

Can AI Website Chat Convert for Small Businesses? A 30-Day Factics Playbook

HubSpot's chief marketing officer, Kipp Bodnar, traces one of the company's early turning points to a plain discovery. Prospects preferred to chat with the company about buying the product rather than fill out forms or send emails. Chat then became one of the company's most valuable channels for customer satisfaction and acquisitions (Bodnar, 2024). A small business faces the same preference with far fewer people available to answer it. AI website chat is one way to close that staffing gap. For an owner with one to five people, the open question is whether it holds up at their size and traffic, under the rules that now apply.

The way to answer that question is older than the technology. A method that pairs every fact with a tactic and every tactic with a measurable outcome, which the author calls [Factics](#), was first stated publicly in October 2012. It went into print in *Digital Factics: Twitter* that November. From 2013 onward, goals, outcomes, and KPIs turned the pairing of fact and tactic into a measured cycle (Puglisi, 2012, 2026). The principle that governs this playbook comes from that method: every fact must lead to a tactic, and every tactic must leave evidence.

This playbook applies that principle in the order a small business needs it. It starts with what HubSpot actually reported and what independent research adds, then asks whether the business has a problem big enough for a bot. From there it sets out what the law requires for disclosure and what the research says disclosure costs. It then covers accuracy, liability, privacy, and what retail-scale agents show. It closes with a 30-day plan that ends in a decision, and one of the valid decisions is not to deploy. One limit runs through all of it. Every outcome figure in the vendor sources comes from a company reporting on its own product. The research reviewed for this playbook also located no independent field study of autonomous generative chat in businesses of one to five people.

What did HubSpot actually report?

HubSpot integrated AI into its website chat and named three KPIs before scaling: conversion rate, value per chat, and customer satisfaction (CSAT). The team placed the first chatbot on high-traffic, lower-risk knowledge base pages, where customers ask practical, straightforward questions. HubSpot reported a 43% increase in chat conversion rates by the end of the experiment and an improvement of more than 50% in value per chat. It also reported that CSAT scores matched human-led interactions (Bodnar, 2024).

The early phase matters as much as the ending. CSAT fell after launch while the new model still needed training, and HubSpot treated the dip as expected while the model learned. A team member then annotated chat transcripts by hand, editing responses so they were more accurate and better matched to what users asked. Emmy Jonassen, HubSpot's VP of Marketing for Demand Generation, put the trade plainly. "The annotation piece is really one of the most important pieces through all of this," she said, "but it's also the most time-intensive" (Bodnar, 2024). Her advice to other teams favored speed over polish: "Get your AI experiment to a good enough place, get it out in the wild, and then iterate based on real-world feedback."

The post leaves out what a reader would need to judge the result independently. It does not publish the chat or conversion denominators, the absolute rates before and after, the length of the test, the control design, or the CSAT sample sizes. It also does not separate the effect of annotation from the other

changes made during the experiment. Read that way, HubSpot's account is an operating example from a well-staffed software company, and a small business is better served treating it as one. Its next test, on the pricing page with touchless purchases in some cases, was still running when the post was last updated.

An older HubSpot experiment sharpens the routing lesson, with a caution about its age. Ari Echt-Wilson's account, first published in July 2019 and updated in 2025, describes a rules-based routing bot built before generative AI chat existed. HubSpot's sales team had found that 15% of chat questions on the website were technical support requests. The company built a bot that sent customers to support resources and passed product and pricing questions to sales. Rolled out across primary pages, that bot raised sales team efficiency by 80%, which HubSpot defined as meetings booked divided by chats handled (Echt-Wilson, 2019).

The same team then noticed that the support-shaped bot was built for the 30% of its audience who were customers, while visitors and leads made up about 70%. In an A/B test that ran for a month, HubSpot sent half the traffic on one product page to a cloned page. The clone carried a visitor path that offered free CRM signup alongside sales chat. In the visitor variant, sales team efficiency rose by almost 70% and CRM signups rose 7%. Fewer chats reached the sales team, yet meetings booked held at a similar level. The customer bot also performed 52% better when shown only to customers than when shown to all visitors (Echt-Wilson, 2019). The post does not name the metric behind that figure.

