AI-Era Attribution Blueprint

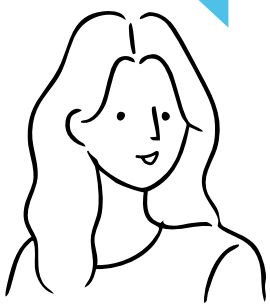
November 2026

Table of Contents

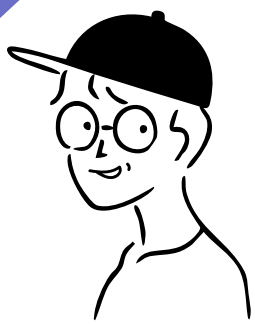
Executive Summary..... pg. 3
Contributing Industry Leaders..... pg. 6
The Attribution Challenge..... pg. 8
AI Influence Pathways..... pg. 11
AI Credit Claim Framework..... pg. 17
Closing the Measurement Gap..... pg. 23

Maximize this Blueprint

How can I use AI to get the most out of this blueprint?



Ask the AI platform of your choice questions like: **A platform says AI drove 20% of my sales. What data should I ask for to validate it?**



Executive Summary

iab.

AI-Era Attribution Blueprint

Executive Summary

AI is changing who, or what, ignites interest, makes a selection, and completes a purchase. The attribution trail the industry built on impressions, clicks, and referrals thins or disappears in these journeys.

> The consequence is not that AI-influenced sales vanish from reporting. They are often reported under the wrong source of influence.

A sale shaped by an AI recommendation may arrive as search or direct traffic. Dashboards credit what they can observe, and the next marketing investment may then follow the wrong channel.

This blueprint gives the industry a shared way to classify those conversions and evaluate claims for credit. It does two things.

1. **It classifies the role AI played in a conversion.** Three questions are applied at the point of sale:
   - Who or what created the intent?
   - Who or what made the decision?
   - Who or what executed the transaction?

   These questions classify the role played in the conversion, but they do not imply that intent or demand has a single causal source.

   Every conversion receives one of three pathways:
   - **Human-Owned**
   - **AI-Assisted**
   - **AI-Executed**

   Within AI-Assisted, the framework distinguishes between AI-Influenced conversions, in which AI shapes intent but the human makes the decision, and AI-Recommended conversions, in which AI shapes the choice itself. Within AI-Executed, it distinguishes between specific instructions, in which the human has already chosen what to buy, and open instructions, in which AI makes the selection. These pathways describe what actually happened, not what current measurement tools can observe, so the categories stay fixed as measurement catches up. Whether a transaction is on- or off-platform, and whether a recommendation is paid or organic, adds context but does not change the pathway.

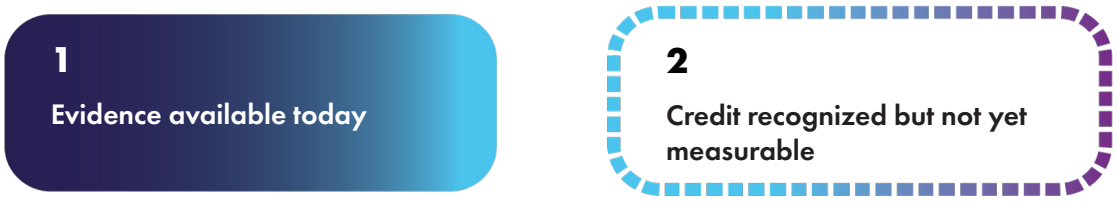
2. **It defines who is eligible to claim credit and what evidence that claim requires.** When AI does not shape either intent or the final choice, there is no new AI claim. When AI shapes intent but not the final choice, AI and the parties that informed its answer may be eligible for contributory claims. When AI determines the choice, the AI surface or agent is eligible for the decision claim, while parties whose content, data, or paid inputs inform that decision may be eligible for contributory claims. Execution alone earns no claim.

The framework does not divide conversions into percentages. **A claim is not credit; it is the right to be evaluated for credit.** The pathway determines claim eligibility, evidence determines whether that claim is substantiated, and the parties evaluating it determine what credit follows. In other words: **pathway (which sets who is eligible to claim), then evidence, then credit.** Existing attribution systems still operate alongside this framework and the blueprint defines the additional claims created by AI involvement.

Evidence falls into **two categories:**

1. **Evidence available today**
2. **Credit recognized but not yet measurable**

Two Categories of Evidence



Today, referral and citation data, seller-side purchase records, agent transaction records, incrementality tests, and shopper-reported data can substantiate parts of an AI claim. Other influence is real but cannot yet be fully substantiated because the necessary evidence is either not captured or remains inside AI platforms.

