

Appendix A: Literature Review

Table A.1 Literature Review on Conversational AI and Social Bots

Study	AI Type and Purpose	Methods	Social Media	Platform-owned	Research Questions
Schuetzler et al. (2020)	Conversational AI	• Online experiments			Investigates the impact of CA conversational skill (generic, varied and tailored) on the perception of CA humanness and anthropomorphism
Schanke et al. (2021)	Conversational AI	• Field experiment		✓	Studies the impact of chatbot anthropomorphism on transaction conversion by randomly exposing consumers to humor, communication delays, and social presence
Seeger et al. (2021)	Conversational AI	• Online experiments			Develops and tests a framework to design anthropomorphic CAs with three factors of anthropomorphism: technology design-related, task-related, and individual factor
Chandra et al. (2022)	Conversational AI	• Survey-based studies			Focuses on human-like interactional competencies in conversational AI agents—cognitive, relational, and emotional—in facilitating user engagement
Brendel et al. (2023)	Conversational AI	• Online experiments			Focuses on understanding user aggression towards humanlike and non-humanlike CAs when they produce errors
Gnewuch et al. (2024)	Conversational AI	• Field experiment • Online experiment		✓	Focuses on disclosing human involvement in hybrid service agents (human-AI)
Han et al. (2023)	Conversational AI	• Lab experiments			Explores how, when, and why an AI agent’s expression of positive emotions affects customers’ service evaluations
Sachdeva et al. (2024)	Conversational AI	• Lab experiments • Online experiments			Investigates how collecting reviews using chatbots (vs. forms) impacts user experience and the nature of the reviews
Salge et al. (2022)	Social bot for (mis)information dissemination	• Theoretical framework • Case-based analysis	✓		Proposes a theoretical framework of algorithmic conduit brokerage (algorithmic social alertness and algorithmic social transmission) for understanding information dissemination by bots
Benjamin and Raghu (2023)	Social bot for opinion manipulation	• Deep learning • Design science	✓		Focuses on social bot detection and uses human reactions to social bot messages to augment existing social bot detection capabilities
Delkhosh et al. (2023)	Social bot for upvoting	• PVAR	✓		Investigates the impact of bots’ voting activity on human users’ behavior
He et al. (2024)	Social bot for content moderation	• DID • IV	✓	✓	Focuses on understanding how volunteer human moderators react to the introduction of bot moderators in terms of human moderators’ community-policing and nurturing efforts
Safadi et al. (2024)	Social bot for text summary and content moderation	• PVAR	✓	✓	Investigates how bots (reflexive and supervisory) influence human-to-human interactions in online communities
This study	Social Bot for public engagement	• IV • Online experiment • Policy learning • DID	✓	✓	Investigates the impact of a platform-owned social bot’s public interactions with human posts on user engagement and user posting activity on a social media platform. Proposes strategies for a social bot deployment

Notes. CA stands for conversational agent. PVAR stands for panel vector autoregression. DID stands for the difference-in-differences analysis. IV stands for the instrumental variable analysis.

Appendix B: Robustness Checks

B.1. Reverse Causality

To alleviate the concern that the bot is more likely to comment on posts that already have engagement from peers, we re-estimate the model using only posts where the bot comments first as the treated group. Table B.1 reports the estimation results.

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.0730*** (0.0064)	0.0997*** (0.0063)	0.0932*** (0.0062)	0.2178*** (0.0060)	0.2235*** (0.0058)	0.2187*** (0.0058)
User Controls	No	Yes	Yes	No	Yes	Yes
Post Controls	No	Yes	Yes	No	Yes	Yes
Date FEs	No	No	Yes	No	No	Yes
Hour of Day FEs	No	No	Yes	No	No	Yes
Observations	63,198	63,198	63,198	63,198	63,198	63,198
R-squared	0.002	0.157	0.163	0.019	0.175	0.181

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

For the number of likes, the general public, including researchers, cannot observe the timing of individual likes, making it impossible to exclude likes before bot comments when constructing the dependent variable, *LogLikes*. To alleviate the reverse causality concern, we conduct a subsample analysis focusing on posts where bot comments appeared very shortly after the original post. The rationale is that human likes are unlikely to occur before an extremely fast bot comment. Given that the average bot comment speed in our sample is 68 seconds, we consider two cutoffs: 60 seconds and 30 seconds. Columns (1) to (3) in Table B.2 report the results with posts receiving bot comments in less than 60 seconds as treated and Columns (4) to (6) report the results with posts receiving bot comments in less than 30 seconds as treated. Posts without receiving bot comments served as the control group in both cases. As shown, the coefficients of *BotComment* remain significantly positive, thereby alleviating the concern about reverse causality.

