Proceedings of ICFTBA 2026 Symposium: Driving Organizational Strategic Change for Financial Gains
DOI: 10.54254/2754-1169/2026.BH37993

Why Customers Welcome or Resist AI Salespeople: Formation Mechanisms of Algorithm Aversion and Algorithm Appreciation in Large Language Model Sales

Dasen Yin

Hohhot No.2 High School, Hohhot, China
17601615850@163.com

Abstract. Large language models (LLMs) now answer product questions, recommend items and help close deals in online shops, banks and telecom services. Customers react to these AI sales agents in two opposite ways. Some avoid them, a pattern researchers call algorithm aversion, while others prefer them to human advice, a pattern called algorithm appreciation. This paper reviews the main empirical studies on both patterns and asks how each one forms when the "algorithm" is a conversational sales agent. Drawing on research in judgment and decision making, consumer behavior and human computer interaction, it proposes a framework with three layers. The first layer contains the cues that an LLM sales agent gives off. The second layer contains five psychological pathways that operate during the conversation: competence belief, error sensitivity, empathy and uniqueness concern, social presence, and perceived control. The third layer contains situational conditions such as product type, the timing of AI identity disclosure, the customer's emotional state and prior AI experience. Six propositions are derived. The paper argues that aversion and appreciation are not fixed traits of customers but outcomes of how these layers combine in a particular sales encounter, and it discusses implications for firms and for future student research.

Keywords: Algorithm Aversion, Algorithm Appreciation, Large Language Models, AI Sales, Chatbots

1. Introduction

Buying something has always involved a conversation. A shop assistant asks what the customer needs, explains the options and tries to clear up doubts before the customer pays. Since generative AI tools became widely available to the public, many companies have started to hand part of this conversation to large language models. An LLM sales agent can talk in fluent natural language, keep track of what the customer said earlier in the chat, compare products in seconds and work at any hour.

Whether customers accept this new kind of salesperson is far from obvious. Research on AI in service work shows real gains. In a study of 5,172 customer support agents, access to a generative AI assistant raised the number of customer issues resolved per hour by 15% on average [1]. Yet a field experiment with more than 6,200 customers of a financial services company found that when a

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.BH37993

sales chatbot revealed that it was a machine before the call began, purchase rates fell by more than 79.7%, even though the undisclosed chatbot sold as well as proficient human staff [2]. The same technology can be welcomed in one situation and rejected in another.

Psychologists have names for these two reactions. Algorithm aversion describes the tendency to avoid an algorithm after seeing it make mistakes, even when it performs better than a human [3]. Algorithm appreciation describes the opposite finding that people often give more weight to advice when they believe it comes from an algorithm rather than a person [4]. Much of this research, however, used forecasting tasks with numbers, such as predicting the success of MBA applicants or estimating a person's weight from a photograph. A sales conversation with an LLM is different. It is social, emotional and persuasive, and the customer has money at stake.

This paper therefore asks one question: through what mechanisms do algorithm aversion and algorithm appreciation form when customers deal with LLM sales agents? To answer it, the paper first defines the key concepts, then reviews the empirical literature, and finally builds a theoretical framework with testable propositions. The aim is not to report new data but to organise existing findings in a way that a later survey or experiment could test.

2. Key concepts

2.1. Algorithm aversion

In the study that gave the idea its name, Dietvorst asked participants to make forecasts, for example about how well MBA applicants would perform [3]. Participants could rely on their own forecasts or on a statistical model. After watching the model make some errors, many of them chose to rely on themselves, even though the model was more accurate than they were. The central point is that people forgave human mistakes more easily than machine mistakes. Later reviews show that aversion appears in many domains and has several kinds of causes. Burton reviewed 61 articles and grouped the proposed causes and solutions into five themes: expectations and expertise, decision autonomy, incentivization, cognitive compatibility, and divergent rationalities [5]. Mahmud reviewed 80 empirical studies and sorted the influencing factors into four groups: factors of the algorithm, of the individual, of the task, and of the wider environment [6].

2.2. Algorithm appreciation

Logg challenged the idea that people always distrust algorithms [4]. In a series of experiments, participants received the same advice labelled either as coming from a person or from an algorithm. They moved their own estimates further toward the advice when it was labelled as algorithmic. The effect was weaker when people had to choose between the algorithm and their own judgment, and it was weaker among experienced professionals. Jussupow later argued that aversion and appreciation should be treated as two directions of one underlying preference, the preference for a human or an algorithm in a decision, and that which direction appears depends on what is being compared in the study [7].

