - but that framework was superseded by replacement guidance in March 2026, and case law (notably the 2023 Ripple decision) remains genuinely contested, with classification turning on the specific transaction rather than the token as an object[45]. "Utility token" is not a formal US statutory category; labeling does not by itself determine treatment. ConConAI positions $CON as a utility token on the substance set out in §15-17 - a consumptive means of payment on a functioning network - but the regulatory landscape is uncertain and evolving, and participants bear their own compliance responsibility (§25). Counsel review is mandatory before any live sale.

PART VI

Execution and roadmap

How the business runs and where it goes next - a project-wide view: the reinvestment flywheel that funds expansion, and the concrete path from today's working foundation to the long-term decision layer.

23 · EXECUTION: TURNING FEES INTO AN EXPANDING NETWORK

The reinvestment flywheel.

Capital - presale proceeds and, increasingly, recurring network fees - is reinvested to widen and deepen the network, which raises fee volume, which funds the next expansion.

A token economy is only as strong as the business beneath it, so capital allocation is a first-order concern rather than an afterthought. The operating loop is **build → usage → fees → reinvest → build**: proceeds fund product and engineering, which expand the advisors and add verticals; better and broader advice draws more consumers and partners; their activity generates fees settled in $CON; and those fees fund the next cycle. This is the network expression of the reinvestment doctrine Amazon articulated - prioritize long-run free cash flow and reinvest it rather than optimize reported profit (Bezos, 1997/2004)[46] - and it is why network-effect businesses have accounted for a disproportionate share of value creation in technology (NfX, 2019)[47].

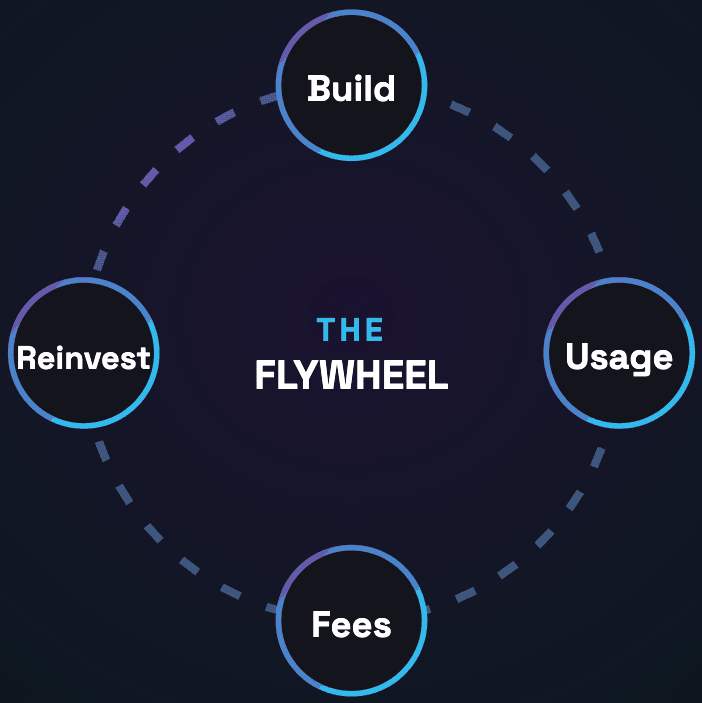

The reinvestment flywheel: fees settled in $CON fund the build, which grows usage, which generates more fees - one motion, not two.

Where the capital goes

The allocation follows the stage-appropriate pattern established by software and marketplace benchmarks. Early on, spend is weighted toward *product and engineering* - deepening the advisors, building the machine-learning and recommendation systems, and standing up each new vertical; below roughly $5M in recurring revenue the majority of operating expense typically goes to R&D, with sales and marketing overtaking only past scale (Scale Venture Partners; ICONIQ)[48]. At maturity the mix converges toward the SaaS median of roughly 26% R&D and 48% sales-and-marketing (Blossom Street Ventures)[49]. Three destinations recur: *development and engineering* (advisors, ML/recsys teams, new verticals, native apps); *growth and user acquisition* (more buyers → more listings → more fees); and *operations and liquidity* (running the network and supporting healthy token markets).