That measurement gap creates a concrete agenda for the ecosystem. AI platforms can close many of the largest gaps within their control by noting whether a recommendation was paid or organic, which sources informed it, what the AI system surfaced, recommended, or executed, and by providing consistent, durable labeling of AI-originated traffic and referrals. Brands can measure or estimate the share of their conversions that fall into Human-Owned, AI-Assisted, and AI-Executed pathways. Agencies can apply the same claim and evidence standard across platforms. Publishers can identify the contributory claims their content creates and the evidence needed to substantiate them.

The framework applies a taxonomy to how AI influences a conversion and sets out how a claim would be identified. But the larger task is to close the gap between the role AI is already playing in commerce and the industry's ability to observe, verify, and measure it. That means identifying where evidence is missing, what greater transparency is needed, and what new measurement capabilities the ecosystem must build.

IAB invites brands, agencies, AI platforms, publishers, and measurement providers to apply the framework to live categories, surface those gaps, and work together to close them.

Contributing Industry Leaders

iab. AI-Era Attribution Blueprint

Contributing Industry Leaders

The IAB AI-Era Attribution Blueprint Working Group is a cross-industry team of leaders from brands, agencies, publishers, AI platforms, and measurement providers who collaborated to create this framework. This group's goal is to give the industry a shared way to classify conversions that AI influences or completes, define the claims AI involvement creates, and identify the evidence those claims require. Additionally, it names where that evidence does not yet exist, so brands, agencies, publishers, AI platforms, and measurement providers can work from a common starting point to close those gaps.

This working group is led by Caroline Giegerich, VP, AI at IAB, with special thanks to members:



Breeze Zhao
Sr Director, Measurement Science,
Walmart



Caroline Giegerich
VP, AI,
IAB



Christoper Papes
Head of Growth Strategy and Yield Management,
WPP Media



Collin Colburn
VP, Commerce & Retail Media,
IAB



Dave Lewis
Head of Partnerships,
emberos



Floriana Nicastro
Sr Director, Measurement Product,
Criteo



Ihab Rizk
Sr Product Manager,
Microsoft Clarity



Imri Marcus
CEO & Co-founder,
Brandlight



Jing Feng
Co-founder and COO,
Bluefish



Joseph Geritano
Director of Measurement & Agentic Solutions,
WPP Media



Justin Inman
Founder & CEO,
emberos



Luke MacLean
US Total Search Lead,
WPP Media



Priya Bhasin
Principal Product Manager, AI,
Microsoft



Ross Bucholc
Associate Media Director,
Prudential Financial



Simon Poulton
EVP, Innovation & Growth,
Tinuiti



Edward (Teddie) Cowell
Global VP, GSO (Generative & Search Optimisation),
WPP Media

The Attribution Challenge

Digital attribution, in its dominant deterministic form, was built around **three observable signals:**



An ad server logged the impression, a click carried a signal forward, and the conversion landed somewhere that could identify how the visitor arrived. Credit could flow backward as a trail existed.

That said, attribution has never provided a complete or causal view of the customer journey. Signal loss, walled environments, offline behavior, and the gap between observable touchpoints and incremental impact all predate AI. AI exposes and extends the existing attribution challenges by introducing new journeys in which influence, decision, and execution may occur without conventional attribution signals.

AI is increasingly introducing a new journey: prompt, recommendation, action. Consumers are adopting these experiences because they can collapse discovery, research, comparison, and decision-making into fewer steps, often within a single interaction. A consumer asks a question, an AI system recommends a product or brand, and the consumer acts. Research and comparison that once unfolded across multiple trackable pages can now happen inside a single conversation. What follows may be a click, but it may also be a search days later, a typed URL, an in-store purchase, or a transaction an agent completes on the consumer's behalf.

In that journey, the traditional attribution trail weakens or disappears:

No required click. A consumer can leave the conversation informed, persuaded, and decided without clicking anything. Some media has always influenced without clicks: television, display, and zero-click search worked at the awareness end, and the research that followed still generally produced clicks. An AI conversation can carry the consumer through that research to a decision, with no click anywhere along the way.

Fewer referral signals. Even when a consumer clicks through, the signal identifying the AI source may not survive the handoff.

No conventional impression. An organic recommendation inside an AI answer does not generate the kind of served-media impression attribution systems were built to measure, though some emerging AI metrics now count brand or link appearances as impressions. Organic search posed a version of this problem, since a results-page appearance was never a served impression either, but search generally delivered a click that identified the source. AI answers often do not.

Walled transactions. When checkout happens inside the AI surface, critical recommendation and transaction data may remain with the platform. Platforms keeping transaction data inside their walls is an old problem: marketplaces, social checkout, and retail media networks have done it for years. What is new is what sits behind the wall: whether a product was recommended and why, so a brand can be recommended and sold without ever knowing it.