2.3. The LLM sales scenario

In this paper, an LLM sales scenario means a buying situation in which the customer talks, by text or voice, with an agent powered by a large language model whose job includes informing, recommending, persuading and helping to complete a purchase. Three features set this scenario

DOI: 10.54254/2754-1169/2026.BH37993

apart from the algorithms studied in earlier research. First, the agent communicates in natural, humanlike language rather than showing a number or a ranked list. Second, its answers are open ended, and LLMs can produce fluent text that is not supported by facts, which researchers call hallucination [8]. Third, the agent is deployed by the seller, so customers may wonder whether it works for them or for the company. Each feature connects to one or more of the mechanisms discussed in Section 4.

3. Literature review

3.1. Evidence for algorithm aversion

Besides the error effect found by Dietvorst, two lines of consumer research are important for sales [3]. Longoni found that consumers were reluctant to use medical AI, and one reason was uniqueness neglect: the belief that an AI cannot take a person's unique characteristics and circumstances into account [9]. Castelo showed that aversion depends on the task [10]. People trusted and used algorithms less for tasks they saw as subjective, such as dating advice, than for tasks they saw as objective, such as financial advice. In a field study with Facebook advertisements, dating advice ads received more clicks when they were attributed to a human than to an algorithm, while financial advice ads showed no such difference. Both findings suggest that aversion grows when a decision feels personal.

3.2. Evidence for algorithm appreciation

Logg provided the clearest evidence of appreciation. Other studies show when appreciation is likely [4]. Longoni and Cian identified a "word-of-machine" effect: people prefer AI recommenders when utilitarian attributes, such as function and practical value, are important, and prefer human recommenders when hedonic attributes, such as sensory pleasure, are important [11]. They trace this to a lay belief that AI is more competent than humans in the utilitarian realm and less competent in the hedonic realm. The effect disappeared when the recommendation came from a human and an AI working together. Aversion can also be reduced by design. Dietvorst found that people were much more willing to use an imperfect algorithm when they were allowed to modify its forecasts, even by a small amount [12].

3.3. Evidence from sales and service chatbots

A smaller group of studies looks directly at chatbots that talk to customers. The field experiment by Luo is the closest to a real sales setting [2]. Customers who were called by an undisclosed chatbot bought about as often as those called by proficient human workers, and four times as often as those called by inexperienced workers. When the chatbot disclosed its identity before the conversation, customers became curt, calls became shorter and purchases fell sharply. The authors found that customers perceived the disclosed bot as less knowledgeable and less empathetic, although its actual performance had not changed. The negative effect was smaller when disclosure came later and among customers with prior AI experience.

Two studies examine humanlike design. Adam ran an online experiment with a customer service chatbot and found that humanlike verbal cues increased users' compliance with the chatbot's request, an effect explained by a stronger sense of social presence [13]. Crolic, using data from an international telecommunications company and four experiments, found that anthropomorphic

DOI: 10.54254/2754-1169/2026.BH37993

chatbots lowered satisfaction, firm evaluations and purchase intentions among customers who entered the chat already angry, but not among other customers [14]. The cause was expectancy violation: the humanlike chatbot raised expectations that it then failed to meet. Finally, Brynjolfsson showed that generative AI can improve service performance when human agents stay in charge and decide whether to use the AI's suggestions [1]. Table 1 summarises these studies.

Table 1. Key empirical studies on algorithm aversion and appreciation relevant to AI sales

Study	Setting and method	Main finding	Direction
Dietvorst [3]	Laboratory experiments; forecasting tasks	After seeing an algorithm err, people chose it less often than a human forecaster, even though it was more accurate.	Aversion
Dietvorst [12]	Laboratory experiments; forecasting tasks	People were more willing to use an imperfect algorithm when they could slightly modify its forecasts.	Aversion reduced
Logg [4]	Experiments; estimation and forecasting tasks	The same advice received more weight when labelled as algorithmic; the effect weakened for experts and when people compared the algorithm with their own judgment.	Appreciation
Castelo [10]	Experiments and a Facebook advertising field study	Aversion was stronger for tasks seen as subjective; dating advice attributed to a human drew more clicks, financial advice showed no difference.	Depends on task
Longoni [9]	Experiments; medical services	Consumers resisted medical AI, partly because of uniqueness neglect.	Aversion
Luo [2]	Field experiment; more than 6,200 customers; outbound sales calls	Undisclosed chatbots sold as well as proficient workers; disclosure before the call reduced purchase rates by more than 79.7%.	Aversion after disclosure
Adam [13]	Online experiment; customer service chatbot	Humanlike verbal cues raised compliance with the chatbot's request, mediated by social presence.	Appreciation
Crolic [14]	Company data and four experiments; customer service	For angry customers, anthropomorphic chatbots lowered satisfaction, firm evaluation and purchase intentions.	Aversion (conditional)
Longoni & Cian [11]	Experiments; product recommendations	AI recommenders were preferred for utilitarian attributes and humans for hedonic ones; the effect disappeared with human and AI hybrid advice.	Both
Brynjolfsson [1]	Staggered rollout; 5,172 customer support agents	A generative AI assistant raised issues resolved per hour by 15% on average.	Performance benefit

