

The state of ecommerce in 2026

What changed, what didn't, and what
retailers should do next



29.4% OF UK RETAIL SALES WERE MADE ONLINE IN JUNE 2026 [C01]

PUBLISHED

14 August 2026

UltraLabs, Buckingham

WRITTEN FOR

Ecommerce owners and operators

Running an established online store

EVIDENCE

48 sourced claims

Each with a named publisher and date

Introduction

Ecommerce did not get smaller this year. It got harder to be found in, and the reason is not the one most articles gave you.

Demand is strong. Online sales were 29.4% of all British retail in June 2026, the highest share since April 2021 ^[c01, c02]. In the United States, ecommerce grew 9.8% year on year while retail as a whole grew 3.9% ^[c06]. Nothing in the national statistics supports the idea that shoppers are leaving.

Discovery changed. Search engines now answer questions on the results page, and more of the web is read by machines than browsed by people ^[c08, c11]. A shopping assistant reading your catalogue will not add a product to the basket to find out what delivery costs.

This report separates the two. Part one establishes what the demand data actually shows. Parts two and three establish what changed in how products are found, and what regulators now require of the same information. Part four is our own framework, and it is labelled as ours throughout.

You will leave with an audit. Not a trend list. A twenty-cell check you can run on one product page in an afternoon, covering the five things a buyer has to settle before they can order.

What you will be able to audit by the end

01

The right product

02

The total price

03

Stock

04

Delivery

05

Returns

Every figure in this report carries a reference like this ^[c01], resolving to the source list on page 36. Anything labelled UltraLabs analysis is our reading of the evidence, not a finding from it.

HOW TO READ IT

- **Parts one to three are evidence.** Every number names its publisher and its measurement period, and where the evidence is older than 2026 the sentence says so.
- **Part four is ours.** The five questions, the four places and the scoring method are UltraLabs analysis, labelled throughout, and you are meant to argue with them.
- **Parts five and six are the work and the workings,** including every claim we threw out and why.

WHAT THIS REPORT IS NOT

- Not a trend list. Nothing here is included because it was discussed a lot this year.
- Not a forecast. We make none, and where somebody else has, it is attributed to them.
- Not a pitch. The audit at the end works whether or not you ever speak to us.

Key findings

Six measurements that between them describe the year. Each is published by a named organisation on a stated date, and each is checkable from the source list on page 36.

DEMAND

29.4%

of UK retail sales were made online in June 2026, the highest share since April 2021

ONS · C01 C02

DEMAND

9.8%

US ecommerce growth year on year, against 3.9% for retail as a whole

US CENSUS BUREAU · C06

DISCOVERY

8%

of Google visits produced a click on a result when an AI summary appeared, against 15% without one

PEW RESEARCH CENTER, 2025 DATA · C08

DISCOVERY

54%

better conversion from visits referred by AI tools than from other traffic sources

ADOBE ANALYTICS · C16

PRESSURE

£4.7m

in fines in the UK regulator's first year of direct consumer enforcement, much of it for revealing costs late

CMA · C29

PRESSURE

12%

rise in the average price Meta charged for an advert, year on year

META PLATFORMS · C22

What these findings mean for an online retailer

Four conclusions follow from those six numbers. None of them is a prediction, and each one names the findings it rests on.

01 Your traffic problem, if you have one, is not a demand problem

Online retail is near a five-year high in Britain and growing faster than the high street in America ^[C01, C06]. Shopify merchants processed \$115.6bn in a single quarter, up 31.6% ^[C25]. If your sessions are down, the market did not do it.

02 Fewer searches end in a click, but the clicks that survive are worth more

Click-through roughly halved where an AI summary appeared ^[C08], while visits arriving from AI tools converted 54% better than other sources ^[C16]. Both are true at once, and together they describe a filter rather than a decline.

03 That filter reads the same information the law now wants stated early

Total price, delivery, stock and returns. A regulator fining late cost disclosure and a shopping assistant skipping your product are asking for the same thing for different reasons ^[C29].

04 Waiting is getting more expensive, and the fix is unglamorous

The average price of a Meta advert rose 12% year on year ^[C22]. What this points at is not a replatform: it is product data, total price, stock, delivery dates and returns terms, stated in the same words everywhere they appear. Mostly content and configuration, which is why it tends not to get done.

Nothing in the demand data explains a fall in your traffic. What changed is who gets the click.

ULTRALABS ANALYSIS

Contents

01	Ecommerce demand remains strong	06
	Online sales are near a five-year high in the UK	07
	US online sales are growing faster than retail overall	08
02	How people and machines find products has changed	10
	Most web traffic is now machines, not people	11
	AI summaries reduce clicks on search results	12
	Shoppers who arrive from AI tools convert better, not worse	13
	On Shopify, AI-driven orders grew fast and then grew much less fast	14
	What the AI shopping platforms launched, and what Google says to do	15
	What the agent shopping standard actually reads	16
03	The law and your ad costs now push the same way	17
	UK regulators are fining businesses that reveal costs late	18
	Import duty and accessibility are now legal requirements	19
	Advertising is getting more expensive	20
	How much of your demand depends on one channel?	21
	Why shoppers abandon carts	22
04	The product page audit	23
	Five questions every product page should answer	24
	Where your store answers, and how answers go wrong	25
	The twenty checks in one grid	26
	How to score your own product pages	27
	Question one: is this exactly the product I mean?	28
	Question two: what will I pay in total?	29
	Question three: do you have it in stock?	30
	Question four: when will it arrive?	31
	Question five: what happens if it is wrong?	32
05	What to do about it	33
	Fix these in this order	33
	What this audit will not fix	34
06	Method and sources	35
	How we researched this report	35
	Sources	36
	Send us one product page	38

PART ONE

Ecommerce demand remains strong

Before anything about AI, the demand picture. It is less dramatic than the commentary and more useful, because it rules out the explanation most people reach for first.

29.4%

of British retail sales were made online in June 2026 [\[C01\]](#)

16.9%

of American retail sales were ecommerce in the first quarter [\[C05\]](#)

Four years


of a broadly flat British online share since the 2021 peak [\[C04\]](#)

Online sales are near a five-year high in the UK

27.4%

was the online share for the whole of 2025, still below the 2021 peak of 30.7%. The monthly figure reached **29.4% in June 2026**, the highest since April 2021 ^[C01, C02, C04].

FIGURE 1 Internet sales as a percentage of total retail sales, Great Britain, annual



Source. ONS series J4MC, Retail Sales Index (DRSI), released 24 July 2026 ^[C04]. Annual values. The monthly figures in this series retain retail seasonality, which is why they differ from the seasonally adjusted proportion the ONS quotes in its monthly bulletin ^[C01].

What this means. Online share took eleven years to climb from 3.4% to 19.2%, then gained the same ground again in eighteen months. It peaked in 2021 and fell back. Four years later it has still not returned to that peak. Retail volumes overall rose 4.2% on the year ^[C03], so this is a healthy channel growing slowly, not a boom and not a collapse.

