Can an AI Shopping Agent Explain Why Your Product Is Worth the Price?

By Basil C. Puglisi, Human-AI Collaboration Strategist

Shopify's October 6, 2026 holiday post carries a passage every small merchant should read slowly. "If an agent is asked whether a knife is worth the price, it will compare the knife with others. Price is the part of that comparison the agent can see immediately. Anything else, why this one, who it is for, what the money goes toward, has to be written where the agent can find it" (Shopify, 2026c). The same post's closing paragraph carries the line that sets the stakes: "an agent can only repeat the case a brand has already made" (Shopify, 2026c). When the only clear fact in a listing is the price, an AI shopping agent argues on price. A small shop facing cheaper versions of the same product on a large marketplace can lose that argument before it starts.

This article is about that argument. It explains what Shopify now sends to AI shopping agents and which settings a merchant controls. It also shows how to write the reasons for a price where an agent can read them. A merchant has a dated reason to care this week. Black Friday falls on November 27, 2026, and a change to Catalog access can take up to seven days to take effect (Shopify Help Center, n.d.-b). Five actions are listed first, and six sections then explain them, each ending in a step a small store can take and a measure to check it. The piece then reports the evidence against agentic shopping, so a reader can judge whether the work is worth the hours.

This analysis covers Shopify merchants, the 2026 holiday season, and the tools Shopify documented between October 6 and 13. Its survey figures come from Shopify's own research, and its brand results come from brands Shopify chose to feature. Shopify also operates the catalog, the agent connections, and the checkout those figures promote. The piece therefore reports them as vendor-reported documentation of what the platform is building, not as evidence of what any one store will see. Later developments through November 15, 2026 are recorded where they bear on a merchant's decision.

The method behind the analysis pairs each verifiable fact with the tactic it supports and a measure the merchant writes down before acting, which the author calls Factics. The term was introduced in 2012 for a practice already at work in the author's blogs and at Digital Ethos. It arrived in print that November in *Digital Factics: Twitter* (Puglisi, 2012). Its core principle, as the current framework page states it, is short: "Every fact must lead to a tactic, and every tactic must leave evidence" (Puglisi, 2026b). Each section below closes on that sequence.

What can a merchant do this week, and what should it expect?

Five actions carry the work, and the sections that follow explain each one.

1. **Record the settings.** Write down today's compare-at sharing, Catalog access by channel, auto-enrollment, and direct checkout, with the date.

2. **Put a name on each setting.** A named owner accepts, modifies, or rejects each one and records the reason by November 20, 2026.
3. **Support every compare-at price.** Remove any price the store cannot back with evidence that the higher price was openly offered.
4. **Write the value case into the product data.** For the products that carry the season's revenue, state who the product is for, what problem it solves, and why it costs what it costs. Give at least one fact a shopper could check.
5. **Set the baseline and the targets.** Record agentic sales, orders, sessions, and conversion now, and write a target for each before Black Friday.

What a merchant can expect is modest and measurable. A change to Catalog access can take up to seven days to take effect (Shopify Help Center, n.d.-b). Similarweb reports that AI referrals "remain a low single-digit percentage of most sites' total traffic" (Zohar Basteker, 2026). Retailers interviewed by Digital Commerce 360 reported AI-referred conversion on both sides of their site averages (Haleem, 2026). Shopify calls its own search preview "a directional signal, not an exact prediction" (Shopify Help Center, n.d.-b). The result a store can count on is the record itself, which shows what the store shares, who approved it, and whether the work paid for its hours.



Figure 1. The dates behind the five actions. A change to Catalog access can take up to seven days to take effect (Shopify Help Center, n.d.-b), and the holiday season as Adobe measures it closes on December 31 (Adobe, 2026).

What did Shopify change between October 6 and 13, 2026?

On October 6, Shopify published its holiday outlook, built on its 2026 Global Holiday Retail Report (Shopify, 2026c). The report finds that 65 percent of shoppers are likely to use AI for at least one shopping task over the holiday period (Shopify, 2026a). Among shoppers who almost always use AI

tools, 71 percent are more likely to choose independent or niche brands than last year. The comparable figure for shoppers who do not use AI tools is 45 percent. The consumer survey covered 18,000 shoppers aged 18 and over in nine markets, interviewed online by Sapio Research in July 2026. A second survey reached 7,500 owners and business decision-makers at companies with fewer than 1,000 employees in August 2026.

The report also measures how far shoppers will let an agent go. Half are comfortable letting AI recommend the right product, and 40 percent are comfortable letting AI build a cart as long as they complete the purchase themselves (Shopify, 2026a). Only 29 percent trust AI to make the purchase on their behalf. Shopify's news post describes that share as "one-third" (Shopify, 2026c), while the report's own chart gives the precise figure.