The routing lesson outlives the technology that produced it, and the test design teaches as much as the result. A split of traffic for a fixed period is a stronger measure than a comparison of this month against last month. A rules-based routing bot is also a fair baseline for any generative system to beat, since the 2019 gains came without a language model at all.

HubSpot reported that annotated, disclosed AI chat lifted conversion and chat value while matching human CSAT (Bodnar, 2024). Its 2019 test showed that audience-specific paths raised sales efficiency and signups (Echt-Wilson, 2019). The tactic for a small team is to treat both as hypotheses to test on its own pages, with annotation budgeted as labor and routing tested before generative answers. The KPI is the same set HubSpot named, conversion, value per chat, and CSAT, measured against the team's own baseline.

What does research outside the vendors add?

Independent research on AI in customer service exists, and it measures something narrower than the vendor stories. Brynjolfsson, Li, and Raymond (2025) examined the staggered rollout of a generative AI conversational assistant across 5,172 customer support agents, in a study published in *The Quarterly Journal of Economics*. They reported that access to the assistant raised productivity, measured as issues resolved per hour, by 15% on average. Less experienced and lower-skilled workers improved both the speed and quality of their output, and customers were more polite and less likely to ask to speak to a manager.

Zhang and Narayandas (forthcoming) ran a randomized field experiment at a meal delivery company, where some customer service agents received AI-generated reply suggestions. They reported that the assisted agents responded faster, engaged customers more deeply, and saw greater improvement in customer sentiment, with the largest gains for less experienced agents. They also reported a warning that matters for any small business planning a bot. Some customers first dealt with an automated

chatbot that failed to understand them. For those customers, the AI-assisted agent who followed made sentiment worse, because the agent's quick replies made them think they were still talking to a bot.

Both studies measure AI assisting human agents. Neither measures an autonomous bot answering a small business's website visitors, which is the setup this playbook covers. The broader return picture is also mixed. CX Dive reported a Gartner analysis of 432 AI use cases in customer service. Only one quarter produced a return on investment, another quarter delivered negative returns, 11% broke even, and 42% had unclear returns (Doerer, 2026). In the same report, Julie Geller of Info-Tech Research Group warned that "containment is also too often mistaken for success." She named a simpler test: "did the customer get what they needed, with less effort?"

Independent studies report productivity and sentiment gains when AI assists human agents, along with a failure mode when a weak bot hands off to a human (Brynjolfsson et al., 2025; Zhang & Narayandas, forthcoming). Across hundreds of service use cases, positive and negative returns split evenly (Doerer, 2026). The tactic is to write the KPI sheet before launch, so the business can show its return when the pilot ends. An unclear return, the outcome for 42% of the cases Gartner reviewed, is what the sheet is built to avoid. The KPI is a written return measure for the pilot, covering conversion, resolution, and owner time, completed before the bot goes live.

Is the problem big enough for a bot?

Government data places the stated audience at the early edge of adoption. The U.S. Census Bureau reported that business AI use hovered between 17% and 20% from December 2025 to May 2026. Use rose among firms with at least 20 employees but did not change significantly among firms with fewer than 20. Less than 20% of firms with four or fewer employees reported using AI (Grundy et al., 2026). NFIB's 2025 survey of small employers found 24% using AI, including 21% of firms with single-digit employee counts against 48% of firms with 50 or more employees. Only 9% were using or planning to use AI for customer service (NFIB, 2025).

Adoption also differs from results. The Federal Reserve Bank of New York analyzed its 2025 Small Business Credit Survey and found that 46% of firms with at least one employee were using AI tools. The authors also wrote that "only 31 percent of survey respondents reported increased sales from their firm's use of AI" (Aarons & Sarkar, 2026). Those figures cover AI use in general rather than website chat, so they do not settle whether a chatbot pays for itself. They do make one point plainly: using AI and selling more because of it are separate outcomes.