B.2. Alternative Specifications

We conduct four sets of analyses to demonstrate the robustness of our findings to alternative empirical specifications. First, due to the inclusion criterion in the main sample, which only includes posts where the user mentions the bot (i.e., including *@CommentRobot*) for the first time, adding user-fixed effects to model (1) is infeasible. Instead, we compile a list of user characteristics to mitigate concerns regarding the selection of bot comments based on user attributes. To further alleviate this concern, we construct an alternative

Table B.2 Robustness Check — Subsample Analysis with Fast Bot Comments

	<i>LogLikes</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.0663*** (0.0067)	0.1083*** (0.0066)	0.1007*** (0.0065)	0.0644*** (0.0069)	0.1049*** (0.0068)	0.0982*** (0.0068)
User Controls	No	Yes	Yes	No	Yes	Yes
Post Controls	No	Yes	Yes	No	Yes	Yes
Date FEs	No	No	Yes	No	No	Yes
Hour of Day FEs	No	No	Yes	No	No	Yes
Comment Speed	<= 60s	<= 60s	<= 60s	<= 30s	<= 30s	<= 30s
Observations	60,466	60,466	60,466	58,053	58,053	58,053
R-squared	0.002	0.157	0.163	0.002	0.158	0.163

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

sample comprising users who mention the bot in multiple posts. With such a sample, we can add user-fixed effects in the econometric model and provide within-user evidence regarding the effects of bot comments.

The estimation results are reported in Table B.3

Table B.3 Robustness Check — User Fixed Effects

	<i>LogLikes</i>		<i>LogComments</i>	
	(1)	(2)	(3)	(4)
<i>BotComment</i>	0.1279*** (0.0056)	0.1366*** (0.0059)	0.2835*** (0.0061)	0.2892*** (0.0063)
Post Controls	No	Yes	No	Yes
User FEs	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes
Observations	59,038	59,038	59,038	59,038
R-squared	0.708	0.711	0.647	0.650

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Second, to alleviate the concern that post content could be a confounding factor, besides the LSA variables, we further control for post categories in model (1). Table B.4 reports the results with post category as one of the control variables.

Table B.4 Robustness Check — Control for Post Category

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.0951*** (0.0064)	0.1254*** (0.0065)	0.1186*** (0.0064)	0.2475*** (0.0061)	0.2483*** (0.0059)	0.2440*** (0.0059)
User Controls	No	Yes	Yes	No	Yes	Yes
Post Controls	No	Yes	Yes	No	Yes	Yes
Date FEs	No	No	Yes	No	No	Yes
Hour of Day FEs	No	No	Yes	No	No	Yes
Observations	64,365	64,365	64,365	64,365	64,365	64,365
R-squared	0.004	0.168	0.172	0.024	0.188	0.194

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Third, log-transformed engagement metrics serve as the outcome variable in the main analyses due to their high skewness (Kumar et al. 2022). As a robustness check, we use the raw number of likes and comments as the outcome variable and estimate two commonly used count models—Poisson and Negative Binomial. Table B.5 presents the estimation results.

	Poisson		Negative Binomial	
	(1)	(2)	(3)	(4)
	<i>Likes</i>	<i>Comments</i>	<i>Likes</i>	<i>Comments</i>
<i>BotComment</i>	0.9593*** (0.1797)	0.2548*** (0.0349)	0.6146*** (0.0448)	0.3343*** (0.0185)
User Controls	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes
Observations	64,365	64,365	64,365	64,365

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Finally, to alleviate the selection concern, a common approach in the literature is matching posts with and without bot comments based on observable characteristics. We implement the coarsened exact matching (Iacus et al. 2012) to ensure the comparability between posts with and without bot comments regarding observable characteristics. The results are illustrated in Table B.6. Across various model specifications, the coefficients of *BotComment* are consistently positive and statistically significant in Tables B.3 - B.6 aligning with the results obtained in the main analyses.