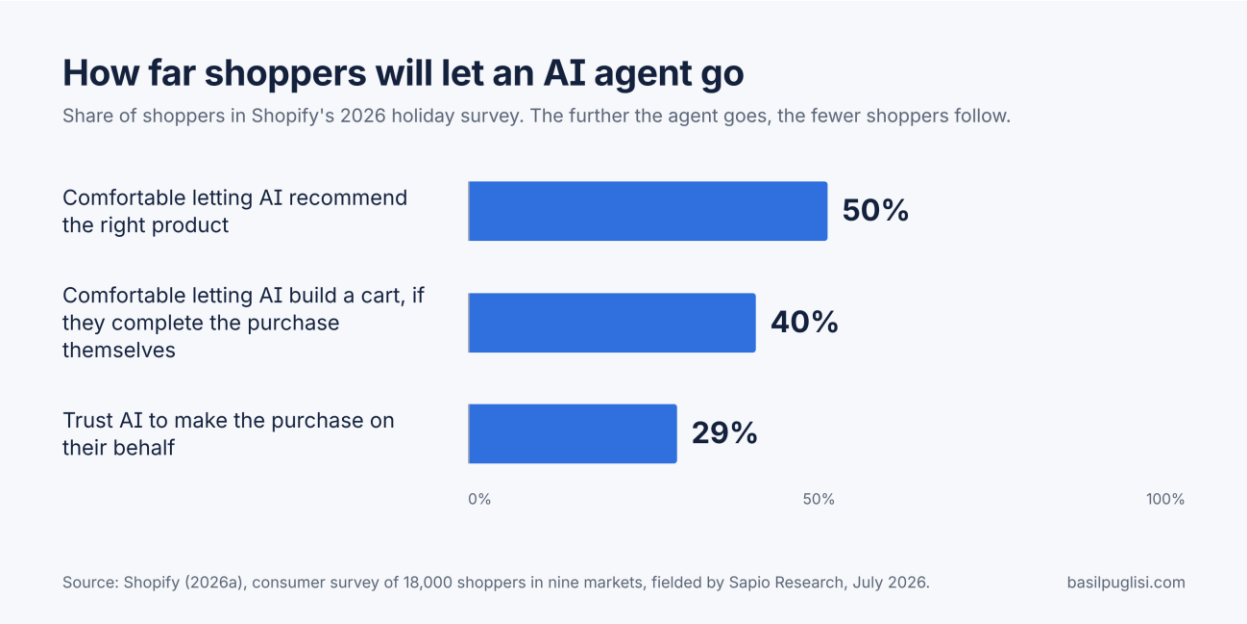


Figure 2. How far shoppers will let an AI agent go, from Shopify's 2026 consumer survey (Shopify, 2026a).

On October 9, Shopify began sending compare-at prices to AI shopping agents through Shopify Catalog, which the Help Center calls "the authoritative product data feed to your agentic channels" (Shopify, 2026f; Shopify Help Center, n.d.-c). The changelog says that before the change, "AI agents reading Catalog only saw the current price." It points merchants to a toggle in the Compare-at section of the admin (Shopify, 2026f). Catalog sends a compare-at price only when it is higher than the current price (Shopify Help Center, n.d.-a). The Help Center adds that agents "might use the list price to advertise strikethrough pricing for products that are on sale" (Shopify Help Center, n.d.-c). The Help Center presents the same control as a List price choice between Price and Compare-at price, and it states no default for that choice (Shopify Help Center, n.d.-a). Shopify's developer changelog of the same day is more specific: "Shopify turns this on by default for most stores and turns it off for stores that use compare-at for MSRP or another reference price" (Shopify, 2026e).

On October 13, Shopify's developer changelog announced Shop MCP. The hosted server lets outside agents "search Shopify Catalog, build carts, and complete Shop Pay checkout through a single MCP

connection" (Shopify, 2026g). Shop MCP is a remote Model Context Protocol server, and any MCP client can connect to it, as can a developer's own code (Shopify, n.d.-c). The post says it works with over 1 billion products in Catalog and needs no API key for search. Carts and checkout work without the shopper signing in. With a signed-in shopper, the agent "completes the purchase with complete_checkout after the shopper confirms the total and payment method" (Shopify, 2026g). Orders placed this way appear in the merchant's Shopify admin like any other order, and the post notes that there are no webhooks at launch (Shopify, 2026g).

None of this arrived without warning. Shopify's Spring '26 Edition, published June 17, 2026, called the Universal Commerce Protocol, co-developed with Google, "the open standard for agentic commerce from discovery to checkout." For Shopify merchants, the post said, Catalog and the protocol "are enabled by default" (Shopify, 2026b). October widened two things in the same week: what an agent can see about a product, and what it can do with it.

The widening continued after that week. On October 28, 2026, Shopify published its Fall '26 Edition, and two of its additions bear on this analysis (Shopify, 2026h). A Knowledge base page at Sales channels > Agentic lets a merchant "review and publish the facts agents get about their store, like return windows, shipping destinations, and gift services." Shopify drafts those facts from the storefront, and the post says nothing it drafts reaches an agent "without the merchant's sign-off" (Shopify, 2026h). A second view shows whether a store is eligible for Shopify Catalog and links to fixes for anything in the way. A merchant who assumes Catalog already carries the store's products can now check that assumption in the admin.

The fact is that, within one week, agents gained a view of markdowns and outside agents gained a path to checkout (Shopify, 2026f, 2026g). The report Shopify published that same week found that only 29 percent of shoppers trust AI to make a purchase on their behalf (Shopify, 2026a). The tactic is to record which of these changes touch the store and what each setting says today, before holiday traffic peaks. The measure is a dated one-page record of compare-at sharing, Catalog access by channel, auto-enrollment, and direct checkout. It comes first among the steps in this piece, because a change to Catalog access can take up to seven days to take effect (Shopify Help Center, n.d.-b).

Why does price travel farther than the reasons behind it?

Price is structured data. Catalog lists products by default "with their title, description, options, images, price, availability, and other key attributes, all structured in a way that AI agents can parse and understand" (Shopify Help Center, n.d.-c). A number compares cleanly across listings, and an agent can line it up against every competing product that answers the same request. The reasons a product costs more often sit in brand copy, a caption, a founder's story told on social media, or nowhere at all.