That gap makes the first decision a question about demand rather than software. A business with 20 relevant website inquiries a month, few repeated questions, and no CRM has a different problem from one that fields the same five questions every day. Before any bot goes live, the owner can count a month of inquiries and sort the repeated ones. The next step is to estimate the minutes each takes and mark how many can be answered from approved information. Cheaper fixes then get a fair hearing, including a better FAQ, a booking form, a rules-based routing bot, or faster human follow-up. A decision not to deploy is a valid result, because the tactic still left evidence.

Small-firm AI use trails larger firms in government data, and reported sales gains reach a minority of users (Grundy et al., 2026; NFIB, 2025; Aarons & Sarkar, 2026). The tactic is a 30-day inquiry count before any purchase, with each repeated question tagged as answerable or not from approved information. The KPIs are repeated inquiries per month, owner minutes per inquiry, the share answerable from approved content, and a written go, no-go, or alternative decision.

What does the law now require for AI chat disclosure?

Disclosure began as a design choice and is now, in several places, a legal floor. The rules differ by jurisdiction, by business size, and by who built the bot. The summary below reports what the cited sources state as of October 2026, and it is not legal advice.

European Union. Under Article 50 of the AI Act, the transparency obligations apply from August 2, 2026. The European Commission's guidance states that providers of AI systems that interact directly with people must design them to inform people they are dealing with AI. That notice must come "from the start of the first interaction" and be clear, distinguishable, and accessible. The exception for cases where AI use is obvious is to be read restrictively. The limited grace period to December 2, 2026 covers only the marking of AI-generated content under Article 50(2), not chatbot disclosure (European Commission, n.d.). Because the guidance assigns this duty to providers, whether a small business counts as a provider depends on how its bot is built.

California. Since July 1, 2019, Business and Professions Code section 17941 has barred using a bot online to mislead a person about its artificial identity in order to incentivize a sale. A person using a bot is not liable under the section if the bot discloses that it is a bot. That disclosure must be "clear, conspicuous, and reasonably designed" to inform the person (Cal. Bus. & Prof. Code § 17941). Governor Newsom signed AB 1609 on September 28, 2026 (Office of the Governor of California, 2026). As summarized by Mahrouyan (2026), the law takes effect January 1, 2027 and applies to businesses with more than \$500 million in annual revenue. It requires a clear and conspicuous disclosure where a reasonable person is likely to be misled into thinking a customer-service chatbot is human. It also requires a good-faith effort to connect a customer to a human within 15 minutes of a request, or to schedule a specific appointment within one business day. Almost no reader of this playbook crosses the revenue line, yet the law shows where customer-service rules are heading.

Colorado. HB 26-1263, signed May 29, 2026 and effective January 1, 2027, requires an operator of a conversational AI service to tell users that the service is AI. The official summary covers services "accessible to the general public" that primarily simulate human conversation (Colorado General Assembly, 2026). It does not settle how far that reaches into ordinary business chat.

Utah. Since May 7, 2025, Utah Code section 13-77-103 has required a supplier using generative AI in a consumer transaction to disclose the AI when a person clearly asks. Proactive disclosure is required only for high-risk interactions in regulated occupations (Utah Code § 13-77-103).

Federal. The research for this playbook located no federal statute that requires a label on commercial chatbots.

The rules described here require AI disclosure from the first interaction in the EU, under Colorado's 2027 law, and for large California businesses where a customer could be misled. They also make disclosure the safe harbor under California's 2019 bot law (European Commission, n.d.; Colorado General Assembly, 2026; Mahrouyan, 2026; Cal. Bus. & Prof. Code § 17941). The tactic is a plain AI label in the bot's first message on every page where it runs, reviewed with counsel by any business with EU visitors. The KPI is disclosure at the first message on 100% of automated chats, checked by a weekly sample of transcripts.

What does the research say disclosure costs?

The published research on disclosure and sales does not point one way, and both sides belong in front of a small business owner. Luo, Tong, Fang, and Qu (2019) reported a field experiment on sales calls with more than 6,200 customers. Undisclosed chatbots performed as well as proficient human workers and four times better than inexperienced ones in generating purchases. When the chatbot's identity was disclosed before the conversation, purchase rates fell by more than 79.7%. Customers perceived the disclosed bot as less knowledgeable and less empathetic, and the authors reported that later disclosure timing and customers' prior AI experience softened the effect.