US online sales are growing faster than retail overall

16.9%

of all American retail sales were ecommerce in the first quarter of 2026, on **\$326.7bn** of adjusted sales. Ecommerce grew 9.8% year on year while retail as a whole grew 3.9% ^[C05, C06].

FIGURE 2 Ecommerce as a percentage of total US retail sales, seasonally adjusted



Source. US Census Bureau, Quarterly Retail E-Commerce Sales, first quarter 2026 (CB26-81), released 18 May 2026, Table 1 ^[C07]. Estimated from a stratified random sample of approximately 10,800 retail firms; employer firms only since the April 2025 benchmark. Q1 2026 is preliminary.

What this means. Five consecutive quarters of share gain, 16.0% to 16.9% ^[C07]. No inflection and no collapse. The Census publishes error bars most reporting drops: that 9.8% carries a 90% confidence interval of plus or minus 1.8 points ^[C06], so the honest statement is that ecommerce grew two to three times faster than retail overall, not exactly 2.5 times.

TWO THINGS WORTH SAYING PLAINLY

The British and American figures are not comparable as levels. Different geographies, different industrial classifications, and the American series excludes businesses with no paid employees. Read each against itself.

The platform data agrees. Shopify merchants processed \$115.6bn of volume in the second quarter of 2026, up 31.6%, on platform revenue up 34% ^[C25, C26].

ULTRALABS ANALYSIS

AI search is changing who gets the click. It is not reducing ecommerce demand.

Drawn from the click behaviour measured by Pew ^[C08] set against the demand measured by the Office for National Statistics ^[C01] and the US Census Bureau ^[C06]. Part two is about the first half of that sentence.

PART TWO

How people and machines find products has changed

A store is now read, at scale, by software that will never add a product to the basket to find out what delivery costs. What it cannot read, it cannot recommend.

Over half

of internet traffic is now non-human ^[C11]

52%

of crawler requests are for AI training, up from 22% a year earlier ^[C12]


8%

of visits produced a click where an AI summary appeared, against 15% ^[C08]

Most web traffic is now machines, not people


Cloudflare reported in July 2026 that more than half of internet traffic is non-human, and that 52% of crawler requests are made for AI training, up from 22% in spring 2025 ^[C11, C12]. Its network is very large and it is not a census of the web. The direction is not in doubt.

FIGURE 3 Share of AI crawler requests made for model training



Source. Cloudflare, "Content Independence Day, one year on", 1 July 2026 ^[C12]. Crawl purpose is derived from declared user agents and Cloudflare's own bot classification. Mixed-use crawlers, which blend search, agent use and training, are over a third of activity, so the categories are not clean.

FIGURE 4 Pages crawled for every page referral sent back




Source. Cloudflare, "A deeper look at AI crawlers", 28 August 2025, measuring the week of 1 to 7 August 2025 ^[C14]. **This is 2025 evidence, shown as an order of magnitude rather than a current retail metric.** No shopping or retail cut was published in that post.

What this means. Reading and visiting have come apart. In that week one crawler requested roughly fifty thousand pages for every visitor it sent. Google still accounts for around 88% of referral traffic ^[C13], so this is not yet where your customers are. It is where the descriptions of your products are being formed.

AI summaries reduce clicks on search results

The strongest available measurement is a Pew Research Center panel study. It tracked the real browsing of 900 American adults across 68,879 Google searches in March 2025, 12,593 of which produced an AI summary ^[C08]. **It is 2025 data.** We use it because the 2026 alternatives are agency studies of their own client sets with undisclosed sampling, and their headline figures disagree with each other by a factor of two.

FIGURE 5 What happened on a Google results page, with and without an AI summary



Source. Pew Research Center, 22 July 2025 ^[C08, C10]. Browsing data collected from Pew's KnowledgePanel in March 2025; search results collected 7 to 17 April 2025. Percentages are of visits to a Google results page.

What this means. Where a summary appeared, clicks on a normal result fell from 15% of visits to 8%, and the share of visits that ended there rose from 16% to 26% ^[C08, C10]. Links inside the summary itself were clicked on 1% of visits ^[C09].


An answer shown above the results is not a new referral channel. It is a decision made without you.

Shoppers who arrive from AI tools convert better, not worse

138%

year-on-year growth in AI-referred traffic to US retail sites to May 2026, and **1,324%** since measurement began in October 2024 ^[c15].

FIGURE 6 How AI-referred visits compared with everything else, May 2026



Source. Adobe Analytics, May 2026, based on more than a trillion visits to US retail sites, as reported by Digital Commerce 360 on 17 June 2026 ^[c15, c16, c17]. Adobe's own pages could not be reached from our build environment, so this is trade reporting of a primary dataset rather than the primary itself, and it is classified accordingly in the source list.

What this means. A year earlier this traffic converted at roughly half the rate of everything else. In May 2026 it converted 54% better, with 53% more time on site and 23% more pages per visit ^[c16, c17]. The visits are not just more numerous. They are longer, deeper and worth more.

ULTRALABS ANALYSIS


Nobody has published a causal test, so treat the mechanism as a reading rather than a finding. Ours is that this is not a better class of shopper. It is a shopper whose questions were largely settled before the click, by something reading published facts. That makes the channel a filter, and it means an unanswerable store loses sales it never sees: no session, no bounce, no line in the analytics.

On Shopify, AI-driven orders grew fast and then grew much less fast

3×

the year-earlier level for both AI-driven traffic and AI-driven orders to Shopify stores in the second quarter of 2026. A quarter earlier the same two measures were **8× and nearly 13×** ^[C35, C36].

FIGURE 7 AI-driven traffic and orders to Shopify stores, as a multiple of the year before



Source. Shopify Q1 2026 earnings call, 5 May 2026 ^[C35]; Q2 2026 earnings call, reported 5 August 2026 ^[C36]. Both are company operating metrics. Shopify has not published how AI-driven traffic is identified, the figures are not independently audited, and the Q1 orders multiple was stated as "nearly" 13×.

What this means. Each multiple is measured against the same quarter a year earlier, and a year earlier the base was small, so a fall from 13× to 3× is arithmetic as much as it is demand. Shopify has not published the size of the channel in orders or in currency, so the conclusion available here is narrow and still worth having: a platform that cleared \$115.6bn in a quarter ^[C25] now reports AI-driven orders as a line worth reporting, and the rate at which that line multiplies is coming down.

WHAT THE SAME CALL SAID ABOUT ORDINARY SEARCH

Traditional search sessions were up 1.3× over two years and still about a third of all storefront sessions, and Shopify described AI as a complement to search rather than a substitute for it ^[C36]. Set that beside part one and the shape is two routes into one market, one growing from a large base and one from a small one. Neither reading requires the other to be failing.