	<i>LogLikes</i>		<i>LogComments</i>	
	(1)	(2)	(3)	(4)
	<i>BotComment</i>	0.0981*** (0.0072)	0.1000*** (0.0068)	0.2473*** (0.0074)
User Controls	No	Yes	No	Yes
Post Controls	No	Yes	No	Yes
Date FEs	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes
Observations	46,749	46,749	46,749	46,749
R-squared	0.015	0.122	0.044	0.160

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

B.3. Instrumental Variables

One potential concern is that users who had prior interactions with the bot may be more likely to make AI-related posts, which could confound our results. To alleviate this concern, we compile a list of AI-related

keywords, including artificial intelligence, language model, deep learning, algorithm, neural net, image generation, text generation, large model, and reinforcement learning. As shown in Row 1 of Table B.7, the number of AI-related posts in our sample is minimal, and, more importantly, there is no significant difference in AI-related posting behavior between users who had and had not previously interacted with the bot. Moreover, we compare average engagement levels between users who had prior bot interactions and those who did not, using data from November 1, 2023, to January 1, 2024. As illustrated in Rows 2 and 3 of Table B.7, there is no significant difference in engagement, reinforcing the validity of *BotInteractionBefore* as an instrumental variable.

Table B.7 AI-related Posts and User Engagement

Variable	<i>BotInteractionBefore</i> = 0	<i>BotInteractionBefore</i> = 1	Diff.
Number of AI-related posts	0.020 (0.001)	0.027 (0.004)	-0.007 (0.006)
Average number of likes	8.928 (1.231)	6.416 (0.919)	2.512 (5.628)
Average number of comments	2.749 (0.277)	4.400 (0.209)	-1.651 (1.265)

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

The bot’s capacity to comment is limited, with approximately 40% of posts in our sample receiving bot comments, as shown in Table 1. Due to this limitation, an increase in posts requesting bot interactions reduces the likelihood of receiving a bot comment. Leveraging this feature, we propose an alternative IV, *BotTraffic*, which measures the number of posts mentioning the bot immediately before and after the focal post. Two variables, *BotTraffic60* and *BotTraffic90*, capture bot mentions within 60 and 90 seconds of the focal post, respectively. Over 85% of bot comments in our sample occur within 60 seconds, making this an appropriate threshold. Given this short time window and the platform’s large user base, these posts are unlikely to compete for attention and engagement, satisfying the exclusion restriction. The results of the 2SLS analysis with these alternative IVs are reported in Table B.8. Columns (1) and (4) show that *BotTraffic* is negatively associated with the likelihood of receiving a bot comment, confirming the instrument’s relevance with F-statistics exceeding 10. The significantly positive coefficients of *BotComment* in Columns (2), (3), (5), and (6) align with the results in Table 2, further supporting the robustness of our findings.

To further alleviate the concern regarding external issues that could affect both the likelihood of receiving bot comments and the post engagement, we construct a variable, *Influencers* (added to User Controls) to capture the number of posts from Weibo influencers (for whom the number of followers exceeds 100K) that

Table B.8 Robustness Check — Alternative Instrumental Variables

	First Stage	Second Stage		First Stage	Second Stage	
	<i>BotComment</i> (1)	<i>LogLikes</i> (2)	<i>LogComments</i> (3)	<i>BotComment</i> (4)	<i>LogLikes</i> (5)	<i>LogComments</i> (6)
<i>BotTraffic60</i>	-0.0004*** (0.0001)					
<i>BotTraffic90</i>				-0.0003*** (0.0001)		
<i>BotComment</i>		3.9292*** (1.1279)	2.8447*** (0.8097)		3.7011*** (1.0080)	2.6552*** (0.7189)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	13.59			15.48		
Observations	64,365	64,365	64,365	64,365	64,365	64,365

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

mentioned CommentRobot within 12 hours of the focal post. Influencers discussing CommentRobot might represent external issues because their discussions could draw attention from Weibo users, which may increase the number of CommentRobot requests and attract more user engagement for these requests. The results, presented in Table B.9, confirm the robustness of the IV estimates to external influences. We also test a stricter threshold of 500K followers, and the results remain qualitatively unchanged.