Newer research reports a different result. Habel, Zawadzki, and Freise (2026) ran two field experiments with AI voice agents selling by phone, reaching 1,337 and 366 customers, plus an online study of 238 participants. They reported that disclosure no longer reduced conversion compared with nondisclosure, in part because customers largely recognized the AI agents whether or not they were told. The same paper reported that the AI agents still substantially underperformed human sales agents, which the authors traced to weaker sales process discipline.

Context changes the answer again. Mozafari, Weiger, and Hammerschmidt (2021) reported that chatbot disclosure lowered trust, and through trust lowered retention, for services of high importance to the customer. They found no effect on trust for low-criticality services. When the chatbot failed to resolve an issue, disclosure did not hurt retention and even raised it, working much like an explanation for the failure.

HubSpot chose early to be 100% transparent that customers in chat were speaking with an AI assistant. Bodnar's stated reasoning is that transparency builds trust, manages expectations, and can ease concerns about privacy and data use (Bodnar, 2024). That reasoning is a company's operating position, and the post reports no test of disclosed against undisclosed chat. The marketers HubSpot surveyed showed the same concern from their side. Among them, 42% said data privacy concerns had kept their teams from adopting new AI tools in the past year (Santiago, 2026).

One mitigation from the older research does not travel. Luo and colleagues found that later disclosure reduced the penalty, but the EU rule calls for notice from the first interaction, and Colorado's 2027 law and California's large-business rule point the same way. A small business serving those markets cannot buy back conversion by disclosing late. This playbook carries its own AI-assistance mark at the end for the same reason it asks readers to label their bots: the reader is owed the information either way.

The research reports a large disclosure penalty in 2019 sales calls and no penalty in 2026 voice-agent sales (Luo et al., 2019; Habel et al., 2026). It also reports effects that depend on how much the service matters to the customer (Mozafari et al., 2021). The tactic is to disclose at the first message and keep a clear human route on high-stakes pages. Any conversion cost then becomes a number to measure rather than an argument to win. The KPI is conversion on disclosed AI chats against the team's own pre-launch baseline, read alongside CSAT, so a disclosure cost, if one appears, shows up in the sheet.

How do accuracy, liability, and privacy change the setup?

A wrong answer can cost more than a missed sale. In *Moffatt v. Air Canada*, decided February 14, 2024, the British Columbia Civil Resolution Tribunal held the airline liable for negligent misrepresentation. Its website chatbot had told a customer he could apply for a bereavement fare reduction within 90 days of ticketing. The airline's linked policy page said the policy did not apply after travel. Air Canada argued

that the chatbot was a separate legal entity responsible for its own actions. The tribunal rejected that argument. It held that the company bore responsibility for all the information on its website, whether it came from a static page or a chatbot (American Bar Association, 2024).

For a small business, the bot's most dangerous answers concern price, refunds, terms, and availability. An approved reference sheet for those topics, kept current and loaded as the bot's source, limits what the bot can invent. A weekly sample of transcripts on those same topics, marked right or wrong by the owner, turns accuracy into a number. A single wrong statement of price or policy is a reason to pause the bot the same day.

Privacy belongs in the first version for the same reason disclosure does. Chat transcripts hold customer names, contact details, and sometimes account or health information. Three steps cover the basics for a team of one to five. The team writes a short rule on how long transcripts are kept, reads the chat vendor's data processing terms, and keeps sensitive customer details out of general-purpose AI tools.

The human route has to be honest about capacity. A solo owner can't promise an instant human at midnight, and a bot that implies one sets up the frustration the Zhang and Narayandas study described. A visible callback request, a booking link, or a stated response window gives the customer a real path to a person. That design also mirrors what California now requires of large businesses (Mahrouyan, 2026).

A tribunal held a company liable for its chatbot's wrong answer and rejected the claim that the bot answered for itself (American Bar Association, 2024). The tactic is an approved reference sheet for price, refund, and policy topics, with a weekly accuracy sample on those topics. It adds a written transcript retention rule and a human route the business can actually staff. The KPIs are the accuracy rate on sampled high-risk answers, with any wrong price or policy answer triggering same-day review, and a completed privacy rule before launch.