What the AI shopping platforms launched, and what Google says to do

29 Sep 2025

■ Instant Checkout and the Agentic Commerce Protocol

Stripe and OpenAI launch in-chat purchasing, starting with American Etsy sellers, with over a million Shopify merchants named as next ^[C18].

8 Jan 2026

■ Copilot Checkout

Microsoft launches in-conversation purchasing in the United States with PayPal, Shopify and Stripe as payment partners ^[C20].

11 Jan 2026

■ Google and Shopify publish an open standard

The Universal Commerce Protocol is announced at NRF, co-developed with Etsy, Wayfair, Target and Walmart and endorsed by more than twenty other companies ^[C39].

6 Mar 2026

□ OpenAI moves checkout into merchant apps

Instant Checkout is deprioritised as a consumer feature. About a dozen Shopify merchants had gone live with it before it was scaled back ^[C19, C43].

24 Apr 2026

□ Amazon, Meta, Microsoft, Salesforce and Stripe take seats on the standard

The protocol's technical council expands to sixteen seats. The companies competing over discovery are governing the same specification ^[C41].

10 Jul 2026

■ Google states the rules for its own AI surfaces

Structured data is not required for generative AI features, there is no special markup for them, and llms.txt files are ignored ^[C21].

THE CORRECTION WORTH PRINTING

A large amount of 2026 advice sells markup as the route into AI answers. Google's own documentation says the opposite: structured data is not required for its generative AI features, there is no special schema for them, and content rewritten for AI systems does not help ^[C21]. What the same page does recommend is publicly accessible, crawlable content and a Merchant Center product feed.

What the agent shopping standard actually reads

Through 2025 this was a set of separate platform experiments. Since January 2026 there has been a published standard with a governing council and a specification anyone can read [\[C39, C41\]](#). That matters less for what it adds than for what it confirms: the inputs are the ordinary facts about a product, held somewhere a machine can fetch them.

50bn+

product listings in Google's Shopping Graph, including inventory, prices and reviews [\[C40\]](#)

GOOGLE · JANUARY 2026

2bn

of those listings refreshed every hour, which is the freshness the rest of the market is measured against [\[C40\]](#)

GOOGLE · JANUARY 2026

1bn+

products Shopify says it has structured with clean attributes and real-time pricing [\[C37\]](#)

SHOPIFY · MAY 2026

WHAT THE SPECIFICATION COVERS

- ❑ Published on GitHub under the Apache 2.0 licence, so the working parts can be read rather than inferred from a press release [\[C42\]](#).
- ❑ Four capabilities in this release: checkout, identity linking, order events and payment token exchange [\[C42\]](#).
- ❑ **Discovery is not one of them.** Being found still runs on crawlable pages and product feeds, which is what Google's own guidance also says [\[C21, C42\]](#).

YOUR PRODUCT DATA IS PROBABLY ALREADY PUBLISHED

- ❑ Eligible Shopify products are listed in Shopify Catalog by default, with nothing for the merchant to configure [\[C38\]](#).
- ❑ From there they are syndicated to ChatGPT, Microsoft Copilot, AI Mode in Google Search and the Gemini app [\[C38\]](#).
- ❑ Shopify reports that searches served from that structured data convert twice as well as searches served from general web data [\[C37\]](#), which is a company statement about its own product and is used here as one.

WHERE THE PLATFORMS ARE COMPETING, AND WHERE THEY ARE NOT

- ❑ **The specification's four capabilities are all transaction-side:** checkout, identity linking, order events and payment token exchange [\[C42\]](#). Being found is not in it.
- ❑ **Shopify Payments handled 68% of the platform's global merchandise volume** in the second quarter of 2026, three points more than a year earlier [\[C45\]](#).
- ❑ **Shopify now sells the agent channel to businesses that are not on Shopify**, who can sync products into its catalogue without replatforming [\[C46\]](#).
- ❑ **The one in-chat checkout that launched reached about a dozen Shopify merchants** before it was scaled back [\[C43\]](#).

PART THREE

The law and your ad costs now push the same way

Every fact a shopping assistant wants early is a fact a regulator now also wants early. That is not a coincidence. Both are asking the same question: is the commitment true at the point it is made?

£4.7m

in fines in the regulator's first year of direct consumer enforcement [\[C29\]](#)

29 Aug 2025

duty-free de minimis suspended for all countries into the US [\[C33\]](#)

28 Jun 2025

the European Accessibility Act starts applying to ecommerce [\[C34\]](#)

UK regulators are fining businesses that reveal costs late

Under the Digital Markets, Competition and Consumers Act the Competition and Markets Authority can act against consumer law breaches directly, without going to court first. Drip pricing, where part of the price appears only later in the journey, is treated as automatically unfair.

£4.7m

in fines in the first year of direct enforcement, across 14 investigations [\[C29\]](#)

CMA · YEAR ONE

£760k

in consumer refunds ordered, alongside 157 advisory letters and 46 information notices [\[C29\]](#)

CMA · YEAR ONE

£4.2m

fined against the AA's driving school businesses for drip pricing. The business admitted liability [\[C30\]](#)

CMA · APRIL 2026

READ THE ENFORCEMENT AS A DESIGN INSTRUCTION

Two of the regulator's three stated priorities are about **when** a fact appears rather than whether it appears at all. Drip pricing is a sequencing offence. Online choice architecture, meaning defaults and pressure tactics in the journey, is a sequencing offence. The third is fake reviews: 54 review publishers were contacted and 90% took compliance action [\[C31\]](#).

New subscription contract rules are expected in force in spring 2027 [\[C31\]](#), so the direction of travel is not finished.

- Any charge a customer cannot avoid belongs in the headline price, not in a later step.
- A pre-ticked extra in the payment journey is a default the regulator has named [\[C31\]](#).
- Scarcity and countdown devices have to be literally true, and provable if asked.
- Review displays need a policy behind them, not just a widget.
- "From" pricing needs its condition attached where the price is shown.
- None of this is legal advice. It is what the published enforcement suggests operators check first.

Import duty and accessibility are now legal requirements

Two changes made cross-border totals harder to state and more important to state.

The United States suspended duty-free de minimis treatment for all countries

From 29 August 2025, goods valued at or below \$800 from anywhere lost de minimis eligibility and became subject to applicable duties, taxes and fees ^[C33]. The suspension was continued by a further presidential action in February 2026.

The European Union is removing the €150 customs duty exemption

The European Commission announced in November 2025 that the threshold would go, with an interim method of calculating duty until the EU Customs Data Hub is operational, and an ecommerce handling fee mandated from November 2026 ^[C32].

The European Accessibility Act has applied to ecommerce services since 28 June 2025

Microenterprises providing services are exempt, and services lawfully provided before that date have transitional relief to 2030 ^[C34]. The obligation is not a badge. It is that a person using a screen reader, a keyboard or a magnifier can complete the same purchase.