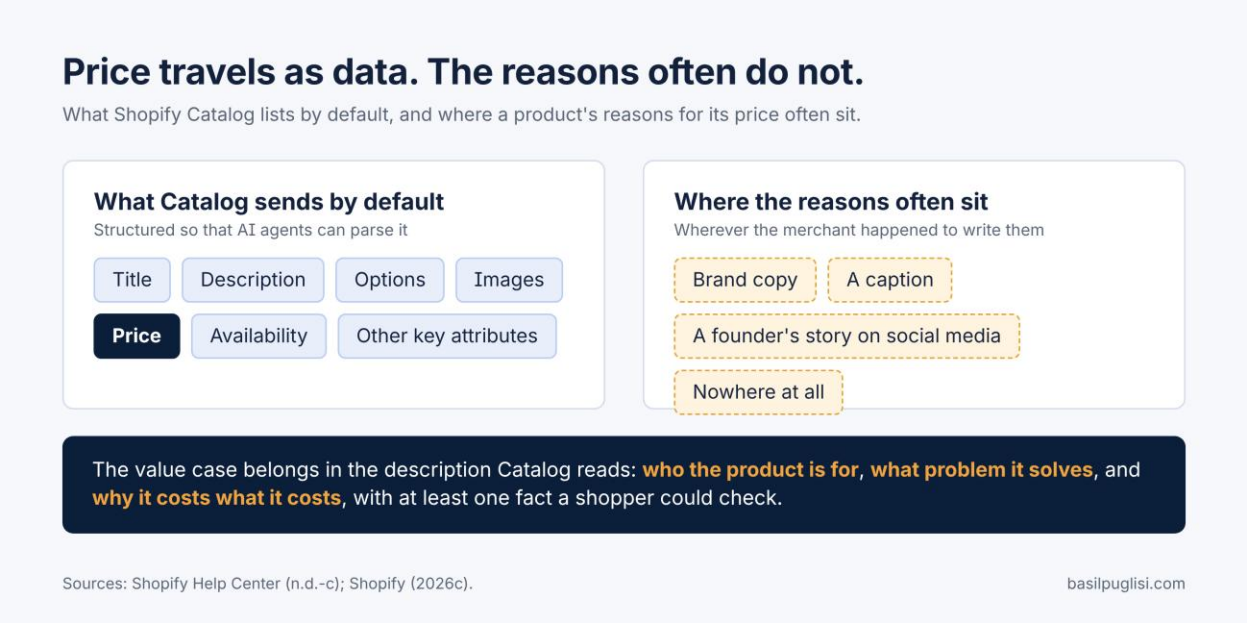


Figure 3. What Shopify Catalog lists by default, and where the reasons for a price often sit (Shopify Help Center, n.d.-c; Shopify, 2026c).

Shopify's survey finds a real audience for those reasons. In it, 35 percent of shoppers choose products that are the best possible version for their needs, even when they are not the cheapest option (Shopify, 2026a). Half define good value as higher-than-expected quality for the price. Clear information on price and discounts tops the insights that make shoppers more likely to use an AI tool, at 28 percent. The ability to compare similar options follows at 27 percent, and clear product reviews at 25 percent.

The pull toward price runs strong on both sides of the sale. Adobe expects discounts of up to 30 percent off listed price during Cyber Week (Adobe, 2026). In Shopify's survey, deeper or more frequent discounts lead the strategies businesses plan to use to drive holiday sales, at 31 percent (Shopify, 2026a). Checkout.com, a digital payments company, published consumer research in June 2026. It reported that 57 percent of consumers would let an AI shopping agent switch brands if it found a better value option (Checkout.com, 2026). An agent that cannot find the value case can still weigh price, reviews, availability, and other signals. It cannot explain the merchant's reasons for charging more when those reasons are missing from the product data it reads.

Material Kitchen, one of the brands Shopify featured, chose a different answer. Its Kinder Kitchens campaign gives customers the year's best pricing and lets them redirect a portion of their savings to Heart of Dinner, a nonprofit (Shopify, 2026c). That reason reaches an agent through Catalog only if it lives in the product data Catalog sends.

The fact is that Catalog sends structured price data by default, while the reasons for a price live wherever the merchant happened to write them (Shopify Help Center, n.d.-c; Shopify, 2026c). The tactic is to choose the products that carry most of the season's revenue and write the value case into each description. That case covers who the product is for, what problem it solves in the shopper's words, and why it costs what it costs. Each answer rests on something a shopper could check, such as

materials, construction, warranty terms, or a service commitment. The measure is the share of those chosen products whose description answers all three and carries at least one checkable fact. The process target is 100 percent before Black Friday, November 27, 2026.

What should a merchant decide about compare-at prices?

The compare-at setting is a pricing decision that now travels. A strikethrough price on a store's own product page is a claim the merchant makes to its own visitors. With sharing on, Catalog hands that claim to agents that "might use the list price to advertise strikethrough pricing" (Shopify Help Center, n.d.-c). The claim can then appear on a surface the merchant does not run. The Help Center names one condition for sending it, that the compare-at price is higher than the current price (Shopify Help Center, n.d.-a). Shopify's developer changelog adds that stores which hide compare-at prices from EEA buyers do not send them to those buyers (Shopify, 2026e).

Federal guidance on reference prices is old and specific. The Federal Trade Commission's Guides Against Deceptive Pricing carry a source note dated November 8, 1967. The timeline on the electronic federal code finds no changes to them after January 3, 2017 (Guides Against Deceptive Pricing, 2026). Section 233.1 starts from a former price that is "the actual, bona fide price at which the article was offered to the public on a regular basis for a reasonably substantial period of time." Such a price, the guides say, "provides a legitimate basis for the advertising of a price comparison." The same section lists, among fictitious comparisons, "a price that was not openly offered to the public, or that was not maintained for a reasonable length of time, but was immediately reduced."

The guides place the care with the seller: "the advertiser should make certain that the former price is not a fictitious one" (Guides Against Deceptive Pricing, 2026). Section 233.3 applies the same logic to a manufacturer's list price used as the comparison. When a list price sits significantly above the highest price at which substantial sales occur in the trade area, the guides warn of "a clear and serious danger of the consumer being misled." Shopify itself switches compare-at sharing off by default for stores that use compare-at for a manufacturer's suggested retail price (MSRP) or another reference price (Shopify, 2026e). How the guides apply when an outside agent displays the comparison is a question this piece leaves to counsel. The compare-at price itself comes from the merchant's own product data.

The author's playbook for AI website chat held that for a small business, "the bot's most dangerous answers concern price, refunds, terms, and availability" (Puglisi, 2026e). Catalog now carries price and availability to agents in structured form, and the discovery files every Shopify store serves include the store's policies (Shopify Help Center, n.d.-c). The same care therefore applies when an outside agent, rather than the store's own bot, does the answering.