What do branded agents show at retail scale?

Salesforce's analysis of the 2025 holiday season, covering November 1 through December 31, shows where enterprise retail practice has gone. Salesforce reported record online sales of \$1.29 trillion globally and attributed \$262 billion of holiday spend, or 20% of retail sales, to AI and agents. It reported that retailers that deployed their own AI agents saw a 59% higher growth rate, averaging 6.2% year-over-year sales growth against 3.9% (Salesforce, 2026).

Salesforce also reported that shoppers used retailers' AI and agents for customer service 126% more during the holiday rush than in the two months before it. Agentic service conversations rose 66% from November to December, and agents handled 142% more tasks, such as updating delivery addresses and starting returns, than in the prior two months. For Pandora, Salesforce reported that the jewelry brand deflected 60% of service cases through Clara, its customer service agent serving North America and the United Kingdom since January 2025. Salesforce also credited Clara with contributing to a 10% increase in Net Promoter Score (Salesforce, 2026).


Three limits come with those numbers. Salesforce sells the agent platform it measured, and Pandora is its customer. Salesforce built the figures from aggregated activity on its own platform with extrapolation factors for the broader industry. The growth gap therefore describes an association among retailers rather than a measured cause. The same report shows average selling prices up 7% and online traffic up 13% globally that season (Salesforce, 2026). Rising prices and traffic may account for part of that growth.

A small business will not copy Pandora's stack, yet it can copy the job design. One job answers repeated service questions, a second qualifies sales or booking requests, and both are measured against the human backlog and customer satisfaction.

Retailers running branded AI agents grew faster in Salesforce's holiday data, in a season of rising prices and traffic (Salesforce, 2026). The vendor also reports Pandora's service results for its own product. The tactic is two chat jobs for the first quarter, service questions and sales or booking qualification, with no third job until both clear their KPIs. The KPIs are the resolution rate on service questions with CSAT flat or better, and meetings booked or checkouts started per eligible session on the sales path.

The 30-day Factics playbook

The plan runs four weeks after a gate, and each week leaves evidence the next week depends on.



Before week one: the gate. The owner completes the inquiry count described earlier and writes a go, no-go, or alternative decision. A no-go or an alternative such as a better FAQ ends the plan here, and the count stays on file as the baseline for the next review.

Week one: custody. No technology goes live this week. The team picks one URL family, usually help content or a product explainer, and seldom the only pricing page the business has. It writes the KPI sheet, the disclosure line for the bot's first message, and the human route. It also builds the approved reference sheet for price and policy answers and sets the transcript retention rule. If any of those six items lacks an answer, the tactic has no way to leave evidence, and under the Factics principle it is not ready to run.

Week two: launch. The bot goes live on the chosen pages with the disclosure in its first message. Where traffic allows, half the visitors see the bot and half keep the current experience, the same design

HubSpot used in 2019. Where traffic is too low to split, the team compares against the pre-launch baseline over a longer window and records that the comparison is weaker. Every Friday, the owner annotates failed answers and marks the accuracy sample.

Week three: routing and cost. If the CRM or chat tool can tell customers from prospects, the team adds separate paths. That applies HubSpot's lesson that one bot for everyone underserves both groups (Echt-Wilson, 2019). The owner logs actual annotation hours, the software cost, and the time spent on human handoffs.

Week four: decision. The team compares results against the sheet written in week one and chooses one of four outcomes: continue, modify, stop, or inconclusive. A low-volume page may not produce enough chats in 30 days to support any conclusion, and recording that honestly is better than reading noise as a trend. The pattern moves to a second page only if conversion, CSAT, and accuracy all cleared the bar.

Some failures do not wait for week four. A wrong statement of price or policy, a privacy exposure, or advice in a regulated area the bot should never answer calls for same-day review and, where warranted, suspension.