WHAT THIS CHANGES ON A STOREFRONT

For a British store selling internationally, "plus any local duties" is no longer an acceptable answer. It was already a poor one. It is now a poor answer at the moment a comparison engine is deciding whether to quote your price at all.

ULTRALABS ANALYSIS

Every question in this report has an accessibility condition attached: **a fact that cannot be perceived has not been stated**. That is why accessibility is not a separate chapter here. It runs through all five questions.

Advertising is getting more expensive

26%

growth in Amazon's advertising revenue year on year, faster than its own online stores at 15% and its marketplace services at 16% ^[C23, C24].

FIGURE 8 Year-on-year growth in the businesses that sell attention, Q2 2026



Source. Meta Platforms, second quarter 2026 results, 29 July 2026 ^[C22]; Amazon.com, second quarter 2026 results, 30 July 2026 ^[C23, C24]. Meta figures are Family of Apps advertising metrics. Amazon figures are reported segment net sales.

What this means. Meta delivered 14% more adverts and charged 12% more for each one ^[C22]. At Amazon, the fastest-growing line is neither the shop nor the marketplace. It is advertising. The largest marketplace in the world is growing fastest by charging sellers to be found on it, and the largest social advertiser is raising its price per unit of attention.

WHAT THIS COSTS AN OPERATOR WHO DOES NOTHING

Both of these businesses price access to demand, and both raised that price over the year ^[C22, C23]. A store whose product facts are thin has one way to stay visible, which is to keep paying the higher price. A store whose facts are complete has two, because the routes that read published data rather than bids cost nothing per click ^[C21]. That is the commercial argument for the rest of this report, and it does not depend on any platform getting worse.

How much of your demand depends on one channel?

A question we are asked often enough to answer in print: is paid social going to get harder?

We have no forecast, and the published numbers do not support one in either direction.

They do support a reading of the risk and a list of things to watch.

28%

growth in Meta's revenue year on year in the second quarter of 2026 ^[C44]

META PLATFORMS · Q2 2026

55%

growth in Meta's total costs and expenses over the same quarter, about a quarter of it one-off legal and severance charges ^[C44, C48]

META PLATFORMS · Q2 2026

\$130bn

the floor of Meta's own guidance for 2026 capital spending, topping out at \$145bn, against \$60.80bn of revenue in the quarter ^[C44]

META PLATFORMS · GUIDANCE

WHY SOME OPERATORS EXPECT IT TO GET HARDER

- ❑ **Costs are growing faster than revenue**, and capital spending is guided well above a quarter's revenue ^[C44]. A business funded by advertising has reasons to want more from each unit of attention. **Read the 55% carefully:** it includes \$2.40bn of legal charges and \$1.18bn of severance, so roughly a quarter of the increase is one-off rather than recurring ^[C48].
- ❑ **The price is already moving.** Meta's average price per advert rose 12% year on year ^[C22].
- ❑ **Some of the context an advertiser used to buy is moving to the assistant.** Google began letting shoppers connect Gmail and Google Photos to AI Mode in January 2026, so answers can use what the model already knows about them ^[C47].

WHAT THE PUBLISHED NUMBERS SHOW INSTEAD

- ❑ **The business is growing.** Meta delivered 14% more adverts, charged 12% more for each, and grew revenue 28% ^[C22, C44]. That is not the shape of a channel in retreat.
- ❑ **Advertising and assistants do different jobs.** Advertising creates demand that did not exist. Most of what an assistant does is serve demand that already does.
- ❑ **Meta is inside the standard, not outside it.** It took a seat on the technical council of the agent commerce specification in April 2026 ^[C41].
- ❑ **The personal-context feature is small so far.** It is opt-in, runs as a Labs feature, and is limited to paying subscribers in one country ^[C47].

ULTRALABS ANALYSIS


We are not forecasting a slowdown, and nothing here should be read as one. Our confidence that paid social gets structurally harder is low, because the evidence runs both ways. Four things would change that reading, and all four are public: the split between volume and price in Meta's quarterly release ^[C22]; whether AI-referred retail traffic keeps compounding or settles ^[C15]; whether Shopify's AI multiples keep falling towards ordinary channel growth ^[C35, C36]; and whether the agent standard extends past the transaction into discovery ^[C42].

The number that decides this for you is not any of Meta's. It is the share of your revenue arriving through one paid channel, which you can calculate this week.

Why shoppers abandon carts

Cart abandonment is the most quoted number in ecommerce and one of the weakest. Baymard's often-cited 70.22% is an average across fifty studies published between 2006 and 2025 ^[C27]. Treat it as a reminder that most sessions do not end in an order, not as a benchmark you can move. The reasons underneath it are far more useful.

FIGURE 9 Why shoppers abandoned, excluding those who were only browsing



Source. Baymard Institute, page updated 22 September 2025 ^[C28]. Respondents could give more than one reason. **Baymard does not publish the sample size or fielding date for this study**, so it is indicative. The right-hand column is our mapping to the five questions, not Baymard's.

What this means. The two largest reasons and the eighth are the same failure stated three ways: the total was not knowable early enough ^[C28].

Four of the eight most common reasons for abandoning a cart are facts the page could have stated earlier.

PART FOUR

The product page audit

Everything from here is UltraLabs analysis: a framework built on the evidence in parts one to three, not a finding from any source in them. It is meant to be argued with, and it is meant to be used on a real page this week.

Five

questions a buyer has to settle before they can order

Four

places your store can answer them

Twenty

cells, each one answered, contradicted or missing

Five questions every product page should answer

Not benefits, and not reasons to buy. These are the five facts that have to be true and findable before an order can happen, and they are the same five whether the reader is a person, a price comparison engine or a shopping assistant.

1	Is this exactly the product I mean? The right product	GOES WRONG AS Wrong variant, missing specification, unstated compatibility, ambiguous condition or edition.
2	What will I pay in total? The total price	GOES WRONG AS Delivery, tax, duty or handling revealed after the shopper has committed attention.
3	Do you have it in stock? Stock	GOES WRONG AS Stock state absent, stale, or true in the warehouse and false on the page.
4	When will it arrive? Delivery	GOES WRONG AS No date before checkout, no tracking, no plan for the delivery that goes wrong.
5	What happens if it is wrong? Returns	GOES WRONG AS Returns terms unfindable, conditional in ways the page does not say, or contradicted on arrival.

WHY FIVE AND NOT FIFTEEN

Because these are the five that stop a transaction. Brand, story, photography and reviews all change how much someone wants the thing. None of them changes whether the order can be placed. A store can be weak on the first list and still sell. It cannot be silent on the second and be chosen by something that reads before it clicks.

THE NAME FOR THE PROBLEM

Every store answers these five questions somewhere. The problem is where. Most answer them in the order the funnel needs rather than the order the buyer wants: the product on the page, the total at the basket, stock at the variant, delivery in the confirmation email, returns in a footer link.