Agent autonomy. An agent acting on standing instructions can create a purchase without a conventional browsing session at all.

The practical consequence is not that AI-influenced sales disappear from reporting. They are often reported under the wrong source of influence. A sale shaped by an AI recommendation may arrive as branded search, direct traffic, or an ordinary retail media conversion, and existing dashboards assign credit to what they can observe. The conversion is visible. The influence that produced it is not. As AI absorbs activities that once occurred across multiple upper- and lower-funnel touchpoints, some of those intermediate influences may disappear from the observable attribution trail altogether.

The scale of that gap is becoming measurable. Adobe Analytics reports that traffic from AI sources to U.S. retail sites grew 393% year over year in the first quarter of 2026, and by March converted 42% better than non-AI traffic. This reflects only the AI activity measurement can currently see. Similarweb found that consumers receiving a brand recommendation from an AI assistant were 2.5 times more likely to visit the brand's site within seven days. Yet only 8.8% of those visits arrived by clicking through from the AI itself. The rest arrived wearing other labels. Most came through search in which the consumer heard about the brand in an AI answer, searched for it by name days later, and arrived from the results, where web analytics logged a typical search visit. Others arrived as direct traffic in which the consumer typed the site's address, which leaves no trail at all.

9 in 10 AI-influenced visits are credited to something else

How AI-influenced visits show up in analytics



Source: Similarweb, The Downstream Impact of AI Visibility, June 2026. U.S. desktop visits

The scale of that gap The AI recommendation created the influence, though search and direct collected the observable visits. For a brand, this has a direct downstream effect as these are the analytics that may decide where the marketing investment goes. A dashboard that credits search and direct sends the next dollar to search and direct, while the channel that actually created the demand may not get enough of the credit, thus diminishing the potential effectiveness of that investment.

Two scenarios illustrate the attribution problem.

Scenario one. A consumer asks an AI assistant for the best running shoe under $150. The assistant recommends a specific Brand X model. Three days later, the consumer searches for Brand X and buys the shoe at a retailer. Analytics credit search. The retailer sees a search conversion, and Brand X sees a sale with no attributable AI touchpoint. The AI recommendation that initiated the purchase is absent from the measurement record.

Scenario two. A consumer researches products in an AI assistant, arrives at a retailer already considering a short list, then encounters retailer search and display ads before purchasing. The retail media trail supports a claim that those ads contributed to the sale. But the AI may have determined which brands entered consideration in the first place. Did the media create demand, influence the final choice, or simply capture demand that already existed? Current attribution systems cannot reliably distinguish among those roles.

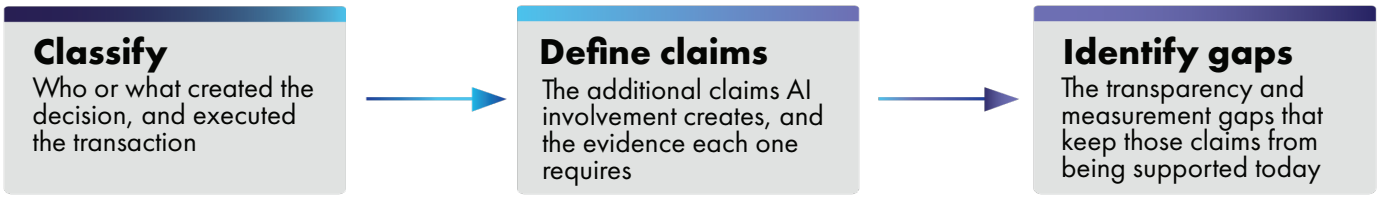
Both scenarios expose the same problem: observable touchpoints and actual influence are no longer the same thing. Before the industry can decide who receives credit, it needs to establish what role each participant legitimately played.

This blueprint addresses the additional attribution questions created by AI involvement in a conversion. It does not look to replace existing attribution systems. Instead, it provides a common model for identifying AI's role within an outcome and where today's measurement environment falls short. It does that by:

1. Classifying who or what created intent, made the decision, and executed the transaction
2. Defining the additional claims created by AI involvement and what evidence those claims require
3. Identifying the transparency and measurement gaps that prevent those claims from being substantiated today

What this blueprint does

It works alongside existing attribution. It covers only the additional questions AI involvement creates.



AI Influence Pathways



iab. AI-Era Attribution Blueprint

AI Influence Pathways

AI is changing not only how people discover products, but also who, or what, ignites interest, makes a selection, and completes a purchase. Before the industry can assign credit consistently, it needs a shared way to classify the role AI played in each conversion.

To begin this classification work, we start with 3 key questions:

Three Questions

Intent
Who or what created the intent or interest?

Decision
Who or what made the decision?

Execution
Who or what executed the transaction?