Mature-stage operating-expense split benchmark (R&D 26% / S&M 48% / G&A 21% / other). Source: Blossom Street Ventures SaaS analysis. Early-stage mixes are more R&D-weighted.

Efficiency is disciplined against a recognized guardrail - the Rule of 40 (revenue growth plus profit margin ≥ 40%) - with the honest caveat that only about a third of software companies achieve it in a given period (McKinsey; Bessemer)[50], so it is a bar to aspire to rather than a promise. Precedent for reinvesting token proceeds into building exists and is instructive: the Ethereum Foundation spent roughly $48M in 2021 (about $21.8M on core R&D and $19.6M in grants), and ecosystem programs from the Uniswap Foundation ($26M committed in grants in 2025), Optimism (RetroPGF), and Polkadot ($45M Decentralized Futures) show fees and treasuries funding developers at scale[51]. The essential point is structural: because network fees are denominated in $CON, growth of the product and demand for the token are not two objectives but one motion - every reinvested dollar that expands the network expands the volume of commerce that must be settled in the token.

24 · ROADMAP

From foundation to the decision layer.

Where this network has come from, what the presale funds next, and the long-horizon ambition the entire thesis is built toward. The phases below are forward-looking statements of plan and intent, not guarantees, fixed dates, or a promise of value or return.

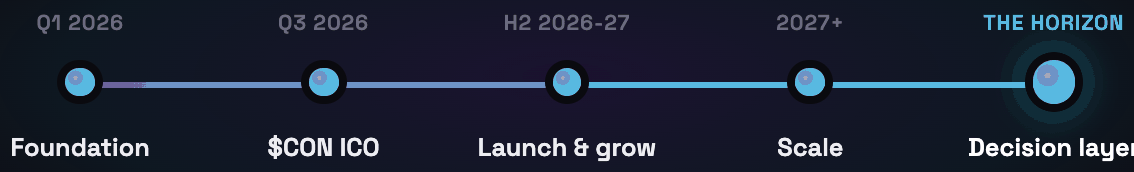

The trajectory: a working foundation built through 2026, funded to launch by the $CON presale, growing into a multi-vertical network, and aimed at becoming the trusted decision layer over the consumption economy - settled in $CON.

Q1 2026 · Foundation (in progress)

The groundwork is already built, not merely promised. Seven specialist AI advisors are live on the web; the ConConAI Listings system and partner console exist, with an organic-first ranking engine that surfaces the best fit rather than the highest bidder (§9-13). On-chain, the fixed-supply $CON token is deployed on Ethereum, independently audited by SolidProof, and verified on Etherscan (§16-18). This phase is about proving the machine works before asking anyone to fund its expansion.

Q3 2026 · The $CON ICO

The presale opens 1 August 2026 on a five-phase price ladder from $0.005 to $0.009, followed by a targeted DEX listing at $0.01 with locked liquidity and tokens claimable at listing (§19-21). Proceeds are allocated across product, apps, the listings platform, marketing, and operations. $CON is a utility token - the means of paying for listings and placements on the network - not an investment instrument (§22).

H2 2026 - 2027 · Launch and grow

Native iOS and Android apps ship to the App Store and Google Play. Listings and advertising roll out across every advisor, staying organic-first: promoted slots are always clearly labelled and never outrank a genuinely better match. Marketing brings buyers and supply partners onto the network in parallel, so both sides of the flywheel thicken together (§13, §23). The partnerships pursued in this phase fall into two tracks: supply and distribution partners by vertical (real-estate portals, OTAs and airlines, car marketplaces, labels and ticketing,

regulated brokers, fintechs and telehealth providers), and crypto trust and reach (audit trust badge, DEX and liquidity locker, CoinGecko and CoinMarketCap listings, wallet and oracle infrastructure, and later centralized-exchange coverage).