KPI	What it measures	Note
Conversions per eligible session	Qualified leads, bookings, or purchases per session on the URL family	Counting sessions rather than chats avoids crediting the bot with visitors who would have converted anyway
CSAT on AI chats	Customer satisfaction after the conversation	Compared with a human baseline where one exists; otherwise held to a fixed floor the owner sets before launch, with the sample size recorded
Resolution rate	Questions fully resolved without a human, confirmed in transcripts	Replaces deflection, which counts escaped humans rather than solved problems
Accuracy on high-risk answers	Share of sampled price, refund, and policy answers that were correct	Any wrong answer triggers same-day review
Disclosure coverage	Share of automated chats with the AI label in the first message	Target 100%
Owner hours	Annotation and handoff time logged each week	An input measure, recorded to price the pilot rather than to prove it worked
Net result	Conversions gained against software cost and owner hours	Decides continue, modify, stop, or inconclusive

The Facts loop stays visible all month as three checks. The fact check asks what HubSpot, the independent research, and the law actually say. The tactic check asks what the team did on one set of pages with disclosure, annotation, accuracy checks, and two jobs. The KPI check asks what changed against targets written before launch, and a pilot that cannot answer all three is still a demo.

Where do small teams waste the AI chat advantage?

The first waste is a scoreboard built on deflection alone. In HubSpot's 2019 routing test, fewer chats reached the sales team while meetings booked held steady, so the number that mattered was meetings rather than chat volume (Echt-Wilson, 2019). A team counting only deflection would have read that result backward, which is the confusion Geller described when she warned that containment is mistaken for success (Doerer, 2026).

The second waste is skipped annotation. HubSpot treated transcript annotation as central work and read the early CSAT dip as a training signal. It also began building an annotation interface so more team members could take part (Bodnar, 2024). A small team that leaves Friday review unowned will watch answers drift and may blame the model for a staffing decision.

The third waste is a single bot for every audience, which HubSpot's support-bot experience showed leaves most visitors without the right next step (Echt-Wilson, 2019). The fourth is an unchecked answer on a money page, the exposure Air Canada carried into a tribunal (American Bar Association, 2024).

Channel confusion adds a quieter loss. HubSpot's AI Trends report lists email, text-based social media, and blogs or long-form content as the most common channels for AI-assisted creation. In each channel, 63% to 68% of marketers using AI reported at least a somewhat positive return. Only 4% use AI to write entire pieces of content, while most use it for inspiration, outlines, or drafts they build on (Santiago, 2026). That content work sits beside conversational conversion without replacing it. A polished AI email sequence does not answer a buyer's question on a pricing page in the moment the buyer asks it. An instrumented chat path may serve that buyer where a content campaign cannot.

The losses in this record come from deflection-only scoring, unowned annotation, one bot for every audience, and unchecked answers on money pages (American Bar Association, 2024; Bodnar, 2024; Doerer, 2026; Echt-Wilson, 2019). The tactic is a named owner for weekly annotation and a named owner for the KPI review. Both names may belong to the same solo operator, on different calendar blocks. The KPI is the number of weeks with annotation, accuracy sampling, and KPI review all logged. The target is eight out of the next eight before any expansion to a second set of pages or a second language.

What to watch

Annotation hours are the first number worth tracking, since they are the hidden cost behind every promise of set-and-forget chat. No source reviewed here shows whether they shrink, hold, or grow as a small business's content changes. The disclosure research is moving too, since the 2026 voice-agent results differ from the 2019 sales-call results. The next field studies will show which pattern holds for website chat. Two effective dates arrive on January 1, 2027, when California's large-business chatbot rules and Colorado's conversational AI law take effect. HubSpot's pricing page test of touchless purchases is the enterprise result to watch for small business tools. The useful question is whether its numbers arrive with the denominators the first post left out.