Table B.9 Robustness Check — IV Analysis with Controlling for Influencer Posts

	First Stage	Second Stage		First Stage	Second Stage	
	<i>BotComment</i> (1)	<i>LogLikes</i> (2)	<i>LogComments</i> (3)	<i>BotComment</i> (4)	<i>LogLikes</i> (5)	<i>LogComments</i> (6)
<i>BotTraffic60</i>	-0.0004*** (0.0001)					
<i>BotTraffic90</i>				-0.0003*** (0.0001)		
<i>BotComment</i>		4.0159*** (1.1607)	2.9031*** (0.8317)		3.7794*** (1.0352)	2.7075*** (0.7369)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	13.31			15.19		
Observations	64,365	64,365	64,365	64,365	64,365	64,365

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

The second IV, *BotTraffic*, may suffer from the insignificant temporal patterns of bot commenting. To alleviate this concern, we conduct two subsample analyses. First, we construct a binary variable *Late* to indicate whether a post was published during the least active periods on Weibo. We restrict the sample to posts made outside these low-activity periods (i.e., $Late = 0$), and the results, reported in Table B.10, remain consistent with those in Table B.8. Second, we consider the busiest posting hours in our sample, where the number of posts exceeded 3,900 per hour. Given the bot's 40% response rate, we expect engagement effects

to be unaffected unless bot mentions exceed 1,560 within an hour. Restricting our sample to high-traffic hours (0:00–2:00 AM and 9:00 AM–11:00 PM), we find that the IV estimates remain consistent, as reported in Table B.11.

Table B.10 Robustness Check — Alternative IVs on Subsamples with Late = 0

	First Stage	Second Stage		First Stage	Second Stage	
	<i>BotComment</i> (1)	<i>LogLikes</i> (2)	<i>LogComments</i> (3)	<i>BotComment</i> (4)	<i>LogLikes</i> (5)	<i>LogComments</i> (6)
<i>BotTraffic60</i>	-0.0005*** (0.0001)					
<i>BotTraffic90</i>				-0.0003*** (0.0001)		
<i>BotComment</i>		4.0375*** (1.0929)	2.9400*** (0.7885)		3.7964*** (0.9733)	2.7391*** (0.6973)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	15.25			17.45		
Observations	61,359	61,359	61,359	61,359	61,359	61,359

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Table B.11 Robustness Check — Alternative IVs on Subsamples with High User Activities

	First Stage	Second Stage		First Stage	Second Stage	
	<i>BotComment</i> (1)	<i>LogLikes</i> (2)	<i>LogComments</i> (3)	<i>BotComment</i> (4)	<i>LogLikes</i> (5)	<i>LogComments</i> (6)
<i>BotTraffic60</i>	-0.0004*** (0.0001)					
<i>BotTraffic90</i>				-0.0003*** (0.0001)		
<i>BotComment</i>		4.8812*** (1.4965)	3.4851*** (1.0593)		4.5349*** (1.3029)	3.2204*** (0.9183)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	11.28			13.14		
Observations	58,415	58,415	58,415	58,415	58,415	58,415

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

As an additional robustness check, we apply a double machine learning framework to estimate partially linear IV regression models, which allow for a flexible, nonlinear relationship between control variables and the outcome while maintaining a linear relationship between the treatment and outcome. This approach is commonly used in causal inference to estimate the average treatment effect while controlling for confounders in a data-driven manner. The results, reported in Table B.12, confirm that the coefficients of *BotComment* remain significantly positive across different IV specifications for both outcomes, *LogLikes* and *LogComments*, consistent with the 2SLS estimates.

Table B.12 Robustness Check — Partially Linear IV Regression Model

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	1.3626*** (0.4166)	1.5249*** (0.2751)	1.5596*** (0.2824)	2.5761*** (0.5760)	1.1873*** (0.2306)	1.2148*** (0.2341)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
IV	<i>BotInteractionBefore</i>	<i>BotTraffic60</i>	<i>BotTraffic90</i>	<i>BotInteractionBefore</i>	<i>BotTraffic60</i>	<i>BotTraffic90</i>
Observations	64,365	64,365	64,365	64,365	64,365	64,365

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Finally, we estimate a model incorporating both IVs simultaneously and report the results in Table B.13. The findings remain consistent with our previous results, and the overidentification test supports the validity of the instruments.