2027+ · Scale

Self-serve partner onboarding opens the network to supply at scale. New advisor categories extend beyond the first seven; matching and recommendations deepen as the data-learning flywheel compounds (§13). $CON utility widens to more of the surfaces where it is the unit of payment, liquidity and exchange coverage grow deliberately rather than hastily, and the product expands into more languages and markets - with privacy-safe demand-intelligence and a partner API turning the network into infrastructure others build on.

THE HORIZON · THE MEGA-PROJECT

The endpoint the whole thesis points toward is singular: **ConConAI as the trusted decision layer over the consumption economy itself.** Advice matures - always with the user's consent - from "here is the right pick" toward completing the purchase, positioning the network to capture a share of the multi-trillion-dollar agentic-commerce flows projected for the end of the decade (§6-8). The listings backend and matching API become **infrastructure that other products embed** - an advice-and-decision layer for AI-guided commerce, spanning dozens of verticals and many geographies. And beneath all of it, **$CON is the unit of account**: every listing, placement, and transaction across the network settles in the token, so its utility demand scales with the size of the network rather than with market sentiment (§22). That is the mega-project - a global, open, honestly-run layer for how the world decides what to consume, settled in $CON.

Forward-looking statement. This roadmap reflects current plans and intentions and is subject to change; it is not a guarantee, a fixed schedule, or a promise of value, utility timing, or return. Partnerships named are categories being pursued, not confirmed relationships. See the risk & compliance note in §25.

PART VII

Close

The opportunity quantified in numbers, the argument drawn together into a single conclusion, and the sources in full.

25 · THE OPPORTUNITY, QUANTIFIED

What the take-rate becomes at scale.

The model of §11 expressed in numbers - transparent, scenario-based, and built from the market sizes already cited rather than from optimism. It runs in two phases: an advertising platform today, direct transactions tomorrow.

ConConAI monetizes in two distinct phases, and the difference between them is the difference between a large business and a category-defining one. **Today it is an advertising and listings platform:** the network influences the purchase decision and earns a take on the guided GMV, while the transaction itself completes on the partner's own site. That is the live model of §11 - it captures a slice of *advertising* economics. **The second phase is additive and optional:** as trust, integrations, and payment rails deepen, ConConAI can process the purchase directly - earning a transaction take on the GMV that flows *across* its rails rather than a placement fee on the GMV it merely influences. That reaches into the far larger pool of *transaction* economics, the highest rung of the take-rate ladder (§4). The roadmap in §24 is the path between the two.

The scenarios below are illustrative and deliberately transparent. They express guided GMV as a modest share of the AI-guided commerce slice sized in §6 ($3-5T globally by 2030), apply a blended advertising take-rate of roughly 4% consistent with the network's position on the take-rate ladder (§4, §11), and - because every fee settles in $CON - read network revenue as a floor on annual $CON settlement volume (§16-18). They are scenarios that show the shape of the opportunity, not forecasts or promises.

MILESTONE	GUIDED GMV	NETWORK REVENUE (~4%)	$CON SETTLED / YR
Year 1 · launch	~$200M	~$8M	**~$8M**
Year 3	~$3B	~$120M	**~$120M**
Year 5	~$18B	~$720M	**~$720M**
At scale · ~1% of the SAM	~$50B	~$2.0B	**~$2.0B**

Illustrative base-case scenario (Phase 1, advertising). Guided GMV as a rising share of the $3-5T AI-guided slice; ~4% blended take; all fees settled in $CON. Figures are rounded scenario values, not forecasts or a promise of value.