The distance between where your store answers a question and where the buyer asks it is what we call the resolution gap.

Where your store answers, and how answers go wrong

A store answers those questions in four places, and they are not equivalent. The first is the one most stores treat as an afterthought, and it is the only one a machine deciding whether to mention you can read.

S1 Listing and feed Everything readable without interacting: the page HTML, the category listing, the product feed, whatever a crawler or shopping agent can fetch. Read by machines first	S2 Product page The product and collection pages a person actually reads once they have arrived and are considering the purchase. Read by the shopper	S3 Basket and checkout Where the number stops being indicative and becomes the amount the customer is charged. Where it becomes binding	S4 After the order Order confirmation, tracking, delivery exceptions, returns and the support conversation. Where it is kept or broken
---	--	--	---

THE RULE

Every question must be answerable on the listing and feed, and the answer must be the same in all four places. If the fact is not in the listing, the store is not in the consideration set, and no analytics package will tell you it happened.

FAILURE ONE: A MISSING ANSWER

A question first answered further down the journey. Delivery cost that only appears in the basket. Stock state that only appears once a variant is chosen.

What it costs. Consideration you never observe. The reader that could not settle the question never became a session, so there is no bounce and no exit page to find.

FAILURE TWO: A CONTRADICTIONARY ANSWER

The same question answered differently in two places. £4.99 delivery on the listing and £7.95 at checkout. "In stock" on the category page and a three-week lead time on the product page.

What it costs. The sale at the point of highest intent. In the UK it is also the exact shape of the practice the regulator has been fining ^[C30].

Most teams fix the missing answers first, because a gap looks like work that has not been done and a contradiction looks like a detail. We would reverse that. A contradiction is usually cheap to correct, it is the one with legal exposure, and it is destroying the most expensive traffic you have.

The twenty checks in one grid

Five questions down, four places across. The shaded column is the one that decides whether you are in the consideration set at all.

WHAT A GOOD ANSWER LOOKS LIKE IN EACH CELL · ULTRALABS ANALYSIS

	S1 Listing and feed Read by machines first	S2 Product page Read by the shopper	S3 Basket and checkout Where it becomes binding	S4 After the order Where it is kept or broken
Q1 The right product	☐ ☐ ☐ Title, variant, key specification and condition present in the markup and the feed, not only in a photograph.	☐ ☐ ☐ Enough detail to tell two similar products apart without contacting anyone.	☐ ☐ ☐ The basket line names the exact variant that was chosen.	☐ ☐ ☐ The confirmation and the packing note match the page.
Q2 The total price	☐ ☐ ☐ Price, delivery and any duty or fee knowable before interaction, per destination.	☐ ☐ ☐ The full total stated on the product page, not deferred to the basket.	☐ ☐ ☐ The amount charged equals the amount stated. Nothing new appears.	☐ ☐ ☐ The receipt breaks down what was charged and why.
Q3 Stock	☐ ☐ ☐ Stock state per variant, current, and honest about backorder or made to order.	☐ ☐ ☐ Availability shown where the variant is chosen, not only at the basket.	☐ ☐ ☐ No silent substitution, no quantity surprise at payment.	☐ ☐ ☐ A short-shipped order is explained before the customer notices.
Q4 Delivery	☐ ☐ ☐ A delivery expectation attached to the product and destination, not a generic policy page.	☐ ☐ ☐ A date or a range on the product page, with the cut-off that produces it.	☐ ☐ ☐ The choice of service, its cost and its date, in one view.	☐ ☐ ☐ Tracking that works, and a route when it does not arrive.
Q5 Returns	☐ ☐ ☐ Return window, who pays and any exclusions, readable without a click.	☐ ☐ ☐ The same terms next to the buy button, in the shopper's own words.	☐ ☐ ☐ No new condition introduced at the moment of payment.	☐ ☐ ☐ A return that can be started without an email, and a refund with a stated time.



Answered

2 points

Present, specific and current in that place.



Contradicted

0 points

Present, but it does not match another place.



Missing

1 point

Not there at all in that place.

How to score your own product pages

One product, twenty cells, three states. Print this page or copy the grid. It is deliberately empty: a scorecard filled in with somebody else's store teaches nothing.

PRODUCT	S1 LISTING AND FEED	S2 PRODUCT PAGE	S3 BASKET AND CHECKOUT	S4 AFTER THE ORDER
Q1 The right product	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐
Q2 The total price	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐
Q3 Stock	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐
Q4 Delivery	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐
Q5 Returns	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐	☐ ☐ ☐
SCORE OUT OF 40 · DATE · WHO CHECKED · PRODUCT URL				

HOW TO RUN IT IN AN AFTERNOON

1. Pick one product that matters commercially and has more than one variant. Your worst performer teaches more than your best.
2. Score the listing and feed first, from the page source and the product feed only. No clicking. That is the column everything else is judged against.
3. Then walk the other three as a customer, from a private window, in the country you sell most into. Mark it contradicted the moment a number changes.
4. Repeat for two more products. Three is enough to tell a product problem from a platform problem, and that distinction decides what the fix costs.

TWO HONEST LIMITS

This is a diagnostic, not a metric. It has no benchmark, no industry average and no published distribution, because we have not measured one and will not invent one.

Its use is comparative: the same product before and after, or five products against each other. Score it internally and do not publish it as a marketing claim.

1

Question one: is this exactly the product I mean?

Identity failures rarely look like design problems. They look like a catalogue that grew, a supplier feed imported once, and variant names that made sense to whoever created them.

WHERE IT USUALLY GOES WRONG

- ❑ Variant names describing the internal SKU rather than the difference a customer cares about.
- ❑ Specifications living in a PDF, an image or a paragraph rather than in structured fields.
- ❑ Compatibility and fit stated in the reviews because it was never stated on the page.
- ❑ Two near-identical products with no page explaining which one to buy.

WHAT A GOOD ANSWER LOOKS LIKE

- ❑ The distinguishing attribute is in the title, the variant name and the feed, in the words a customer would use.
- ❑ Specifications are fields, not prose, so they survive being read by something other than a browser.
- ❑ Condition, edition, quantity and units are explicit. "Pack" is not a unit.
- ❑ Where two products overlap, one page says plainly which is for whom.

THE FIVE-MINUTE CHECK

Take a product with more than four variants. Read only the feed and the page source. If you cannot tell the variants apart from that alone, neither can anything else.

EVIDENCE BEHIND IT

Google recommends a Merchant Center product feed for visibility across its AI answers and its ordinary results, while stating that no special AI markup exists ^[C21].

THE RIGHT PRODUCT ACROSS THE FOUR PLACES

S1 LISTING AND FEED

Title, variant, key specification and condition present in the markup and the feed, not only in a photograph.

S2 PRODUCT PAGE

Enough detail to tell two similar products apart without contacting anyone.