4. A theoretical framework

Figure 1 summarises the framework developed in this section. Its main idea is that a customer's reaction to an LLM sales agent form in three layers. The first layer is the set of cues the agent gives off: its language and persona, its product knowledge, the risk of hallucination, and the fact that it serves the seller and may or may not disclose that it is an AI. The second layer is a set of psychological pathways that turn these cues into trust or distrust during the conversation. The third layer is the situation, which strengthens or weakens each pathway. In this view, aversion and appreciation are two outcomes of the same evaluation process, and they differ only in which pathways dominate.

DOI: 10.54254/2754-1169/2026.BH37993

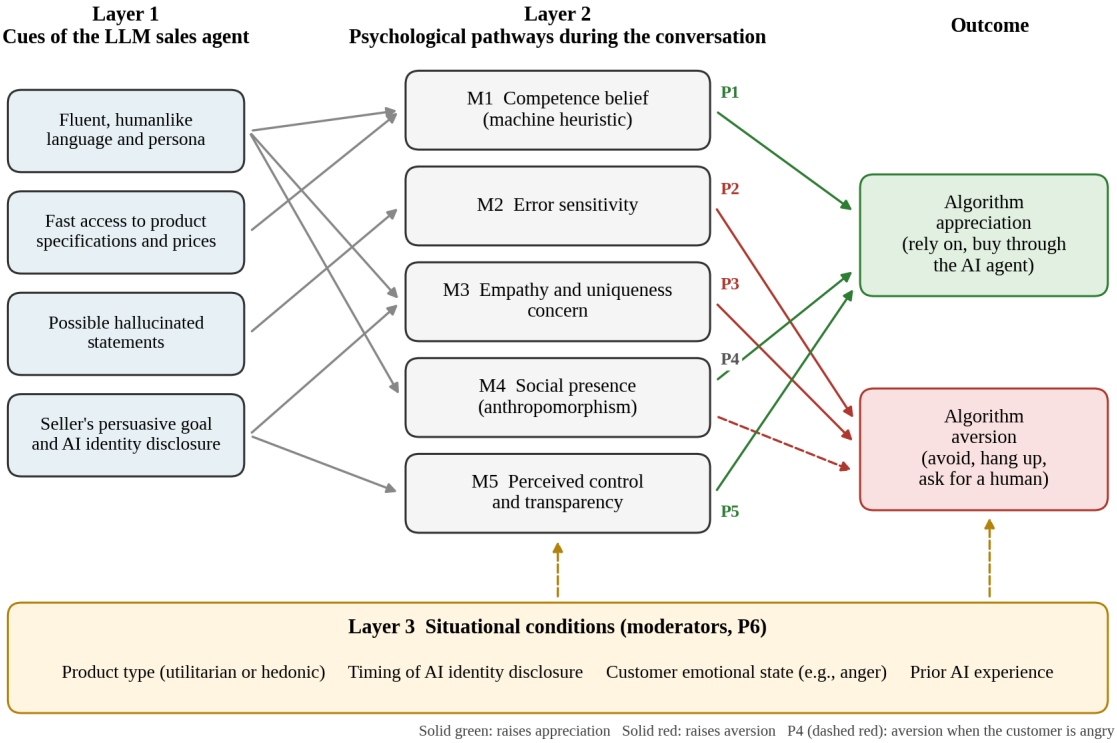


Figure 1. A three-layer framework of how algorithm aversion and appreciation form in LLM sales

4.1. Competence belief pathway

Many people hold a lay theory that machines are objective, consistent and well informed. Sundar called this the machine heuristic: when a message is seen as coming from a machine, people may judge it as more neutral and accurate [15]. Longoni and Cian found a related belief in consumer choice, namely that AI is better than humans at judging utilitarian attributes [11]. An LLM sales agent can retrieve specifications, prices and comparisons quickly, so when a customer is comparing phone batteries or insurance terms, the agent's apparent knowledge fits this belief well.