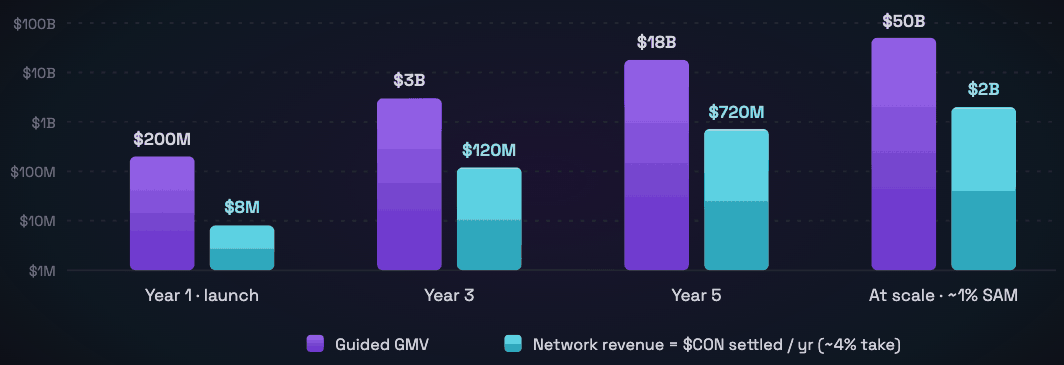

Guided GMV and the network revenue settled in $CON, base case, on a log scale. Advertising phase only; the transaction phase is additive. Illustrative scenario values.

Two things make these figures a floor rather than a ceiling. **First, the "at scale" line is deliberately modest:** ~1% of a $3-5T slice for a category-defining decision layer already implies a multi-billion-dollar revenue base and billions in annual $CON settlement against a *fixed* 100M supply - the mechanical link between network growth and token demand set out in §16-18. **Second, the transaction phase multiplies the pool.** If ConConAI later processes even a fifth of the GMV it guides at a ~6% transaction take, the "at scale" line alone adds on the order of **$600M** in further annual revenue - and routes an order of magnitude more *gross* volume across $CON rails, because settlement then scales with processed value, not merely with advertising fees. For orientation only: guiding a single percent of Gartner's projected $15T of AI-intermediated spend (§6) would be a $150B-GMV network. The ceiling here is set by execution and adoption, not by the size of the prize.

RISK & COMPLIANCE, IN BRIEF

These are illustrative scenarios, not forecasts, guarantees, or a promise of value or return; actual results depend on execution, adoption, and market conditions, and could be materially lower. $CON is a utility token that settles network fees - not an investment - and carries a risk of loss; its supporting value depends on real fee volume and on the velocity dynamics analyzed candidly in §16-18. The regulatory treatment of utility tokens is evolving (§22), and smart contracts carry inherent risk that an audit reduces but does not remove (§21). The full Token Sale Terms and Risk Disclosure are at conconai.com; nothing here is financial, investment, or legal advice.

26 · CONCLUSION

The case, in one thread.

Not "the best product," but a chain of evidence that leads to a single, defensible conclusion.

The argument of this paper is cumulative, and each link rests on cited evidence rather than assertion. Choosing what to consume is a costly, trust-dependent problem - a conclusion that follows directly from the economics of search, asymmetric information, and credence goods, and from the conditional-but-real behavioral literature on choice under complexity. Historically, value has concentrated wherever a trusted intermediary lowers that cost and owns the moment of decision, and the take-rate rises the more completely that moment is owned. That moment is now migrating to artificial intelligence, but a wide gap between researching with AI and delegating to it leaves the guided, trusted middle unclaimed - a middle that sits over a ~$63.2T consumption economy whose AI-mediated slice is projected in the trillions by 2030.

ConConAI is built to occupy that middle: a network of honest, category-specialist advisors, already live across seven markets, positioned at the highest-value point of the funnel, and compounding through two-sided-market and across-user data-learning dynamics that we have bounded rather than exaggerated. Its token, $CON, is the settlement layer for the commerce the network creates - a fixed-supply utility token whose demand is tied to real, prepaid network fees, designed explicitly against the velocity trap and the misallocation failures that have undone comparable tokens, and reinvested through a build-usage-fees-reinvest loop in which growing the product and growing token demand are the same motion.