Intent rarely has a single source. Brand experience, advertising, word of mouth, prior consideration, and other influences may accumulate before an AI interaction occurs. The framework does not attempt to identify the original cause of demand. It classifies the role AI played in shaping the intent and decision associated with the conversion.

These questions are applied at the point of sale. AI may play several roles over the course of a shopping journey, but each conversion receives one pathway based on the roles played ***when the purchase occurs***. Earlier AI interactions may provide evidence for that classification, but they do not create separate or competing claims on the same conversion.

For example, a shopper may first hear about a product through an AI-generated answer, later research it independently, and eventually buy it through a retailer. Those earlier AI interactions may help establish how interest developed, but the final conversion is classified once, based on who created the intent, who made the decision, and who completed the purchase.

These pathways describe the roles actually played in a conversion, not the roles that current tools happen to capture. A conversion that AI shaped is AI-Assisted even when no observable signal survives to show it. Defining the pathways this way keeps the classification stable as what a conversion is does not change based on what a given platform, retailer, or analytics tool was able to see.

The alternative, which would be defining pathways by what can be tracked, would build today's measurement limits into the categories themselves. 'AI-Assisted' would quietly come to mean 'AI-Assisted with a surviving referral signal,' and every conversion without one would default to search or direct. The taxonomy would undercount AI's role by design, and the categories would change meaning every time tracking changed. Naming the pathway by what happened keeps the categories fixed while measurement catches up to them.

Defining pathways this way does not require anyone to know what happened in every conversion. Classification and evidence are separate steps. The pathways therefore establish what measurement should ultimately be able to distinguish, even where the evidence to do so does not yet exist. For a single conversion, a party should assert a pathway only where evidence supports it. Where it does not, no AI claim is asserted. Across many conversions, pathway mix can be estimated through shopper surveys, lift tests, and modeling.

Additional Context for Each Pathway

Two additional distinctions apply within every pathway.

The first is **transaction location:**

1. **On-platform:** The transaction completes within the environment where the AI operates. This includes a standalone AI surface like ChatGPT or Microsoft's Copilot, a retailer's own assistant such as Amazon's Rufus or Walmart's Sparky, and a third-party AI embedded in another party's website or app.
2. **Off-platform:** The shopper or agent leaves that environment and completes the transaction somewhere else, such as a different retailer's website, app, marketplace, or physical store.

Transaction location is assessed relative to the AI system whose role is being classified. For example, a shopper who leaves one AI surface and buys at a retailer has transacted off-platform relative to that AI, even if the retailer operates an AI assistant of its own.

The second is **commercial relationship:**

1. **Paid:** The product's appearance is connected to a paid placement or sponsored recommendation.
2. **Organic:** The product appears without a paid placement connected to that recommendation.

A commercial relationship that provides access to grounding content or data does not make a recommendation paid unless that relationship influences the product's appearance or selection.

These distinctions do not create additional pathways. Instead, they provide important context for understanding how the conversion occurred, what information may be available, and which parties may seek credit.

Defining AI by Its Role

This framework defines AI based on the role it played, not the type of product or interface involved.

An AI system is one that generates a recommendation or takes an action that replaces part of the normal search-and-click journey. This may happen through a chatbot, an AI search experience, a conversational assistant, an autonomous shopping agent, or another AI-enabled interface.

The definition therefore depends on what the system did in the conversion, not what the system is called. The same product may support different pathways in different transactions. A chatbot may answer a factual question in one journey, recommend a specific product in another, and complete a purchase in a third.

Using a behavior-based definition avoids treating every interaction with an AI-enabled product as an AI-influenced conversion. The relevant question is whether the system shaped intent, shaped the decision, or executed the transaction, whether in response to a current user prompt or proactively based on prior context, preferences, behavior, or inferred needs.

This definition excludes conventional retailer recommendation features, such as product carousels, 'customers also bought' modules, and personalized shelves within a retailer's own site or app. These operate inside the traditional search-and-browse journey rather than replacing it, and existing retail media and site attribution already cover them. This exclusion is behavioral, not locational. An AI experience within a retailer's own site or app that understands a shopper's need, narrows options, or recommends a specific product is classified the same way as an equivalent experience anywhere else.

Human-Owned



A conversion is Human-Owned when the human created the intent, made the decision, and executed the purchase.

AI may still appear somewhere in the broader journey, but it did not take over any of the three roles used to classify the conversion at the point of sale.

Human-Owned does not mean the journey was free of technology. A consumer may have used search engines, retailer tools, comparison sites, media, or AI-enabled services while gathering information. The pathway remains Human-Owned when the human retains ownership of intent, decision, and execution whether or not evidence exists to demonstrate it.