References

- Aarons, W., & Sarkar, A. (2026, October 8). AI adoption and employment expectations: Evidence from a survey of small business owners. *Liberty Street Economics*. Federal Reserve Bank of New York. <https://libertystreeteconomics.newyorkfed.org/2026/10/ai-adoption-and-employment-expectations-evidence-from-a-survey-of-small-business-owners/>
- American Bar Association. (2024, February). BC tribunal confirms companies remain liable for information provided by AI chatbot. *Business Law Today*. https://www.americanbar.org/groups/business_law/resources/business-law-today/2024-february/bc-tribunal-confirms-companies-remain-liable-information-provided-ai-chatbot/
- Bodnar, K. (2024). *AI experiment: Hacking HubSpot chat to improve the customer experience*. HubSpot Blog. Updated May 8, 2025. <https://blog.hubspot.com/marketing/ai-chat-experiment>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889 to 942. <https://doi.org/10.1093/qje/qjae044>
- Cal. Bus. & Prof. Code § 17941. <https://law.justia.com/codes/california/code-bpc/division-7/part-3/chapter-6/section-17941/>
- Colorado General Assembly. (2026). *HB 26-1263: Conversational artificial intelligence service operator requirements*. <https://leg.colorado.gov/bills/HB26-1263>
- Doerer, K. (2026, August 17). Only one-quarter of AI customer service use cases produce ROI. *CX Dive*. <https://www.customerexperiencedive.com/news/only-one-quarter-of-ai-customer-service-use-cases-produce-roi/827951/>
- Echt-Wilson, A. (2019, July). *How HubSpot personalized our chatbots to improve the customer experience and support our sales team*. HubSpot Blog. Updated May 13, 2025. <https://blog.hubspot.com/marketing/chatbots-improve-customer-experience-experiment>
- European Commission. (n.d.). *Transparency obligations under Article 50 of the AI Act*. Shaping Europe's Digital Future. <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>
- Grundy, A., Breaux, C., & Khatiwoda, D. (2026, May 26). *Large firms with at least 20 employees biggest AI users*. U.S. Census Bureau. <https://www.census.gov/library/stories/2026/05/ai-use-businesses.html>
- Habel, J., Zawadzki, J., & Freise, P. (2026). *Selling with AI voice agents: Does disclosure still matter?* SSRN. <https://doi.org/10.2139/ssrn.7013259>
- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, 38(6), 937 to 947. <https://doi.org/10.1287/mksc.2019.1192>
- Mahrouyan, O. (2026, September 30). *California enacts AB 1609: What the new human customer-service and chatbot rules actually require*. <https://mahrolaw.com/insights/california-ab-1609-chatbot-human-customer-service-law>

Mozafari, N., Weiger, W. H., & Hammerschmidt, M. (2021). Trust me, I'm a bot: Repercussions of chatbot disclosure in different service frontline settings. *Journal of Service Management*.
<https://doi.org/10.1108/JOSM-10-2020-0380>

NFIB. (2025, June 25). *New NFIB report: How small businesses incorporate tech and AI advancements*.
<https://www.nfib.com/news/press-release/new-nfib-report-how-small-businesses-incorporate-tech-and-ai-advancements/>

Office of the Governor of California. (2026, September 28). *Governor Newsom signs commonsense legislation to make your life easier*. <https://www.gov.ca.gov/2026/09/28/governor-newsom-signs-commonsense-legislation-to-make-your-life-easier/>

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

- Latest update: Puglisi, B. C. (2026, September 2; updated September 15). *Facts: The method behind valuable content and trusted AI*. basilpuglisi.com.
<https://basilpuglisi.com/facts/>

Salesforce. (2026, January 8). *Holiday season rakes in record \$1.29T for retailers, Salesforce data shows*. Salesforce News. <https://www.salesforce.com/news/stories/2025-holiday-shopping-data/>

Santiago, E. (2026, September 8). *The HubSpot Blog's AI Trends for Marketers report [key findings from 1,000+ marketing pros]*. HubSpot Blog. <https://blog.hubspot.com/marketing/state-of-ai-report>

Utah Code § 13-77-103. <https://law.justia.com/codes/utah/title-13/chapter-77/part-1/section-103/>

Zhang, S., & Narayandas, D. (forthcoming). Engaging customers with AI in online chats: Evidence from a randomized field experiment. *Management Science*. <https://doi.org/10.1287/mnsc.2022.03920>

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.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

[Spotify](#) | [Apple Podcasts](#) | [Amazon Music](#) | [YouTube Playlist](#)

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.