P1. When utilitarian attributes of a product are salient, stronger beliefs in machine competence lead customers to show algorithm appreciation toward an LLM sales agent.

4.2. Error sensitivity pathway

The belief that machines should be accurate also makes their mistakes more costly. Dietvorst showed that seeing an algorithm err reduced people's use of it more than seeing a human make errors [3]. Burton link this to expectations: if people expect an algorithm to be nearly perfect, any error looks like proof that it cannot be trusted, whereas a human error looks like a normal slip [5]. LLMs make a particular kind of error, because they can state false information fluently [8]. In a sales chat, an invented product feature or a wrong return policy is easy for a customer to discover later, and because the error came in a confident and polished voice, the gap between what was expected and what happened may be large.

P2. A visible factual error by an LLM sales agent reduces customers' willingness to rely on it more than a comparable error by a human salesperson.

DOI: 10.54254/2754-1169/2026.BH37993

4.3. Empathy and uniqueness pathway

Selling is not only about facts; customers also want to feel understood. Luo found that once customers knew the caller was a chatbot, they judged it as less knowledgeable and less empathetic, even though nothing about its performance had changed [2]. Longoni described the related worry of uniqueness neglect, and Castelo showed that aversion is stronger when a task feels subjective [9, 10]. Choosing a gift, a wedding outfit or a holiday involves personal taste and emotion, so these concerns are more likely to be triggered in such purchases than in the purchase of a charging cable.

P3. When hedonic attributes or the customer's personal uniqueness are salient, perceived lack of empathy and uniqueness neglect lead to algorithm aversion toward an LLM sales agent.

4.4. Social presence pathway

LLMs make it easy to give a sales agent a name, a friendly tone and a personality. Research on computers as social actors shows that people respond to computers socially, applying norms such as politeness and reciprocity, even though they know the computer is not a person [16]. Epley explain that people are more likely to see human qualities in an agent when humanlike cues are present and when they want to understand or connect with it [17]. Humanlike design can therefore raise social presence and cooperation, as Adam found [13]. Crolic show the other side: for angry customers, the same design raised expectations that the chatbot could not meet. Anthropomorphism therefore cuts both ways [14].

P4. Humanlike design of an LLM sales agent increases algorithm appreciation among customers in a neutral emotional state but increases algorithm aversion among customers who are angry.

4.5. Control and transparency pathway

People accept algorithms more readily when they keep some control over the final decision. This is shown by the modification effect in Dietvorst, by the disappearance of the word-of-machine effect under human and AI collaboration in Longoni and Cian, and by the setting studied by Brynjolfsson, in which human agents could accept, edit or ignore the AI's suggestions [1, 11, 12]. In sales, control could mean that customers can ask the agent to explain a recommendation, change the filters it uses, or reach a human at any time. Transparency about AI identity belongs to the same pathway. Luo found that the penalty for disclosure was smaller when it came later and when customers already had experience with AI, which suggests that familiarity shapes how disclosure is received [2].

P5. Giving customers control over an LLM sales agent's recommendation, including an easy route to a human, increases algorithm appreciation.

P6. The negative effect of AI identity disclosure on purchase is weaker among customers with more prior AI experience and when the conversation has already shown the agent's competence.

Table 2 links each pathway to its theoretical roots, the cues that are likely to activate it in LLM sales, and its expected outcome.

Table 2. Five pathways to algorithm aversion and appreciation in LLM sales

Pathway	Theoretical roots	Typical trigger in LLM sales	Expected outcome	Prop.
Competence belief	Machine heuristic; word-of-machine effect [11, 15]	Fast answers on specifications, prices and comparisons	Appreciation	P1
Error sensitivity	Aversion after seeing errors; expectations [3, 5]	Hallucinated features, prices or policies	Aversion	P2

DOI: 10.54254/2754-1169/2026.BH37993

Table 2. (continued)