For example, a shopper decides to replace their running shoes with the same model they already own. They ask an AI assistant whether the newest version runs true to size, compare prices across two retailers, and buy from one of them. Technology appeared throughout the journey, including AI. But the shopper created the intent, made the decision, and executed the purchase. The AI answered a question the shopper brought to it, but it did not shape what they wanted or what they chose.

AI-Assisted

A conversion is AI-Assisted when the human executes the purchase, but AI shapes the shopper's intent, decision, or both by creating awareness, framing the available options, narrowing the choice, or recommending a specific product.



When a shopper follows a link from an AI-generated answer and buys at that destination, the conversion remains AI-Assisted. A clickthrough does not erase AI's role in surfacing the product or source.

Within AI-Assisted, the framework identifies an important distinction between

1. AI-Influenced conversions
2. AI-Recommended conversions

In an AI-Influenced conversion, AI contributes to intent or interest, but the human independently evaluates the options and makes the decision. AI may introduce a category, product, brand, seller, or need, including when the AI raises the need unprompted (e.g., an agent noticing a product is running low and suggesting a repurchase), without determining the final choice. The resulting credit logic is similar to upper-funnel media: AI made a contribution, but its effect may need to be established through lift or other evidence rather than through a direct claim on the full conversion.

In an AI-Recommended conversion, AI materially shapes the final choice, whether or not it also shaped the shopper's original intent. The system narrows the available options to a specific product, brand, seller, or course of action, and the human accepts that recommendation before completing the purchase.

Because AI influenced the final choice, it becomes important to understand why that product was recommended. That includes whether the recommendation was paid or organic, which sources informed it, and how the system selected it. *The AI Credit Claim Framework* addresses how those factors may affect credit.

Transaction Location Within AI-Assisted

When an AI-Assisted transaction occurs off-platform, the shopper leaves the AI experience and completes the purchase elsewhere. In many cases, the AI interaction leaves little or no usable signal at the destination. The resulting sale may appear in reporting as direct traffic, search, retail media, an affiliate visit, or another conventional source even when AI materially shaped the journey.

When a referral signal is available, emerging measurement tools may be able to connect an AI interaction to on-site activity and conversions. When no signal survives, like when a shopper later searches for the brand or purchases elsewhere, AI's role may need to be inferred through other evidence. The absence of a referral signal does not mean AI played no part in the conversion.

When an AI-Assisted transaction occurs on-platform, the purchase takes place inside the AI surface. This creates the potential for a more direct and trackable handoff between recommendation and transaction. However, the platform controls much of the available information. Credit therefore depends in part on what the platform chooses to disclose about why the product appeared, whether money influenced the recommendation, and which sources or signals contributed to it.

AI-Executed

A conversion is AI-Executed when a human has given an AI system permission to buy on their behalf ahead of time, and the system completes the purchase in the moment. AI-Executed also includes agent-to-agent transactions, such as a buyer's agent transacting directly with a seller's agent. Where one agent delegates part of the journey to another, the same role-based logic applies across the agent chain: claims follow the role each agent actually played, not simply the last agent to touch the transaction.



An agent that suggests a purchase and waits for the shopper's approval is not AI-Executed, as the in-the-moment approval makes it AI-Assisted. AI-Executed requires that permission came in advance rather than at the moment of purchase.

AI-Executed conversions divide into two types based on how much of the choice the human made in advance.

1. **Specific instruction:** the human selected the product or clearly defined the purchase in advance. AI's role is to monitor for a human-defined condition and execute the purchase when that condition is met. Because the human retains control of the intent and decision, credit follows the same logic as Human-Owned, even though AI performs the final action. For example, a shopper may tell an agent to reorder the same dog food whenever the supply falls below a certain level. The agent may detect that condition and place the order without a new prompt, but the human already determined both what to buy and under what conditions to buy it.
2. **Open instruction:** the human defines a broader objective but leaves the product choice to the AI. The system selects the product and completes the transaction. Because AI makes the decision as well as executing the purchase, credit follows the same logic as AI-Recommended, with AI treated as the decision-maker.

The distinction between specific and open instructions prevents all agent-completed purchases from being treated the same. Execution alone does not determine credit. The framework must also consider whether the human or the AI selected what was purchased.

Three pathways, defined by who played each role

An agent that suggests a purchase and waits for approval is AI-Assisted. AI-Executed requires permission given in advance.

Pathway	Intent	Decision	Execution
Human-Owned	Human	Human	Human
AI-Assisted **AI-Influenced** AI shapes intent only	AI	Human	Human
AI-Assisted **AI-Recommended** AI shapes the final choice	Human or AI	AI	Human
AI-Executed **Specific instruction** Human chose what to buy in advance	Human	Human	AI
AI-Executed **Open instruction** AI chooses and buys	Human	AI	AI

Measurement Gap

The hardest distinction in this framework to observe is the line between AI shaping intent only and AI shaping the decision. The two roles lead to different credit treatment, but they may look identical in the data. A retailer may know that a shopper arrived after an AI interaction without knowing whether AI introduced the product, recommended that exact product, or confirmed a choice the shopper had already made.