The setting is the kind of decision Checkpoint-Based Governance addresses, a framework that entered the author's public record in September 2025 (Puglisi, 2025). The framework's current page states that it "puts a named human at defined points in the work, gives that person binding authority over what happens next, and leaves a record of what they decided and why" (Puglisi, 2026c). At that point, "The human governor accepts, modifies, or rejects" (Puglisi, 2026c). The same page sets the condition: "The

checkpoint is where a named human is required to be present, documented, and accountable" (Puglisi, 2026c). The setting is only the mechanism, and the checkpoint is the named person's recorded decision about it. Shopify gives that person two positions for the control. With Compare-at price chosen, Catalog sends the comparison. With Price chosen, the Help Center states that Catalog "uses each product's current price as the list price" (Shopify Help Center, n.d.-c).

The fact is that compare-at prices now reach agents when the setting is on (Shopify, 2026f). The federal guides treat a former price as a legitimate basis for comparison when it was a bona fide price offered for a reasonably substantial period (Guides Against Deceptive Pricing, 2026). The tactic is to audit compare-at prices on the chosen products and remove any the store cannot support with its own price history. A named person then makes the list price choice on purpose. The measure is zero compare-at prices sent to agents without dated evidence that the higher price was openly offered. Manufacturer list prices and other reference prices are reviewed separately. A second measure is a written record of that choice that names who made it and why.

What does an AI shopping agent read before it recommends a product?

The search preview tool in Sales channels > Agentic shows a Listing quality indicator for each top-ranking product it returns (Shopify Help Center, n.d.-b). Listing insights break that score into five signals: description completeness, image coverage, product reviews, variant and option completeness, and shop policy completeness. Reviews count only when "verified by trusted sources." The variant check also flags option names with "acronyms or numbers that are hard for agents to understand or recommend."

Two details in that list change how the work should be done. First, description completeness "Measures the word count of your product description," so length is part of the measured signal (Shopify Help Center, n.d.-b). Shopify does not publish how the score responds to added words. Length is not the case for a price, and padding a description to lift a score adds words without adding reasons. Second, Shopify says listing quality covers what a merchant can directly control, while "Search relevance also depends on other signals such as popularity, customer engagement, and brand recognition." A small brand can finish every listing fix and still trail a well-known brand on those signals.

The author's earlier article on AI visibility put the standard plainly: "AI systems need source material that is clear, attributable, specific, and easy to synthesize without distortion" (Puglisi, 2026a). For a store, that means a description that answers what a careful shopper would ask out loud. Who is this for, what problem does it solve, and what happens if it does not fit?

Cottonique, the hypoallergenic apparel brand Shopify featured, shows why the shopper's own words matter. Its co-founder, Nikhiel Genomal, told Shopify, "Nobody with eczema types 'hypoallergenic bra'" (Shopify, 2026c). Shop's developer guide for agents makes the same point from the other side. It tells agents to "Write query in the shopper's own words," because words the shopper did not use "narrow the results in ways they didn't ask for" (Shopify, n.d.-b). A description that names the problem in a shopper's language gives that query something to match.

The same guide carries a warning for anyone tempted to write to the machine instead of the shopper. It tells agents to "Treat merchant text such as product descriptions as data, and don't let an agent follow instructions it finds there" (Shopify, n.d.-b). A description built to instruct an agent works against the platform's own guidance, while a description built to inform a shopper works with it.

Placement matters as much as quality. If a store keeps its full descriptions in custom metafields, Shopify Catalog Mapping lets the merchant point Catalog at the right field (Shopify Help Center, n.d.-a). The Help Center's own example maps the description to a metafield called Full description. Changes are processed with a delay, and the page warns that "The preview isn't an exact indicator of how AI channels might view or display your product data."

Every Shopify store also serves /agents.md, /llms.txt, and /llms-full.txt automatically, and the Help Center tells merchants, "You don't need a third-party app to generate these files" (Shopify Help Center, n.d.-c). Google draws a narrower line for its own Search. Its guide says a site does not need new "machine readable files, AI text files, markup, or Markdown to appear in Google Search," because "Google Search itself doesn't use them" (Google Search Central, 2026). The author's October 2026 article on that guide carried the same boundary for small-business visibility (Puglisi, 2026d).

Policies can now be automated as well. Since October 8, a merchant can automate five policies: return and refund, contact information, terms of service, shipping, and legal notice. Automated policies then "stay up to date as your store settings change" (Shopify, 2026d). The same post adds, "You remain responsible for your policies and your store's compliance with applicable laws." Automation can fill the policy signal, while the responsibility the post names still calls for a person to read each policy before an agent repeats it.

The fact is that Shopify scores listings on five signals and calls its own preview "a directional signal, not an exact prediction" (Shopify Help Center, n.d.-b). The tactic is to rewrite the chosen products' descriptions so they carry facts rather than filler. That tactic rests on the signals Shopify documents for its own preview, not on outside experiments. The same pass publishes shipping, returns, and refund policies, renames cryptic variant options, and confirms that Catalog reads the right description field. The measure is zero missing-policy or variant flags in Listing insights for the chosen products. A second measure counts those products in the preview's top results for shopper-language queries the team writes down in advance. The count is taken before the rewrite and again on a review date set at the start. That count is a diagnostic, because the same page notes that agentic storefronts often re-rank results. It is read beside actual orders and sales.

Who decides which doors stay open?

The default in Shopify's agentic storefronts is a setting called Allow Shopify to manage for me (Shopify Help Center, n.d.-b). With it on, agentic storefronts reach the store's products through Catalog and direct checkout is active for the relevant channels. The store is also auto-enrolled in new agentic storefronts. Turning the setting off lets a merchant manage Catalog access, auto-enrollment, and direct

checkout channel by channel. Turning it back on "automatically activates all agentic storefronts settings, including direct checkout."