S3 BASKET AND CHECKOUT

The basket line names the exact variant that was chosen.

S4 AFTER THE ORDER

The confirmation and the packing note match the page.

2

Question two: what will I pay in total?

This is the question with the most evidence behind it and the most reasons to get right. It is the largest stated cause of abandonment, it is the practice the regulator is fining, and it is the fact a comparison engine most needs before it will quote you.

40%

of shoppers who abandoned for a reason other than browsing blamed **extra costs being too high**. A further 12% said they could not calculate the total up front ^[C28].

WHERE IT USUALLY GOES WRONG

- ❑ Delivery cost calculated only once an address has been entered.
- ❑ A free-delivery threshold advertised but never shown against the current basket.
- ❑ Duty and import charges described as "may apply" on international orders.
- ❑ Surcharges for size, weight, remote areas or payment method appearing at the last step.

WHAT A GOOD ANSWER LOOKS LIKE

- ❑ A total a shopper in a named destination can find before interacting with anything.
- ❑ Delivery and any fee expressed against the product, not only in a policy page.
- ❑ For cross-border orders, either duty-paid pricing or an explicit, accurate estimate.
- ❑ The amount charged equals the amount shown, including currency rounding.

THE FIVE-MINUTE CHECK

Open your own product page in a private window from the country you sell most into. Time how long it takes to learn the total. More than one interaction is a gap.

THE TOTAL PRICE ACROSS THE FOUR PLACES

S1 LISTING AND FEED Price, delivery and any duty or fee knowable before interaction, per destination.	S2 PRODUCT PAGE The full total stated on the product page, not deferred to the basket.
S3 BASKET AND CHECKOUT The amount charged equals the amount stated. Nothing new appears.	S4 AFTER THE ORDER The receipt breaks down what was charged and why.

3

Question three: do you have it in stock?

Stock is the question most likely to be present and wrong, because it depends on operations rather than on the website. A stock figure is a claim about a warehouse, and a website is the last place to find out the claim stopped being true.

WHERE IT USUALLY GOES WRONG

- ❑ Stock state shown at product level while the variants underneath it differ.
- ❑ Availability reflecting a nightly sync rather than the current position.
- ❑ Backorder and made-to-order presented as ordinary stock, with the truth arriving by email.
- ❑ A pre-order with no stated dispatch date.

WHAT A GOOD ANSWER LOOKS LIKE

- ❑ Per-variant state, current enough that a shopper acting on it is not misled.
- ❑ Backorder, pre-order and made to order named as themselves, with the wait attached.
- ❑ Low stock stated because it is true, not because it is persuasive.
- ❑ An out-of-stock page that stays available and offers a real alternative or a notification.

THE FIVE-MINUTE CHECK

Pick the variant you sell out of most often. Compare what the page says, what the feed says and what the stock system says, at the same moment.

EVIDENCE BEHIND IT

This is the one question with no directly citable 2026 figure in this report. It is here because the other four cannot be honoured without it.

STOCK ACROSS THE FOUR PLACES

S1 LISTING AND FEED

Stock state per variant, current, and honest about backorder or made to order.

S2 PRODUCT PAGE

Availability shown where the variant is chosen, not only at the basket.

S3 BASKET AND CHECKOUT

No silent substitution, no quantity surprise at payment.

S4 AFTER THE ORDER

A short-shipped order is explained before the customer notices.

4

Question four: when will it arrive?

Delivery speed is the second largest stated reason for abandonment, and most stores answer it with a page rather than a date. A policy page describes a service. A shopper is asking about one order.

20%

of shoppers who abandoned for a reason other than browsing blamed **delivery being too slow**, the second largest stated reason of the eight ^[C28].

WHERE IT USUALLY GOES WRONG

- ❑ A shipping page listing services and prices, with no date anywhere near the product.
- ❑ Cut-off times described in the abstract rather than counted down against now.
- ❑ International delivery given as a range so wide it carries no information.
- ❑ Tracking that stops updating at the border, and a support inbox that finds out last.

WHAT A GOOD ANSWER LOOKS LIKE

- ❑ A date or a tight range on the product page, with the cut-off that produces it.
- ❑ The same date carried into the basket, the checkout and the confirmation.
- ❑ Service choices priced and dated together, so the trade-off is visible in one view.
- ❑ A stated route for a delivery that goes wrong, before the customer has to find one.

THE FIVE-MINUTE CHECK

Read your delivery promise as a machine would. Is there a date attached to this product and this destination, or only a page describing your carriers?

DELIVERY ACROSS THE FOUR PLACES

S1 LISTING AND FEED A delivery expectation attached to the product and destination, not a generic policy page.	S2 PRODUCT PAGE A date or a range on the product page, with the cut-off that produces it.
S3 BASKET AND CHECKOUT The choice of service, its cost and its date, in one view.	S4 AFTER THE ORDER Tracking that works, and a route when it does not arrive.

5

Question five: what happens if it is wrong?

Returns are the cheapest reassurance in ecommerce and the most commonly buried. This is also the question most likely to be contradicted, because the terms are usually written by one part of the business and the reassurance by another.

WHERE IT USUALLY GOES WRONG

- ❑ A returns window stated in the header and qualified in the terms.
- ❑ "Free returns" that means free to send back but not free to have refunded.
- ❑ Exclusions for sale items, hygiene items or bulky goods discovered at the point of return.
- ❑ A process that begins with an email to an address nobody watches on Fridays.

WHAT A GOOD ANSWER LOOKS LIKE

- ❑ Window, who pays, and exclusions, in one place a machine can read and a shopper can find.
- ❑ The same terms next to the buy button and inside the checkout, in identical words.
- ❑ A return that can be started without a conversation.
- ❑ A refund with a stated time, and a message when it is sent.

THE FIVE-MINUTE CHECK

Ask someone outside the business to find your returns terms and tell you who pays for postage. Time it. Then check that answer against your terms page and your checkout.

EVIDENCE BEHIND IT

13% of non-browsing abandonment is attributed to an unsatisfactory returns policy, and 19% to distrust of payment security, which is the same instinct wearing a different coat ^[C28].

RETURNS ACROSS THE FOUR PLACES

S1 LISTING AND FEED

Return window, who pays and any exclusions, readable without a click.

S2 PRODUCT PAGE

The same terms next to the buy button, in the shopper's own words.

S3 BASKET AND CHECKOUT

No new condition introduced at the moment of payment.

S4 AFTER THE ORDER

A return that can be started without an email, and a refund with a stated time.

Fix these in this order

Plotted by what it costs you to leave alone against what it costs to fix. UltraLabs analysis, and a starting point rather than a plan: your own scorecard decides which cells you are actually in.