Resolving that distinction may require information beyond standard transaction data, such as the shopper describing how the decision was made or the platform disclosing what it recommended and why. The framework preserves the distinction as the pathways describe what happened, not only what current measurement systems can observe. The categories should not blur simply because today's data cannot always separate them. The sections that follow define the claims those roles create, the evidence available to support them today, and the gaps the ecosystem would need to close to make those claims consistently measurable.

AI Credit Claim Framework



iab. AI-Era Attribution Blueprint

AI Credit Claim Framework

The previous section classified the role AI played in a conversion. This section addresses what follows from that classification: who can claim credit for the conversion, and what evidence each claim requires. The pathway classification can be applied today. The claim framework that follows defines eligibility rules and evidence standards against which current measurement capabilities can be assessed. Some claims can be substantiated now and others identify evidence the ecosystem does not yet provide.

This framework does not divide a conversion into percentages. There is no defensible basis for assigning fixed shares of credit across parties in journeys that leave partial or no trails, and inventing one would undermine the classification work the framework rests on. Instead, this framework defines the claim each party can make under each pathway, and the evidence needed to support that claim.

The framework establishes whether a party has a legitimate claim of influence within a conversion. It does not establish the incremental economic value of that influence, which requires experimentation, modeling, or other causal measurement.

Claims should also be distinguished by how independently their supporting evidence can be verified. Where the party making a claim also controls the recommendation, transaction, or underlying evidence, that claim should be distinguished from one supported by independently verifiable evidence. As AI platforms take on more roles across the journey, independent verification becomes more important.

The industry has taken a similar approach before. Viewability standards did not measure attention; they defined what counts as a viewable impression and what evidence substantiates it. AI attribution presents a harder implementation challenge because much of the relevant evidence may sit inside AI platforms and depend on platform disclosure. The same principle still applies: define the claim and the evidence required to support it, even when that evidence is not yet consistently available.

Three Types of Claim

Every credit claim under this framework is one of three types:



1. **A decision claim** applies when an AI surface or agent determines the final purchase choice. It identifies AI's role in making or materially shaping that choice; it does not mean AI caused the sale, had the greatest overall influence, or deserves the most credit. Where the human made the choice, no decision claim arises under this framework and existing attribution applies unchanged. Whether AI's role created incremental economic value must be established separately through causal measurement, including incrementality testing where feasible.
    - For example, a shopper may already intend to buy sunscreen for summer and ask an AI assistant which brand to choose. If the AI recommends a specific product and the shopper accepts that recommendation, AI may hold the decision claim even though the underlying demand for sunscreen came from elsewhere.
2. **A contributory claim** applies when the AI answer helped move the shopper toward the purchase but did not determine the final choice. Eligibility for a contributory claim attaches to the conversion, but the claim cannot be proven one sale at a time. It is proven by comparing two groups of shoppers: one group that saw the AI's

answer and one group that did not. If the shoppers who saw the answer buy more often, the difference shows AI contributed, and how much. That difference only shows up across many shoppers. It cannot be seen in any single sale.

- This comparison is feasible today most cleanly with paid AI placements, where exposure can be withheld from some shoppers and tested through established lift methodologies. Organic AI answers require a different approach because responses are dynamic, personalized, and generally cannot be withheld for a controlled group. Evidence from surveys, before-and-after trends, and modeling can help estimate organic AI's contribution, but today it is generally directional rather than causal. Proxy measures such as AI presence, rank, citation frequency, share of prompt, and related visibility metrics can also serve as inputs to trend analysis and MMM, but they are directional indicators rather than standalone proof of a claim.

3. **No claim** means AI's involvement did not rise to either a decision or contributory role. A party with no claim under this framework may still hold credit under existing attribution practices, such as search, retail media, or affiliate attribution. This framework does not displace those systems. It defines the additional claims that AI involvement creates.

The framework separates three steps that are often collapsed. The pathway determines which claims each party is eligible to make. Evidence determines whether an eligible claim is substantiated. The parties evaluating the claim, such as a brand, an agency, or an auditor, determine what credit follows. Pathway (which sets who is eligible to claim), then evidence, then credit.

Where parties disagree about which pathway applies, the party asserting the claim must provide evidence supporting that classification. A decision claim should not be substantiated unless the evidence supports that AI materially shaped the final choice. If the available evidence establishes AI influence but cannot distinguish between AI-Influenced and AI-Recommended, it supports only a contributory claim.

Claims by Pathway

More than one party can be eligible to claim. Only one can hold the decision claim. Execution alone creates no claim.