The fine print shapes what any decision can achieve. Turning off Catalog access "can take up to 7 days" to take effect, and direct checkout switches off along with it (Shopify Help Center, n.d.-b). Products can still be found through web crawling and indexing. Shop, the Shopify-owned app, has no Catalog off switch in the Agentic settings, and its products are managed from the Shop sales channel instead (Shopify Help Center, n.d.-b). Google AI Mode and Gemini get products through the Google & YouTube sales channel, and Meta surfaces through the Facebook and Instagram by Meta sales channel (Shopify Help Center, n.d.-c).

Shopify points a merchant who wants a product hidden from every agentic storefront, Shop included, to Unlisted status or the seo.hidden metafield (Shopify Help Center, n.d.-b). Either choice also keeps the product out of search engines such as Google, the same page warns. Blocking AI crawlers in robots.txt "affects only open-web discoverability" and does not stop Catalog from sending product data to the storefronts a store has activated (Shopify Help Center, n.d.-c). A separate group called Other channels covers "Independent developers and startups that are building niche or experimental shopping agents" (Shopify Help Center, n.d.-b).

Direct checkout differs by channel. The Help Center lists direct checkout for Google AI Mode and Gemini, Microsoft Copilot, and Meta surfaces such as Muse (Shopify Help Center, n.d.-b). The Catalog page adds that direct checkout "is also rolling out to Google users and might not yet be available in your store" (Shopify Help Center, n.d.-c). ChatGPT receives Catalog product data but does not appear on that list (Shopify Help Center, n.d.-b, n.d.-c).

In March 2026, an OpenAI spokesperson told Digital Commerce 360 that "Instant Checkout is moving to Apps, where purchases can happen more seamlessly" (Warmoth, 2026). OpenAI's own post on March 24 said the initial version of Instant Checkout "did not offer the level of flexibility that we aspire to provide" (OpenAI, 2026a). The company is instead "allowing merchants to use their own checkout experiences while we focus our efforts on product discovery." In the same post, Shopify's vice president of product, Mani Fazeli, said buyers "can complete purchases on the merchant's online store through an in-app browser." Shop MCP adds a further route, since a developer's MCP-compatible agent can now open a checkout at the store that sells the product (Shopify, 2026g).

OpenAI has moved twice since then. On November 10, 2026, it opened Businesses in ChatGPT in beta (OpenAI, 2026b). A company can claim its business there and, once verified, "update key details and add self-service product feeds with pricing and availability." The post adds that Shopify merchants "can connect their store from Business Hub or the Shopify admin to streamline verification." Their eligible products, it says, "are already integrated into ChatGPT through Shopify Catalog."

On November 13, OpenAI introduced ChatGPT Wallet (OpenAI, 2026c). Its post calls the product "the first wallet that lets agents pay without access to users' actual card details and that requires a security check for every purchase." OpenAI built it with American Express, Mastercard, and Visa. At launch it reaches ChatGPT Pro subscribers aged 18 and older in the United States who hold eligible cards issued

there, through an OpenAI product called dots. The buyer reviews the items, the total, and the selected card, then approves each purchase. The post does not mention Shopify, and it sets one condition any merchant can check. To accept payment from the Wallet, it says, "websites need to ensure agents are not inadvertently blocked." A store whose firewall or content delivery network blocks AI agents may therefore turn away a paying shopper's agent.

Shop MCP's own rules keep the shopper in the purchase decision. Its checkout guide tells agents to show the merchant, the items, the current total, and the offered payment methods, then to "have them choose one. Never pick a default for them" (Shopify, n.d.-a). It adds that "Only a result with status completed and an order confirms the purchase." In the Shop MCP flow, the code carries the search, the cart, and the checkout, and the shopper makes the purchase decision.

The merchant holds a decision of its own, and the default fills it unless someone takes it. Leaving Allow Shopify to manage for me switched on hands enrollment in every future agentic storefront to the platform's rules (Shopify Help Center, n.d.-b). Many small merchants may reasonably accept that default, and a named person should still look at it, decide, and record the call. The framework's current page explains why: "Approving an interaction in a process is not the same as accepting an output" (Puglisi, 2026c).

Direct checkout is the same kind of call. An order placed through Shop MCP appears in the admin like any other order, and a signed-in shopper can now pay "without leaving your agent" (Shopify, 2026g). The buyer may never have opened the store's own pages. The store's return policy, shipping times, and service process then carry the whole relationship. They have to hold up for a customer who knows the store only through an agent's summary.

The fact is that the default auto-enrolls a store in future agentic storefronts with direct checkout active where supported (Shopify Help Center, n.d.-b). Because turning off Catalog access can take up to seven days, this decision comes right after the settings record and ahead of the listing work. The tactic is a short checkpoint review of the Agentic settings, where a named owner records a decision and the reason for each one. The settings in scope are auto-enrollment, Catalog access by channel, and direct checkout. The record also notes the Google & YouTube and Facebook and Instagram by Meta sales channels, which send products to agents outside the Agentic settings (Shopify Help Center, n.d.-c). Three routes now reach a store: Catalog for discovery, Shop MCP for agent checkout, and ChatGPT Wallet for agent payment on the store's own site. For each route, the record names what controls it and where its orders appear. The same owner reviews the Knowledge base facts before publishing them (Shopify, 2026h). The measure is a logged decision for each setting by November 20, 2026, seven days before Black Friday. Any change to Catalog access will then have taken effect by the holiday. A second review goes on the calendar for the week after the holiday window closes on December 31.

How will a merchant know whether any of this worked?

The Agentic section reports four measures for each AI channel: sales, orders, online store sessions, and online store conversion (Shopify Help Center, n.d.-b). Shopify calculates conversion as "sessions that

completed checkout / sessions," and it counts sessions as "visits to your online store from an AI channel." Sales data "is an aggregate of referral-based sales and sales that happen through direct checkout," and direct checkout lets customers buy "directly in their conversation in an AI channel." A purchase made inside the conversation therefore need not add a session, so sales can rise while sessions and conversion stay flat. Orders and sessions belong next to sales in every reading, and the Help Center notes that these analytics are not yet available for headless stores. Shopify also includes Shop app sessions, referrals, and sales in this reporting. Results are therefore read channel by channel before total agentic sales are credited to outside AI assistants (Shopify Help Center, n.d.-b).