Empathy and uniqueness	Uniqueness neglect; task subjectivity [9, 10]	Gifts, fashion, travel and other emotional purchases	Aversion	P3
Social presence	Computers as social actors; anthropomorphism [16, 17]	Name, avatar, friendly and humanlike tone	Appreciation in a neutral mood; aversion when angry	P4
Control and transparency	Modifiable algorithms; disclosure effects [2, 12]	Explanations, adjustable filters, human handoff, AI identity disclosure	Appreciation	P5, P6

5. Discussion

5.1. One evaluation process, two outcomes

The review suggests that customers should not be divided into "AI lovers" and "AI haters." The same person may be happy to let an LLM agent pick a laptop charger and still refuse its help with an anniversary gift, or may trust it at the start of a chat and leave after one wrong answer. This is consistent with the argument of Jussupow that the direction of preference depends on the decision setting [7]. The framework also adds a sense of time. Before and at the start of a conversation, competence beliefs and disclosure are likely to matter most. As the conversation continues, errors, perceived empathy and the feeling of being in control update the customer's judgment. Appreciation can therefore turn into aversion within a single chat, and the reverse is also possible.

5.2. Implications for firms

Four practical points follow. First, LLM agents are likely to be most welcome in purchases driven by utilitarian attributes, while hedonic and emotional purchases may benefit from a human and AI combination, which removed the word-of-machine effect in Longoni and Cian [11]. Second, because machine errors weigh heavily, answers about prices, specifications and policies should be checked against verified product data before they reach the customer. Third, the agent's persona could adapt to the customer's mood; a plainer and less humanlike style may suit complaint situations, given the findings of Crolic [14]. Fourth, hiding the AI identity is not an acceptable solution on ethical grounds. A better response to the disclosure penalty found by Luo is to earn trust through visible competence and to offer customers control, including a clear way to reach a human [2].

6. Limitations and future research

This paper is a theoretical review and does not test its propositions. Several of the studies it relies on were carried out before modern LLMs became common, so features such as long, flexible conversations and hallucination need to be studied directly. Most experiments also come from a small number of countries, and cultural differences in trust toward machines and toward salespeople may change the strength of each pathway. Future work could test the framework with a simple online experiment that a student research team could run. For example, participants could read a sales chat about either a utilitarian product or a hedonic product, in which the AI agent either makes a factual error or does not. Measuring willingness to buy through the agent, perceived competence and perceived empathy would test P1, P2 and P3 together. Surveys of real customers who have used AI shopping assistants could then examine P5 and P6.

DOI: 10.54254/2754-1169/2026.BH37993

7. Conclusion

Large language models are changing who, or what, customers talk to when they buy. Earlier research found both algorithm aversion and algorithm appreciation, and the evidence from sales and service chatbots shows that both appear in commercial settings. This paper has argued that the two reactions come from the same process. Cues from the LLM agent activate five pathways, namely competence belief, error sensitivity, empathy and uniqueness concern, social presence, and control and transparency, and situational conditions decide which pathways dominate. Understanding these mechanisms can help firms design AI sales agents that customers trust for good reasons and help researchers test when and why that trust is given or withdrawn.

References

[1] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942.
[2] 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), 937–947.
[3] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126.
[4] Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90–103.
[5] Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making, 33*(2), 220–239.
[6] Mahmud, H., Islam, A. K. M. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. *Technological Forecasting and Social Change, 175*, 121390.
[7] Jussupow, E., Benbasat, I., & Heinzl, A. (2024). An integrative perspective on algorithm aversion and appreciation in decision-making. *MIS Quarterly, 48*(4), 1575–1590.
[8] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), Article 248.
[9] Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research, 46*(4), 629–650.
[10] Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research, 56*(5), 809–825.
[11] Longoni, C., & Cian, L. (2022). Artificial intelligence in utilitarian vs. hedonic contexts: The "word-of-machine" effect. *Journal of Marketing, 86*(1), 91–108.
[12] Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science, 64*(3), 1155–1170.
[13] Adam, M., Wessel, M., & Benlian, A. (2021). AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets, 31*(2), 427–445.
[14] Crolic, C., Thomaz, F., Hadi, R., & Stephen, A. T. (2022). Blame the bot: Anthropomorphism and anger in customer–chatbot interactions. *Journal of Marketing, 86*(1), 132–148.
[15] Sundar, S. S. (2008). The MAIN model: A heuristic approach to understanding technology effects on credibility. In M. J. Metzger & A. J. Flanagin (Eds.), *Digital media, youth, and credibility* (pp. 73–100). MIT Press.
[16] Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues, 56*(1), 81–103.
[17] Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review, 114*(4), 864–886.