Pathway	Who is eligible to claim	Claim type
Human-Owned	No AI claim. Credit follows existing attribution practice.	None
AI-Assisted **AI-Influenced**	The AI surface, and the parties whose content or data informed its answer.	Contributory
AI-Assisted **AI-Recommended**	Decision claim: the AI surface. Contributory: parties whose content, data, or paid inputs informed the recommendation.	Decision, Contributory
AI-Executed **Specific instruction**	No AI claim on the decision. The execution is a service, not an influence.	None
AI-Executed **Open instruction**	Decision claim: the AI agent. Contributory: parties whose content, data, or paid inputs informed its selection.	Decision, Contributory

Three notes on reading the table:

First, the claim types follow directly from the three questions in AI Influence Pathways. When AI determines the final choice in the purchase, the AI surface or agent is eligible for the decision claim. Parties that shaped intent or informed the decision may be eligible for contributory claims. An agent that only completes the purchase, without shaping what the person wanted or choosing the product, is not eligible for a claim.

Second, more than one party can be eligible to claim on the same conversion, but only one party can hold the decision claim. This follows from the one-pathway rule: each conversion is classified once, and the classification names one decision-maker.

Third, a claim is not credit. A claim is the right to be evaluated for credit. Being eligible to make a claim does not prove it. And proving a claim does not set what it is worth, because this framework does not assign amounts. Whether a substantiated claim converts into credit, in a report, a budget decision, or a compensation conversation, is decided by the parties evaluating it.

Two Categories of Evidence

Evidence for a credit claim falls into two categories. The distinction between them is the central, practical finding of this framework.

Evidence Available Today

Relies on Signals and tests available now
Uses referral data, purchase records, and holdout tests

Uses existing tools and records
Leverages citation clicks, affiliate links, and surveys

Supports claims when combined
Offers deterministic or directional insights

Not Yet Measurable

Depends on future disclosure or improved tracking
Requires bridging gaps between sourcing and conversions

Includes model training and the link between sources and sales
Includes how sources shape answers and unverified paid influence

Shows up in trends, not individual sales
Includes sales that leave no trace

Evidence available today. Some claims can be supported now, with existing tools. No single item below proves an AI claim on its own. Each proves a different part of the conversion, and claims are substantiated by combining them:

- **Referral and citation data.** Some AI interactions leave a conventional trail. A referral signal is the information a link or browser passes to the destination site showing where the visitor came from, such as a UTM tag on the link or a referrer showing an AI domain. A citation click is a shopper clicking a source link inside an AI answer. An affiliate link is a tracked link that identifies who sent the shopper, and often pays a commission on the sale. In each case, the AI path is demonstrably in the trail, and the evidence is deterministic.
- **Seller-side purchase records.** Whoever operates the checkout, whether a retailer, a marketplace, or a brand

selling directly, holds its own record of every sale. That record proves the sale, not how the shopper arrived. It anchors the other evidence: group comparisons and platform claims are both checked against it.

- **Agent transaction records.** When an agent buys something, agentic checkout standards create a record of the transaction: the agent identifies itself as an agent, the merchant confirms the order, and the payment carries the shopper's advance permission. This record proves an agent completed the purchase. But under this framework, completing the purchase earns no credit on its own, so the record cannot answer the two questions that decide credit: what did the human tell the agent to do, and who picked the product? The instruction answers the first and separates specific from open. A record of the agent's selection answers the second, and it is what the agent's decision claim rests on under an open instruction. Intent, decision, and execution each need evidence of their own. Emerging agentic commerce protocols may increasingly surface the information needed to substantiate those distinctions as they evolve.
- **Incrementality and holdout tests.** These are the group comparisons described under contributory claims: shoppers who saw the AI's answer against similar shoppers who did not. They are the standard evidence for AI-Influenced claims because AI-Influenced plays the same role as awareness advertising, and awareness advertising has always been proven this way. A TV campaign is never proven to have caused one sale. It is proven by markets that saw it outselling markets that did not.
- **Shopper-reported data.** Post-purchase surveys and panels can capture whether an AI interaction shaped a purchase. This evidence is directional rather than causal, but it is available now and it covers journeys that leave no signal.

Credit recognized but not yet measurable. Some influence is real but cannot be substantiated today, either because the evidence exists and is not disclosed or because it is not captured at all. This framework shows platforms exactly what they would need to disclose for these claims to become substantiated.