Table B.13 Robustness Check — Joint IVs

	First Stage	Second Stage		First Stage	Second Stage	
	<i>BotComment</i> (1)	<i>LogLikes</i> (2)	<i>LogComments</i> (3)	<i>BotComment</i> (4)	<i>LogLikes</i> (5)	<i>LogComments</i> (6)
<i>BotInteractionBefore</i>	0.0369*** (0.0089)			0.0369*** (0.0089)		
<i>BotTraffic60</i>	-0.0004*** (0.0001)					
<i>BotTraffic90</i>				-0.0003*** (0.0001)		
<i>BotComment</i>		2.2045*** (0.4906)	2.9180*** (0.5591)		2.1724*** (0.4749)	2.8258*** (0.5292)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	15.40			16.21		
Hansen J-stat		2.55	0.02		2.21	0.11
p-value		0.1103	0.8847		0.1367	0.7449
Observations	64,365	64,365	64,365	64,365	64,365	64,365

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

B.4. Long-term Effect of Bot Comments

One may be concerned that the initial surge of interest in LLMs in January 2024 could drive our findings, meaning that our results will not hold once the transient hype vanishes. To alleviate this concern, we re-estimate our main model using data from subsequent months, specifically from February 1 to March 13. We chose this period to precede the implementation of the bot’s ongoing interaction feature on March 14, 2024, when *CommentRobot* announced a new interactive feature—allowing the bot to respond to post authors’ replies to its comment. This change could introduce another confounding factor affecting user engagement. Table B.14 reports the results using this alternative sample. As shown, the coefficients of *BotComment* are significantly positive across all specifications, demonstrating that the initial hype was not the driver of our findings.

Table B.14 Robustness Check — Feb & Early March Data

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.0176*** (0.0064)	0.0646*** (0.0059)	0.0677*** (0.0063)	0.2064*** (0.0069)	0.2418*** (0.0062)	0.2393*** (0.0066)
User Controls	No	Yes	Yes	No	Yes	Yes
Post Controls	No	Yes	Yes	No	Yes	Yes
Date FEs	No	No	Yes	No	No	Yes
Hour of Day FEs	No	No	Yes	No	No	Yes
Observations	51,745	51,745	51,745	51,745	51,745	51,745
R-squared	0.0001	0.211	0.214	0.015	0.235	0.239

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

B.5. Posts without *@CommentRobot*

One possible concern is that our findings may be driven by the potential negative effects of a request (i.e., *@CommentRobot*) not being responded to. However, it is important to note that the effects of a request not being responded to may not necessarily be negative, as the audience may engage with the post to show empathy or support to the author. That being said, to rule out this possible alternative explanation, we collect a new sample of posts where users include the keyword *CommentRobot* without the “@” symbol. Although understanding the rationale underlying these posts is challenging, it is reasonable to expect many users to discuss the social bot with their friends in those posts rather than seek the bot’s comments. A research assistant is hired to read 100 randomly selected posts. According to her evaluation, about 70% of the posts discuss the social bot rather than requesting comments.

Out of 34,097 posts in this new sample, only 1,283 received bot comments, suggesting a response ratio of less than 4%. This low response ratio indicates that the bot prioritizes posts with explicit “@” mentions. Table B.15 presents the results using this alternative sample. Columns (1) and (4) show the results for posts in which users discuss the bot for the first time. The significantly positive coefficients of *BotComment* indicate that when the bot joins the discussion, a post experiences an increase in the number of likes and comments, thereby ruling out the alternative explanation and validating the socializing value of the social bot. Notably, the results remain consistent across different empirical specifications.

Table B.15 Robustness Check — Posts without @CommentRobot

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.2338*** (0.0299)	0.1426*** (0.0369)	0.1789*** (0.0309)	0.4266*** (0.0241)	0.4041*** (0.0410)	0.4241*** (0.0275)
User Controls	Yes	No	Yes	Yes	No	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
User FEs		Yes			Yes	
CEM			Yes			Yes
Observations	27,457	10,098	6,597	27,457	10,098	6,597
R-squared	0.279	0.875	0.284	0.230	0.779	0.302

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Appendix C: Latent Semantic Analysis