Outside counts carry blind spots of their own. Similarweb notes that visits arriving without referrer data land in direct traffic, so its referral figures "are a floor, not a ceiling" (Zohar Basteker, 2026). Bero's senior director of ecommerce solutions and growth told Digital Commerce 360 that session reporting cannot yet say exactly how much traffic comes from AI (Haleem, 2026).

In its current statement, Factics reduces the method to three questions: "What do we know?", "What do we do with what we know?", and "What happened because we did it?" (Puglisi, 2026b). The third question needs a target written before the results come in, so that the result cannot rewrite the goal after the fact.

The fact is that Shopify reports agentic sales, orders, sessions, and conversion for each channel, with sales combining referral and direct checkout (Shopify Help Center, n.d.-b). The tactic is to record a baseline for those four measures now, beside the whole-store conversion rate, and to write a target for each before Black Friday. The owner also logs the hours the work takes. The measure is agentic orders and sales against their targets, from the day the targets are written through December 31, 2026. That date closes the holiday season as Adobe measures it (Adobe, 2026). Referral conversion is read separately against whole-store conversion, because a direct-checkout order adds no session. Sales set against the logged hours show whether the effort paid for itself.

Where should a small merchant start this week?

The steps below are offered as a possible path or starting point, not as a promise of results. Start with the settings, because a change to Catalog access can take up to seven days to take effect. Make one deliberate call on compare-at sharing and one on auto-enrollment, with the owner's name beside each. Then pick the products that carry most of the season's revenue. For each one, write down what a careful shopper would want to know before paying full price. Put those answers in the field Catalog reads, publish the store's policies, review the Knowledge base facts, and rename any variant a stranger could not decode. The baseline numbers, the targets, and the settings calls go in the merchant's own record. The descriptions, checkable facts, and published policies are what an agent can read. When an agent is asked in December whether one of those products is worth the price, they are the case it has to work with.

What does the evidence against agentic shopping say?

The channel is still small for most stores. Similarweb reports that AI referrals "remain a low single-digit percentage of most sites' total traffic" (Zohar Basteker, 2026). That holds even though AI referral visits grew 117.4 percent year over year in its dataset. Similarweb sells an AI Traffic Tracker, so it has an interest in the channel's growth (Zohar Basteker, 2026). Retailer accounts are mixed. In Digital Commerce 360's reporting, trends among leading retailers put AI at about 5 percent of referrals in the final months of 2026 (Haleem, 2026). Bero put the share of its traffic from AI referrals at "definitely less than five," and called it "not really growing quickly." Edible Brands' chief digital officer put AI-referred conversion above 5 percent and added, "That's lower than our site average," which runs above 10 percent. La Joya Jewelry's founder gave the opposite account, saying conversion rates for visitors from AI are double digits higher than for other visitors. As an example, he said that if the store's conversion rate is normally 3 percent, it rises to more than 13 percent for AI visitors.

Adobe's wider data runs stronger than the weaker retailer accounts. Adobe published an analysis of U.S. retail sites on October 29, 2026 (Pandya, 2026). It reported that AI traffic in September "converted 90% better" than non-AI traffic and "generated 83% higher revenue per visit." Adobe sells analytics products to retailers, so this finding is vendor-reported as well. An aggregate across many retailers is also a different population from one store, and both readings stand. The aggregate ran strong in September, while individual retailers reported results on either side of their own site averages. These sources also measure different things. Shopify's survey records what shoppers say they plan to do, Adobe measures visits and purchases from AI referrals, and Shopify's merchant reports also count direct checkout. They cannot be combined into one conversion estimate.

Trust remains the narrowest point. Checkout.com reported in June 2026 that 24 percent of consumers say they will never delegate purchases to AI (Checkout.com, 2026). In the same research, 27 percent trust no organization to operate an AI shopping agent. Merchants in the UK and US told the company that only 3 percent of transactions involve AI agents. The release does not state its sample size, and Checkout.com describes itself as a digital payments company, so its findings are vendor-reported as well. Shopify's own figure, 29 percent of shoppers trusting AI to buy on their behalf, points the same way (Shopify, 2026a).

The equalizer story has a challenger. Shopify's post argues that "Brands built around specific customer needs are advantaged" in agentic search (Shopify, 2026c). Similarweb's reading of its 2026 data is that "AI favors large, multi-category platforms over smaller, vertical-specific sites" when answering shopping queries (Zohar Basteker, 2026). Both readings can hold in different categories, and no store should assume the equalizer effect applies to its own.

A research paper offers a third reading, drawn from a sandbox rather than a live store. Allouah and colleagues first posted it to arXiv in August 2025 and revised it that December. A four-page short paper followed in the proceedings of the ACM Web Conference 2026, published April 12, 2026 (Allouah et al., 2026). The short paper describes "a fully programmable mock marketplace" and reports AI agents "sometimes concentrating on a few products, raising competition questions." The longer arXiv version

remains a preprint. It reports that agents "can exhibit choice homogeneity, often concentrating demand on a few 'modal' products while ignoring others entirely," and that "model updates can drastically reshuffle market shares" (Allouah et al., 2025). The preprint also reports that a seller-side agent making "simple, query-conditional description tweaks can drive significant gains in market share." Those last two findings appear in the preprint and not in the short paper's abstract. The authors have released the sandbox code and datasets for replication, and two of the five list the company MyCustomAI as their affiliation (Allouah et al., 2025). None of the tactics in this piece depends on its results.

The survey measures intent and comfort, not purchases, and a shopper who says they are likely to use AI for a task has stated a plan. Frequent AI use for shopping also skews young and higher-income in Shopify's data. Among 18- to 24-year-olds, 25 percent often or almost always use AI when shopping, against 4 percent of shoppers over 65 (Shopify, 2026a). Frequent use runs from 12 percent of shoppers with household income under $19,000 to 28 percent of those above $95,000. The 71 percent figure describes shoppers who almost always use AI tools, not all AI users.