		CHEAP TO FIX	STRUCTURAL
COST OF LEAVING IT	DO IT NOW	— Reconcile the delivery number across listing, product page, basket and checkout. — Move the returns terms next to the buy button, in the same words as the terms page. — Remove any urgency device that is not literally true. — Fix variant naming so the difference is legible without a photograph.	— Duty-paid or accurately estimated international pricing ^[C33, C32]. — Real stock state, per variant, synced often enough to be true when someone acts on it. — An accessibility pass on the buying path, which is a legal obligation in the EU ^[C34].
	WORTH DOING	— Specifications as fields rather than prose. — A product feed that matches the site, checked on a schedule rather than at launch. — Delivery dates on the product page, with the cut-off that produces them.	— Agentic checkout integrations. The transaction leg changed shape twice in twelve months ^[C18, C19]. — Markup added as a growth tactic. Useful for finding a fact, useless if the fact is missing ^[C21]. — Replatforming to fix something a data and content pass would fix.
		LOW EFFORT	HIGH EFFORT

Fix the contradictions first. They are cheap to correct and they cost you the traffic you paid most for.

ULTRALABS ANALYSIS

HONESTY

A framework that claims to solve everything is a framework nobody should use.

What this audit will not fix

A product nobody wants

Answering questions makes a store legible. It does not create demand. If the constraint is the proposition, no amount of clean data moves it.

A price that is not competitive

Being unambiguous about an uncompetitive total gets you compared accurately and rejected faster. Better than being compared inaccurately, but it is not growth.

An audience that has never heard of you

Closing gaps improves your chance of being in a consideration set. It does not put you in front of people who do not know you exist.

Operational reality

Stock and delivery are honest only if the warehouse and the carrier make them honest. A website can stop lying. It cannot make the stock arrive.

The unsettled part of agentic commerce

We take a position on discovery, which is live and measurable ^[C15, C16]. We decline to forecast in-chat checkout, which changed shape twice in twelve months ^[C18, C19].

HOW THIS REPORT COULD BE WRONG

The load-bearing inference is that AI-referred traffic converts better because the questions were settled before the click ^[C16]. A different explanation fits the same data: assistants may simply be sending people who were already close to buying. That would weaken the urgency here. It would not change a single recommendation, because every one is also required by the abandonment evidence ^[C28], the enforcement record ^[C29] or accessibility law ^[C34].

AND ONE THING WE CANNOT TELL YOU

What a closed gap is worth in revenue. We have not measured it, nobody has published a controlled test of it, and a number invented for a report would undermine everything else in it.

How we researched this report

HOW THE EVIDENCE WAS SELECTED

- Official statistics, regulators, company filings, platform documentation and research institutions were preferred over trade press, and trade press over vendor content.
- Every claim had to name a publisher, a publication date and a measurement period. Anything undated or methodologically vague was dropped.
- Where a widely quoted figure traced back to another report, we went to the original or excluded the claim.
- Nothing here is a survey, interview, case study or forecast produced by UltraLabs. We conducted none, so we report none.

WHAT WAS EXCLUDED, AND WHY

- **Nine claims** were dropped because the original was inaccessible or did not exist. Two were the most on-thesis statistics we found all year.
- A widely circulated 2026 figure on the share of orders placed by agents traced only to aggregator write-ups. Excluded.
- Several 2026 click-through studies from search agencies disagreed with each other by a factor of two and did not disclose sampling. Excluded in favour of the Pew panel [C08].
- Return-rate benchmarks circulated in 2026 traced to vendor blogs citing one another. Excluded, which is why returns appear here without a headline number.
- Published survey figures for how many shoppers will buy inside an AI chat ranged from about a sixth to about four fifths depending on how the question was asked. Excluded, which is why no adoption percentage appears anywhere in this report.

HOW OPINION WRITING WAS USED

Commentary and newsletters were read during research and none of it is cited as evidence here. Where an argument in that writing looked material, we went to the filing, the specification or the platform's own documentation and cited that instead, which changed several of the numbers on the way. A widely repeated figure of sixty billion listings in Google's Shopping Graph is over fifty billion in Google's own remarks [C40]. A claim that fewer than thirty merchants went live with in-chat checkout is about a dozen when traced to the named attribution [C43]. A claim that Google's assistant reads Gmail, Photos and Drive is Gmail and Photos, opt-in, in one country [C47]. Opinion is useful for finding the question. It is not a source.

LIMITS YOU SHOULD HOLD AGAINST US

- **The two national series are not comparable as levels.** Different geographies, different classifications, and the American series excludes businesses with no paid employees since April 2025 [C05, C01].
- **Cloudflare and Adobe both report from their own networks.** Both are very large. Neither is a census [C11, C15].
- **Two central studies are from 2025**, labelled everywhere they appear [C08, C14].
- **One central dataset is carried through trade reporting**, because the publisher's own pages were unreachable from our build environment on 13 August 2026 [C15].
- **Anything a platform says about its own product is a company statement**, not an independent measurement [C18, C20, C21, C25].
- **The framework is untested at scale.** It is a reading of the evidence and a diagnostic we use in client work. It has no published distribution and no benchmark, and we have not invented one.

REPRODUCING THIS

The full source ledger, the working notes including the contradictions we kept, and the editorial plan are published alongside this report. Every chart is generated from its cited series at build time rather than drawn by hand, so a figure cannot drift from its source without the build changing.

Sources

Every claim reference in this report resolves to one of these rows. Sources were accessed on 13 and 14 August 2026. The list continues on the next page.

[C01](#), [C02](#), [C03](#)

Retail sales, Great Britain: June 2026. Office for National Statistics, 24 July 2026.
<https://www.ons.gov.uk/businessindustryandtrade/retailindustry/bulletins/retailsales/june2026>

[C04](#)

Internet sales as a percentage of total retail sales (ratio) (%), series J4MC. Office for National Statistics, Retail Sales Index (DRSI), released 24 July 2026.
<https://www.ons.gov.uk/businessindustryandtrade/retailindustry/timeseries/j4mc/drsi>

[C05](#), [C06](#), [C07](#)

Quarterly Retail E-Commerce Sales, 1st Quarter 2026 (CB26-81). US Census Bureau, 18 May 2026.
<https://www2.census.gov/retail/releases/historical/ecom/26q1.pdf>

[C08](#), [C09](#), [C10](#)

Google users are less likely to click on links when an AI summary appears in the results. Pew Research Center, 22 July 2025.
<https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>

[C11](#), [C12](#), [C13](#)

Content Independence Day, one year on: building the business model for the agentic Internet. Cloudflare, 1 July 2026.
<https://blog.cloudflare.com/agentic-internet-bot-report/>

[C14](#)

A deeper look at AI crawlers: breaking down traffic by purpose and industry. Cloudflare, 28 August 2025.
<https://blog.cloudflare.com/ai-crawler-traffic-by-purpose-and-industry/>

[C15](#), [C16](#), [C17](#)