None of this is a promise of returns, and the scenarios in §25 are illustrations rather than guarantees. But the conclusion the evidence supports is not a timid one: a trusted decision layer over the consumption economy is economically warranted, already live, and - at even a low-single-digit share of a multi-trillion-dollar market - a network of genuine scale, with a clear path from influencing purchases to processing them. $CON is the instrument through which that commerce is settled and its expansion funded, and because every fee settles in a fixed-supply token, its demand grows with the network rather than with sentiment. The document has argued throughout by evidence rather than adjective - the same standard on which the product itself asks to be trusted.

BEFORE YOU PARTICIPATE

Read the Token Sale Terms and the Risk Disclosure at conconai.com. $CON is a utility token, not an investment; it carries a risk of total loss. Nothing in this document is financial, investment, or legal advice, and it is not an offer where such an offer would be unlawful.

27 · REFERENCES

Sources.

Primary and peer-reviewed sources are prioritized. Figures dated after 2024 include forecasts; market-research figures are cited to their originating organization. Date-sensitive on-chain figures should be re-verified against live sources.

1. World Bank. *Households and NPISHs final consumption expenditure (current US$)*, indicator NE.CON.PRVT.CD, World. 2024 (~$63.2T; 2023 ~$60.4T). data.worldbank.org/indicator/NE.CON.PRVT.CD
2. MacKenzie, I., Meyer, C., & Noble, S. "How retailers can keep up with consumers." *McKinsey & Company*, 2013. (Origin of the "35% of Amazon sales from recommendations" figure; dated, unverified.)
3. Stigler, G. J. "The Economics of Information." *Journal of Political Economy*, 69(3), 1961.
4. Akerlof, G. A. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3), 1970.
5. Nelson, P. "Information and Consumer Behavior." *Journal of Political Economy*, 78(2), 1970.
6. Darby, M. R., & Karni, E. "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics*, 16(1), 1973.
7. Iyengar, S. S., & Lepper, M. R. "When Choice is Demotivating." *Journal of Personality and Social Psychology*, 79(6), 2000.
8. Schwartz, B. *The Paradox of Choice: Why More Is Less.* HarperCollins, 2004.
9. Scheibehenne, B., Greifeneder, R., & Todd, P. M. "Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload." *Journal of Consumer Research*, 37(3), 2010.
10. Chernev, A., Böckenholt, U., & Goodman, J. "Choice Overload: A Conceptual Review and Meta-Analysis." *Journal of Consumer Psychology*, 25(2), 2015.
11. Hagger, M. S., et al. "A Multilab Preregistered Replication of the Ego-Depletion Effect." *Perspectives on Psychological Science*, 11(4), 2016.
12. Baymard Institute. *Cart & Checkout Abandonment Rate Statistics* (aggregate ~70%). baymard.com/lists/cart-abandonment-rate
13. Industry affiliate-commission benchmarks (physical goods ~5-20%), 2024-2026.
14. Booking Holdings FY2023 10-K; Expedia Group FY2023 results (implied take rates ~14% / ~12%).
15. Marketplace Pulse. "Amazon Takes a 50% Cut of Sellers' Revenue." 2023. (Third-party estimate; Amazon disputes framing.)
16. Amazon.com Inc. third-party seller-services revenue, from 10-K filings (2019 $53.8B → 2025 ~$172B), via Marketplace Pulse.
17. PowerReviews / eMarketer product-search origination surveys, 2023 (~50-63% start on Amazon).
18. European Commission. Antitrust decision, Google Search (Shopping), €2.42B fine, IP/17/1784, 2017.
19. McKinsey & Company. ConsumerWise, "Embracing AI-supported shopping," 2026.
20. Adobe Analytics. "Generative AI-powered shopping rises," 2025.
21. Capgemini Research Institute. "What Matters to Today's Consumer" (71%), 2025 (survey of 12,000, 12 countries).
22. Bain & Company. "Agentic AI Commerce Hinges on Consumer Trust," 2025 (AI Payments Consumer Survey, n = 2,016 US).
23. Dietvorst, B. J., Simmons, J. P., & Massey, C. "Algorithm Aversion." *Journal of Experimental Psychology: General*, 144(1), 2015.
24. Logg, J. M., Minson, J. A., & Moore, D. A. "Algorithm Appreciation." *Organizational Behavior and Human Decision Processes*, 151, 2019.
25. McKinsey QuantumBlack. "The agentic commerce opportunity," 2025 ($3-5T global; up to $1T US, goods only; forecast).