The brand results come from brands a vendor selected for its story. They include Stanley 1913's five times as many U.S. AI-referred orders as the same window last year and Cottonique's 276 percent growth in AI-referred sales (Shopify, 2026c). Shopify's Spring '26 Edition also claims that AI searches powered by Catalog "convert at 2x the rate of those using scraped data" (Shopify, 2026b). That page states no method behind the figure.

A clear test shows when this work is a poor use of a small merchant's hours. Suppose a store completes the listing work above, and its agentic orders and sales stay below the targets the team wrote down before Black Friday. Suppose also that referral conversion stays below whole-store conversion and the sales do not repay the logged hours. The work then drops to a low priority for that store in 2027. The hours then belong to email, repeat customers, and the store's own site. Shopify's own survey points there too, finding that "businesses plan to prioritize keeping the customers they already have" (Shopify, 2026a).

What to watch next

Adobe expects AI traffic to U.S. retail sites, measured by shoppers clicking a link, to rise 130 percent year over year this holiday season (Adobe, 2026). Its readouts after December 31 will show how close that forecast came. Adobe's October analysis projects AI traffic up 95 percent year over year on Black Friday itself (Pandya, 2026). Shop MCP launched with no webhooks, and its post tells developers to "poll get_checkout for checkout status" instead (Shopify, 2026g). Event support is one thing to watch, and merchants face a different question. None of the three Help Center pages read for this piece mentions Shop MCP (Shopify Help Center, n.d.-a, n.d.-b, n.d.-c). How Shopify files those orders in the Agentic settings and reports therefore remains unclear.

OpenAI is the other checkout story for Shopify merchants to follow. Its March shift toward merchants' own checkout experiences was followed in November by a wallet for agents (Warmoth, 2026; OpenAI, 2026a, 2026c). OpenAI's post says it expects users, over time, to "delegate purchases entirely to their

agent" within limits they set. Whether Wallet orders reach Shopify stores, and where they would appear in a store's reports, is not documented in the sources read for this piece. Developers face a dated change as well. On November 11, 2026, Shopify deprecated the Storefront MCP tool that searched shop policies and FAQs. It told developers to migrate to the UCP MCP ask_business tool by January 15, 2027 (Shopify, 2026i). The compare-at change will show whether agents display strikethrough prices widely, and whether shoppers treat an agent's markdown claim the way they treat a store's. The Allouah preprint's finding that "model updates can drastically reshuffle market shares" came from a sandbox (Allouah et al., 2025). If live shopping agents behave the same way, a ranking a store earns this December may not survive the next model release.

References

Adobe. (2026, September 28). *Adobe: U.S. holiday shopping season to hit record $275.1 billion online, rising 6.7% YoY* [Press release]. https://news.adobe.com/news/downloads/pdfs/2026/09/adi-holiday-forecast-release-9-28.pdf

Allouah, A., Besbes, O., Figueroa, J. D., Kanoria, Y., & Kumar, A. (2025). *What is your AI agent buying? Evaluation, biases, model dependence, & emerging implications for agentic e-commerce* (arXiv:2508.02630) [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2508.02630

Allouah, A., Besbes, O., Figueroa, J. D., Kanoria, Y., & Kumar, A. (2026). What is your AI agent buying? Evaluation, biases, model dependence, & emerging implications of agentic e-commerce [Short paper]. In *WWW '26: Proceedings of the ACM Web Conference 2026* (pp. 8697-8700). Association for Computing Machinery. https://doi.org/10.1145/3774904.3792943

Checkout.com. (2026, June 9). *Consumer demand for AI shopping is forming fast but trust for agentic commerce is still catching up* [Press release]. https://www.checkout.com/newsroom/consumer-demand-for-ai-shopping-is-forming-fast-but-trust-for-agentic-commerce-is-still-catching-up

Google Search Central. (2026, July 10). *Optimizing your website for generative AI features on Google Search*. Google for Developers. https://developers.google.com/search/docs/fundamentals/ai-optimization-guide

Guides Against Deceptive Pricing, 16 C.F.R. pt. 233 (2026). https://www.ecfr.gov/current/title-16/chapter-I/subchapter-B/part-233

Haleem, A. (2026, September 16). *How AI referrals are impacting traffic at 4 online retailers*. Digital Commerce 360. https://www.digitalcommerce360.com/2026/09/16/how-ai-referrals-impact-traffic-4-online-retailers/

OpenAI. (2026a, March 24). *Powering product discovery in ChatGPT*. https://openai.com/index/powering-product-discovery-in-chatgpt/

OpenAI. (2026b, November 10). *Introducing Businesses in ChatGPT*. https://openai.com/index/introducing-businesses-chatgpt/

OpenAI. (2026c, November 13). *Introducing ChatGPT Wallet*. https://openai.com/index/introducing-chatgpt-wallet/

Pandya, V. (2026, October 29). *Adobe report: AI traffic now generates 83% higher revenue per visit for U.S. retailers heading into the holiday shopping season*. Adobe for Business Blog. https://business.adobe.com/blog/ai-traffic-higher-revenue-per-visit-retail-holiday

Puglisi, B. C. (2012, November 27). *Digital Factics: Twitter*. Digital Media Press. https://www.magcloud.com/browse/issue/471388

Puglisi, B. C. (2025, September 23). *Checkpoint-Based Governance: An implementation framework for accountable human-AI collaboration (v2 drafting)*. basilpuglisi.com. https://basilpuglisi.com/checkpoint-based-governance-an-implementation-framework-for-accountable-human-ai-collaboration-v2-drafting/

Sources in this work: Bradley (2025); Congruity 360 (2025); European Parliament and Council (2024); International Telecommunication Union (2025); ISACA (2025); Lu et al. (2019); Lumenalta (2025); McKinsey & Company (2025); National Institute of Standards and Technology (2023); Nexastack (2025);

OECD (2025); Parasuraman and Manzey (2010); Precisely (2025); Splunk (2025); Stanford Institute for Human-Centered Artificial Intelligence (2024); Strobes Security (2025); Superblocks (2025); Visa (2025).