Adobe: AI-referred traffic to retail sites doubles in a year. Digital Commerce 360 reporting Adobe Analytics, 17 June 2026.
<https://www.digitalcommerce360.com/2026/06/17/adobe-ai-referred-traffic-to-retail-sites-doubles-in-a-year/>

[C18](#)

Stripe powers Instant Checkout in ChatGPT and releases Agentic Commerce Protocol codeveloped with OpenAI. Stripe, 29 September 2025.
<https://stripe.com/newsroom/news/stripe-openai-instant-checkout>

[C19](#)

OpenAI shifts checkout plans in its agentic commerce strategy. Digital Commerce 360, 6 March 2026.
<https://www.digitalcommerce360.com/2026/03/06/openai-shifts-checkout-plans-agentic-commerce-strategy/>

[C20](#)

Conversations that Convert: Copilot Checkout and Brand Agents. Microsoft Advertising, 8 January 2026.
<https://about.ads.microsoft.com/en/blog/post/january-2026/conversations-that-convert-copilot-checkout-and-brand-agents>

[C21](#)

Google's Guide to Optimizing for Generative AI Features on Google Search. Google Search Central, updated 10 July 2026.
<https://developers.google.com/search/docs/fundamentals/ai-optimization-guide>

[C22](#), [C44](#), [C48](#)

Meta Reports Second Quarter 2026 Results. Meta Platforms, Inc., 29 July 2026.
<https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Second-Quarter-2026-Results/default.aspx>

[C23](#), [C24](#)

Amazon.com Announces Second Quarter Results. Amazon.com, Inc., 30 July 2026.
https://s2.q4cdn.com/299287126/files/doc_earnings/2026/q2/earnings-result/AMZN-Q2-2026-Earnings-Release.pdf

[C25](#), [C26](#)

Shopify second quarter 2026 results (Form 8-K, Exhibit 99.1). Shopify Inc. via SEC EDGAR, 5 August 2026.
<https://www.sec.gov/Archives/edgar/data/0001594805/000159480526000046/exhibit991pressreleaseq220.htm>

[C27](#), [C28](#)

50 Cart Abandonment Rate Statistics. Baymard Institute, page updated 22 September 2025.
<https://baymard.com/lists/cart-abandonment-rate>

Sources, continued

The platform and company sources. Everything a company published about its own product is marked as a company statement in the ledger rather than treated as an independent measurement.

C29, C30, C31

Direct consumer enforcement: one year on. Competition and Markets Authority, 17 April 2026.
<https://competitionandmarkets.blog.gov.uk/2026/04/17/direct-consumer-enforcement-one-year-on/>

C32

E-commerce: 150 EUR customs duty exemption threshold to be removed as of 2026. European Commission, Taxation and Customs Union, 13 November 2025.
https://taxation-customs.ec.europa.eu/news/e-commerce-150-eur-customs-duty-exemption-threshold-be-removed-2026-2025-11-13_en

C33

Suspending Duty-Free De Minimis Treatment for All Countries (Executive Order 14324). The White House, 30 July 2025.
<https://www.whitehouse.gov/presidential-actions/2025/07/suspending-duty-free-de-minimis-treatment-for-all-countries/>

C34

Directive (EU) 2019/882 on the accessibility requirements for products and services. EUR-Lex, European Union, applying from 28 June 2025.
<https://eur-lex.europa.eu/eli/dir/2019/882/oj/eng>

C35, C37

Shopify (SHOP) Q1 2026 Earnings Call Transcript. The Motley Fool, transcribing Shopify Inc., 5 May 2026.
<https://www.fool.com/earnings/call-transcripts/2026/05/05/shopify-s-hop-q1-2026-earnings-transcript/>

C36

Shopify says AI search is driving more traffic and sales, not replacing Google. TechCrunch, 5 August 2026.
<https://techcrunch.com/2026/08/05/shopify-says-ai-search-is-driving-more-traffic-and-sales-not-replacing-google/>

C38, C46

Agentic Commerce on Shopify: How It Works. Shopify, 18 June 2026.
<https://www.shopify.com/blog/how-agentic-commerce-works>

C39

New tech and tools for retailers to succeed in an agentic shopping era. Google, 11 January 2026.
<https://blog.google/products/ads-commerce/agentic-commerce-ai-tools-protocol-retailers-platforms/>

C40

The AI platform shift and the opportunity ahead for retail, NRF 2026 remarks. Google, Sundar Pichai, 11 January 2026.
<https://blog.google/company-news/inside-google/message-ceo/nrf-2026-remarks/>

C41

Announcing new Tech Council members, discussion 379. Universal Commerce Protocol project, GitHub, 24 April 2026.
<https://github.com/Universal-Commerce-Protocol/ucp/discussions/379>

C42

Specification and documentation for the Universal Commerce Protocol. Universal Commerce Protocol project, GitHub, read 14 August 2026.
<https://github.com/universal-commerce-protocol/ucp>

C43

Why OpenAI's checkout retreat spells trouble for its commerce strategy. Forbes, Jason Goldberg, 10 March 2026.
<https://www.forbes.com/sites/jasongoldberg/2026/03/10/why-openais-checkout-retreat-spells-trouble-for-its-commerce-strategy/>

C45

Shopify (SHOP) Q2 2026 Earnings Call Transcript. The Motley Fool, transcribing Shopify Inc., 12 August 2026.
<https://www.fool.com/earnings/call-transcripts/2026/08/12/shopify-s-hop-q2-2026-earnings-call-transcript/>

C47

Google brings Personal Intelligence to AI Mode in Search. Google, 22 January 2026.
<https://blog.google/products-and-platforms/products/search/personal-intelligence-ai-mode-search/>

PUBLISHED ALONGSIDE THIS REPORT

source-ledger.csv carries all 48 claims with publisher, publication and access dates, measurement period, geography, method note, classification and status. **research-notes.md** carries the contradictions we kept and every excluded claim with its reason. Both sit in the same repository as the build that produced this document.

Send us one product page

The audit is only useful if it produces an argument about a real page. So the offer is small on purpose, and it is a diagnostic rather than a pitch.

SEND US TWO THINGS

- One product URL. Your worst one is more useful than your best one.
- The constraint you are actually stuck on, in a sentence. Traffic, margin, returns, a platform you have outgrown, a catalogue nobody can maintain.

YOU GET BACK

- The twenty-cell scorecard for that product, filled in, with what we checked.
- The three gaps costing the most, and what closing each one actually involves.
- An honest answer about whether it needs a project. If nothing needs building, we will say so.

SEND IT TO

ultralabs.digital/contact



ULTRALABS LTD · BUCKINGHAM, UNITED KINGDOM

The state of ecommerce in 2026. Edition one · August 2026. Shopify builds, migrations and focused store work, and lead-generation websites. Operator-led, with no account-management layer between the problem and the people doing the work. No guaranteed revenue, conversion or performance outcome is offered or implied by this document.

ULTRALABS.