- The clearest case is publisher content behind an AI recommendation. A shopper asks an AI system for the best option in a category. The answer draws on a publisher's review. The shopper never clicks the citation and buys at a retailer, where the sale reports as search or direct. The review shaped the sale, and the publisher is eligible for a contributory claim. The evidence differs depending on how publisher content informed the answer. Content used in model training is generally not traceable from outside the model. Grounding and citation are more observable: publishers may be able to see when their content is cited or reflected in an AI response, although citation alone does not establish how much that content influenced the answer. What remains largely missing is the connection between those sourcing signals and downstream conversions. Until that link exists, the contributory claim cannot be fully substantiated.
- Other claims in this category today include:

1. **Whether AI shaped the decision or only the intent.** Without platform disclosure or shopper-reported data, that distinction often cannot be observed.
2. **Whether a paid relationship influenced a recommendation.** Platforms may disclose advertising, but the basis for a recommendation inside an AI answer is not independently verifiable today. Commercial influence may also exist within cited or grounding content, such as sponsored listings or affiliate relationships that the AI system cannot distinguish from editorial content.
3. **Conversions that leave no signal.** A shopper may read an AI answer, then later search for the brand or go directly to its site. That influence can be estimated in aggregate, but not proven for an individual conversion.

The two categories are not fixed. Evidence moves from the second category to the first when platforms disclose more, when checkout protocols spread, and when measurement methods mature. Some gaps close through disclosure alone,

because the evidence already exists inside AI platforms. Others will require new measurement capabilities that connect AI interactions to downstream outcomes, capabilities the ecosystem has not yet built.

Evidence should also be understood along two additional dimensions:

1. **What it establishes**
2. **How independently it can be verified**

Observed evidence can establish that an AI interaction, recommendation, or transaction occurred, whereas incremental evidence can establish whether that activity caused additional outcomes. Evidence may be platform-reported and/or independently verifiable. These dimensions describe the strength and nature of the evidence available within them.

How Transaction Location & Commercial Relationship Affect Claims

Transaction location does not change the claim. It changes the evidence available to support it. An AI-Recommended conversion carries the same decision claim whether the purchase happened inside the AI surface or at a retailer. On-platform, the handoff is trackable but the platform controls what is reported. Off-platform, the claim usually depends on referral data where it survives, and on aggregate evidence where it does not.

Commercial relationship does not change the claim either. It changes what must be disclosed alongside it. A paid recommendation and an organic recommendation carry the same claim structure under the same pathway. But a party asserting credit for a paid recommendation should disclose the paid relationship, because the parties evaluating the claim, a brand, an agency, an auditor, cannot weigh it honestly without knowing whether money shaped the recommendation.

Closing the Measurement Gap



iab.

AI-Era Attribution Blueprint

Closing the Measurement Gap



Every credit claim can be evaluated through three questions:

1. **Which pathway?**
2. **What claim type?**
3. **What evidence category?**

Those questions create a common structure for evaluating AI's role in a conversion. From there, each part of the ecosystem has a different role in making that role more observable, verifiable, and measurable.

Brands should measure or estimate the share of their conversions that fall into Human-Owned, AI-Assisted, and AI-Executed pathways. That shows whether the brand is increasingly influencing people, AI systems, or both, and where investment may need to shift across creative, content, structured data, and publisher relationships. When a platform reports AI-driven conversions, brands should ask what pathway those conversions came from, what credit is being claimed, and what evidence supports it.

AI Platforms hold much of the evidence needed to substantiate those claims. Four disclosures would address significant gaps within their control:

1. Whether a product surfaced through paid placement or organic recommendation
2. What sources informed the recommendation
3. What the AI system surfaced, recommended, or executed, and what influence is observed versus inferred
4. Consistent, durable labeling of AI-originated traffic and referrals

Platforms reporting AI-driven conversions should distinguish claims supported by observable evidence from claims that rest on their own reporting.

Agencies can use the same framework across platforms: pathway, claim type, evidence category. The key distinction is between platform-reported credit and credit supported by independently verifiable evidence such as lift.

Publishers can identify the contributory claims they are eligible to make and the evidence needed to substantiate them. Citation reporting is a start. Connecting the sources that informed an AI recommendation to downstream outcomes is what would make those claims measurable.

This blueprint does not assume that today's measurement environment can substantiate every AI-related claim it identifies. Its purpose is to make the gap explicit: to establish how AI can influence a conversion, what evidence would be needed to account for that influence within attribution, and where the ecosystem cannot yet provide it.

The much larger task is to close that gap. Attribution has never provided a complete view of how influence becomes an outcome, and AI adds more challenges. The goal is to make more of AI's role observable, verifiable, and measurable than it is today. Doing so will require greater transparency, stronger standards, and new measurement capabilities across the ecosystem.

Call to action:

IAB invites brands, agencies, AI platforms, publishers, and measurement providers to apply the framework to live categories, identify where evidence breaks down, and work together to build the transparency and measurement capabilities needed to make AI's role in commercial outcomes more observable, verifiable, and measurable.

Help close the gap

IAB invites the industry to test this framework in practice