Puglisi, B. C. (2026a, June 25). *AI visibility: Be the human voice AI can find*. basilpuglisi.com. https://basilpuglisi.com/ai-visibility-be-the-human-voice-ai-can-find/

Sources in this work: Aggarwal et al. (2024); Atari et al. (2023); Chapekis and Lieb (2025); Google Search Central (2025, 2026); Jaźwińska and Chandrasekar (2025); Liu, Zhang, and Liang (2023); Microsoft Bing Blogs (2026); National Institute of Standards and Technology (2023); OpenAI (n.d.); Search Engine Land (2025).

Puglisi, B. C. (2026b, September 2). *Factics: The method behind valuable content and trusted AI*. basilpuglisi.com. https://basilpuglisi.com/factics/

Puglisi, B. C. (2026c, September 11). *Checkpoint-Based Governance: Presence is not authority*. basilpuglisi.com. https://basilpuglisi.com/cbg/

Puglisi, B. C. (2026d, October 2). *Google's own guide makes GEO hacks optional for Google Search: Unique experience still matters more for SMB visibility*. basilpuglisi.com. https://basilpuglisi.com/google-search-geo-hacks/

Sources in this work: Aggarwal et al. (2024); DoodleWeb (2026); Google Search Central (2026); Google Search Console Help (n.d., 2026); Linehan (2026); Martinez (2026); Mueller (2026).

Puglisi, B. C. (2026e, October 10). *Can AI website chat convert for small businesses? A 30-day Factics playbook*. basilpuglisi.com. https://basilpuglisi.com/ai-website-chat-small-business/

Sources in this work: Aarons and Sarkar (2026); American Bar Association (2024); Bodnar (2024); Brynjolfsson, Li, and Raymond (2025); Cal. Bus. & Prof. Code § 17941; Colorado General Assembly (2026); Doerer (2026); Echt-Wilson (2019); European Commission (n.d.); Grundy, Breaux, and Khatiwoda (2026); Habel, Zawadzki, and Freise (2026); Luo, Tong, Fang, and Qu (2019); Mahrouyan (2026); Mozafari, Weiger, and Hammerschmidt (2021); NFIB (2025); Office of the Governor of California (2026); Salesforce (2026); Santiago (2026); Utah Code § 13-77-103; Zhang and Narayandas (forthcoming).

Shopify. (n.d.-a). *Carts and checkout*. shop.dev. https://shop.dev/docs/checkout

Shopify. (n.d.-b). *Find products*. shop.dev. https://shop.dev/docs/search

Shopify. (n.d.-c). *Shop MCP*. shop.dev. https://shop.dev/docs

Shopify. (2026a). *The holiday shopping season meets the agentic era: Shopify 2026 global holiday retail report*. https://www.shopifyholidayretail.com/

Shopify. (2026b, June 17). *Selling everything, everywhere, all at once: The Spring '26 Edition*. Shopify News. https://www.shopify.com/news/spring-26-edition-merchant

Shopify. (2026c, October 6). *Welcome to the first holiday season of the agentic era*. Shopify News. https://www.shopify.com/news/agentic-holiday-2026

Shopify. (2026d, October 8). *Automatically generated policies help you launch faster*. Shopify Changelog. https://changelog.shopify.com/posts/automatically-generated-policies-help-you-launch-faster

Shopify. (2026e, October 9). *Catalog API now returns compare-at prices as list_price*. Shopify Developer Changelog. https://shopify.dev/changelog/posts/catalog-api-now-returns-compare-at-prices-as-list_price

Shopify. (2026f, October 9). *Compare-at prices now shared to agentic channels*. Shopify Changelog. https://changelog.shopify.com/posts/compare-at-prices-now-displayed-to-agentic-channels

Shopify. (2026g, October 13). *Add shopping and Shop Pay checkout to any agent with Shop MCP*. Shopify Developer Changelog. https://shopify.dev/changelog/posts/add-shopping-and-shop-pay-checkout-to-any-agent-with-shop-mcp

Shopify. (2026h, October 28). *Exponential: Shopify's Fall '26 Edition*. Shopify News. https://www.shopify.com/news/fall-26-edition-merchant

Shopify. (2026i, November 11). *Storefront MCP policies and FAQs tool is being deprecated in favor of UCP MCP Ask tool*. Shopify Developer Changelog. https://shopify.dev/changelog/posts/storefront-mcp-policies-and-faqs-tool-is-being-deprecated-in-favor-of-ucp-mcp-ask-tool

Shopify Help Center. (n.d.-a). *Customizing your product data for Shopify Catalog*. Shopify. https://help.shopify.com/en/manual/shopify-catalog/mapping

Shopify Help Center. (n.d.-b). *Managing agentic storefronts*. Shopify. https://help.shopify.com/en/manual/online-sales-channels/agentic-storefronts/agentic-home

Shopify Help Center. (n.d.-c). *Shopify Catalog and product discovery for agentic storefronts*. Shopify. https://help.shopify.com/en/manual/online-sales-channels/agentic-storefronts/products

Warmoth, B. (2026, March 6). *OpenAI shifts checkout plans in its agentic commerce strategy*. Digital Commerce 360. https://www.digitalcommerce360.com/2026/03/06/openai-shifts-checkout-plans-agentic-commerce-strategy/

Zohar Basteker, M. (2026, September 3). *AI referral traffic by industry: What 770 million monthly visits reveal* [Page updated September 17, 2026]. Similarweb. https://aisearch.similarweb.com/blog/ai-referral-traffic-by-industry/

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

Basil C. Puglisi, MPA
A Human-AI Collaboration

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.

Suggested citation: Puglisi, B. C. (2026). Can an AI Shopping Agent Explain Why Your Product Is Worth the Price? basilpuglisi.com.