26. Bain & Company. "2030 Forecast: How Agentic AI Will Reshape US Retail," 2025 ($300-500B; forecast).

27. Morgan Stanley Research. "Here Come the Shopping Bots," 2025 ($190-385B US; scenario range).

28. Gartner. Top Strategic Predictions, 2025 (20% of transactions programmable by 2030; $15T B2B agent-intermediated by 2028).

29. Salesforce. 2024 holiday shopping data (AI influenced ~19% of global online orders; ~60B recommendations, Cyber Week).

30. Grand View Research. *Affiliate Marketing Platform Market* ($22.6B 2025 → $35.7B 2033, 5.9% CAGR).

31. Rochet, J.-C., & Tirole, J. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1(4), 2003.

32. Hagiu, A., & Wright, J. "Data-enabled learning, network effects, and competitive advantage." *RAND Journal of Economics*, 54(4), 2023.

33. Jin, L., & Coolican, D. "16 Ways to Measure Network Effects." a16z, 2018.

34. Briscoe, B., Odlyzko, A., & Tilly, B. "Metcalfe's Law is Wrong." *IEEE Spectrum*, 2006.

35. Burniske, C. "Cryptoasset Valuations." 2017 (M = PQ/V framework). See also Burniske & Tatar, *Cryptoassets*, 2017.

36. Samani, K. "Understanding Token Velocity." Multicoin Capital, 2017.

37. Buterin, V. "On Medium-of-Exchange Token Valuations." 2017.

38. BNB Chain. Quarterly BNB burn reports, 2026 (auto-burn toward 100M floor; BEP-95 real-time burn).

39. ultrasound.money; Ethereum EIP-1559 (base-fee burn since Aug 2021; net issuance mildly positive in 2026).

40. Uniswap. "UNIfication" governance proposal (protocol fees → burn), 2025.

41. Federal Reserve Bank of Richmond. Economic Brief 22-24, "Terra/UST collapse," 2022 (~$45B erased).

42. Benedetti, H., & Kostovetsky, L. "Digital Tulips? Returns to Investors in Initial Coin Offerings," 2018 (~44% of ICOs active after 4 months).

43. Regulation (EU) 2023/1114 (Markets in Crypto-Assets, MiCA). In force June 2023; utility-token provisions apply from 30 Dec 2024.

44. SEC v. W. J. Howey Co., 328 U.S. 293 (1946).

45. SEC. "Framework for 'Investment Contract' Analysis of Digital Assets," 2019 (superseded 2026); SEC v. Ripple Labs, SDNY, 2023 (contested).

46. Bezos, J. Amazon shareholder letters, 1997 & 2004 (free-cash-flow / reinvestment doctrine), SEC filings.

47. Currier, J. "The Network Effects Manual / ~70% of value in tech." NfX, 2019.

48. Scale Venture Partners (Scale Studio); ICONIQ Growth, SaaS spend-allocation benchmarks, 2024.

49. Blossom Street Ventures. SaaS opex analysis (median R&D 26% / S&M 48% / G&A 21%).

50. McKinsey, "SaaS and the Rule of 40"; Bessemer Venture Partners, "The Rule of X," 2024.

51. Ethereum Foundation Report (April 2022); Uniswap Foundation (2025); Optimism RetroPGF; Web3 Foundation / Polkadot Decentralized Futures.