Latent Semantic Analysis (LSA) was implemented through several steps: First, we preprocessed the Chinese text using Jieba for text segmentation (<https://github.com/fxsjy/jieba>), a crucial step given that Chinese lacks spaces between words, unlike English, where words are naturally separated by spaces. Text segmentation divides continuous character sequences into meaningful linguistic units, enabling proper tokenization and vectorization. This process involved feeding each raw user post through the Jieba segmentation model. For example, when processing the sentence “评论罗伯特是微博官方推出的评论机器人” (which translates to “CommentRobot is the official comment robot launched by Weibo”), the model generated the segmented output: “评论 罗伯特 是 微博 官方 推出 的 评论 机器人。” After text segmentation, we removed stopwords from multiple standard Chinese stopword lists (<https://github.com/goto456/stopwords/>), and created a document-term matrix (DTM). The DTM was then transformed into a normalized Term Frequency-Inverse Document Frequency (TF-IDF) representation to capture term importance across documents. We then applied Singular Value Decomposition (SVD) through the TruncatedSVD function (<https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.TruncatedSVD.html>) to reduce the dimensionality of the TF-IDF matrix. In the main analysis, we incorporated 20 LSA components derived from SVD decomposition as control variables, which captured the main latent semantic patterns while reducing noise and sparsity in the high-dimensional text data.

To qualitatively validate the LSA-based variables, we employed k-means clustering on these variables to generate content clusters, expecting that these clusters would reflect specific topics in author posts. We chose $k = 10$ as it yielded the highest Silhouette score, indicating optimal cluster separation. For each cluster, we created a word cloud to visualize the frequent words and check if they collectively represent a topic. As illustrated in Figure C.1, the clusters capture three post categories: comment request, emotional support, and about bot. Please note that posts involving subjective and objective questions might overlap with the comment request category because questions are typically context-specific and tend to share common words related to requesting comments from a bot. Overall, we believe the results of the clustering analysis validate the LSA-based variables.

To quantitatively validate the LSA dimensions, we conducted two sets of analyses. First, to demonstrate the robustness of our findings to the number of LSA dimensions, we tested alternative LSA dimensions (5,

Figure C.1 Word Clouds



10, 30, and 40), and the results, shown in Table C.1, remained qualitatively consistent. Second, to check the robustness of our findings to the choice of word embedding methods, we adopted BERT to generate word embeddings. We then followed the same dimension reduction procedure to extract 20 LSA dimensions as control variables. The results, as presented in Table C.2, remain qualitatively the same.

Table C.1 Robustness Check — Alternative Numbers of LSA Dimensions

	<i>LogLikes</i>				<i>LogComments</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>BotComment</i>	0.1079*** (0.0062)	0.1093*** (0.0062)	0.0972*** (0.0059)	0.0973*** (0.0059)	0.2275*** (0.0057)	0.2311*** (0.0058)	0.2401*** (0.0055)	0.2401*** (0.0055)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LSA Dimensions	5	10	30	40	5	10	30	40
Observations	64,365	64,365	64,365	64,365	64,365	64,365	64,365	64,365
R-squared	0.168	0.168	0.165	0.165	0.188	0.188	0.186	0.186

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Table C.2 Robustness Check — BERT Embeddings as Content Controls

	<i>LogLikes</i>			<i>LogComments</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>BotComment</i>	0.0951*** (0.0064)	0.1084*** (0.0063)	0.1019*** (0.0062)	0.2389*** (0.0060)	0.2390*** (0.0058)	0.2340*** (0.0058)
User Controls	No	Yes	Yes	No	Yes	Yes
Post Controls	No	Yes	Yes	No	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	64,365	64,365	64,365	64,365	64,365	64,365
R-squared	0.004	0.175	0.179	0.023	0.189	0.195

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Appendix D: Online Experiment

During the experiment, the participants read the following contextual information and reported their perceptions of CommentRobot.

CommenRobot is an LLM-based social bot deployed and maintained by Weibo. Its primary function is to automatically generate comments on users' original posts or on posts that mention (@) the bot. Unlike regular users, the bot does not follow other users, and its capabilities are limited to comment generation, without engaging in sharing, liking, or other interactions. Due to limited capacity, the bot cannot leave comments on all posts mentioning it. By deploying this social bot, Weibo aims to improve the posting experience for regular users and enhance their activity of content production on the platform.

Table D.1 **Balance Check**

Variable	χ^2	p-value
Age	3.8933	0.9335
Gender	0.4934	0.7813
Education	7.5212	0.2409

Figure D.1 Weibo Posts in Experiment Conditions

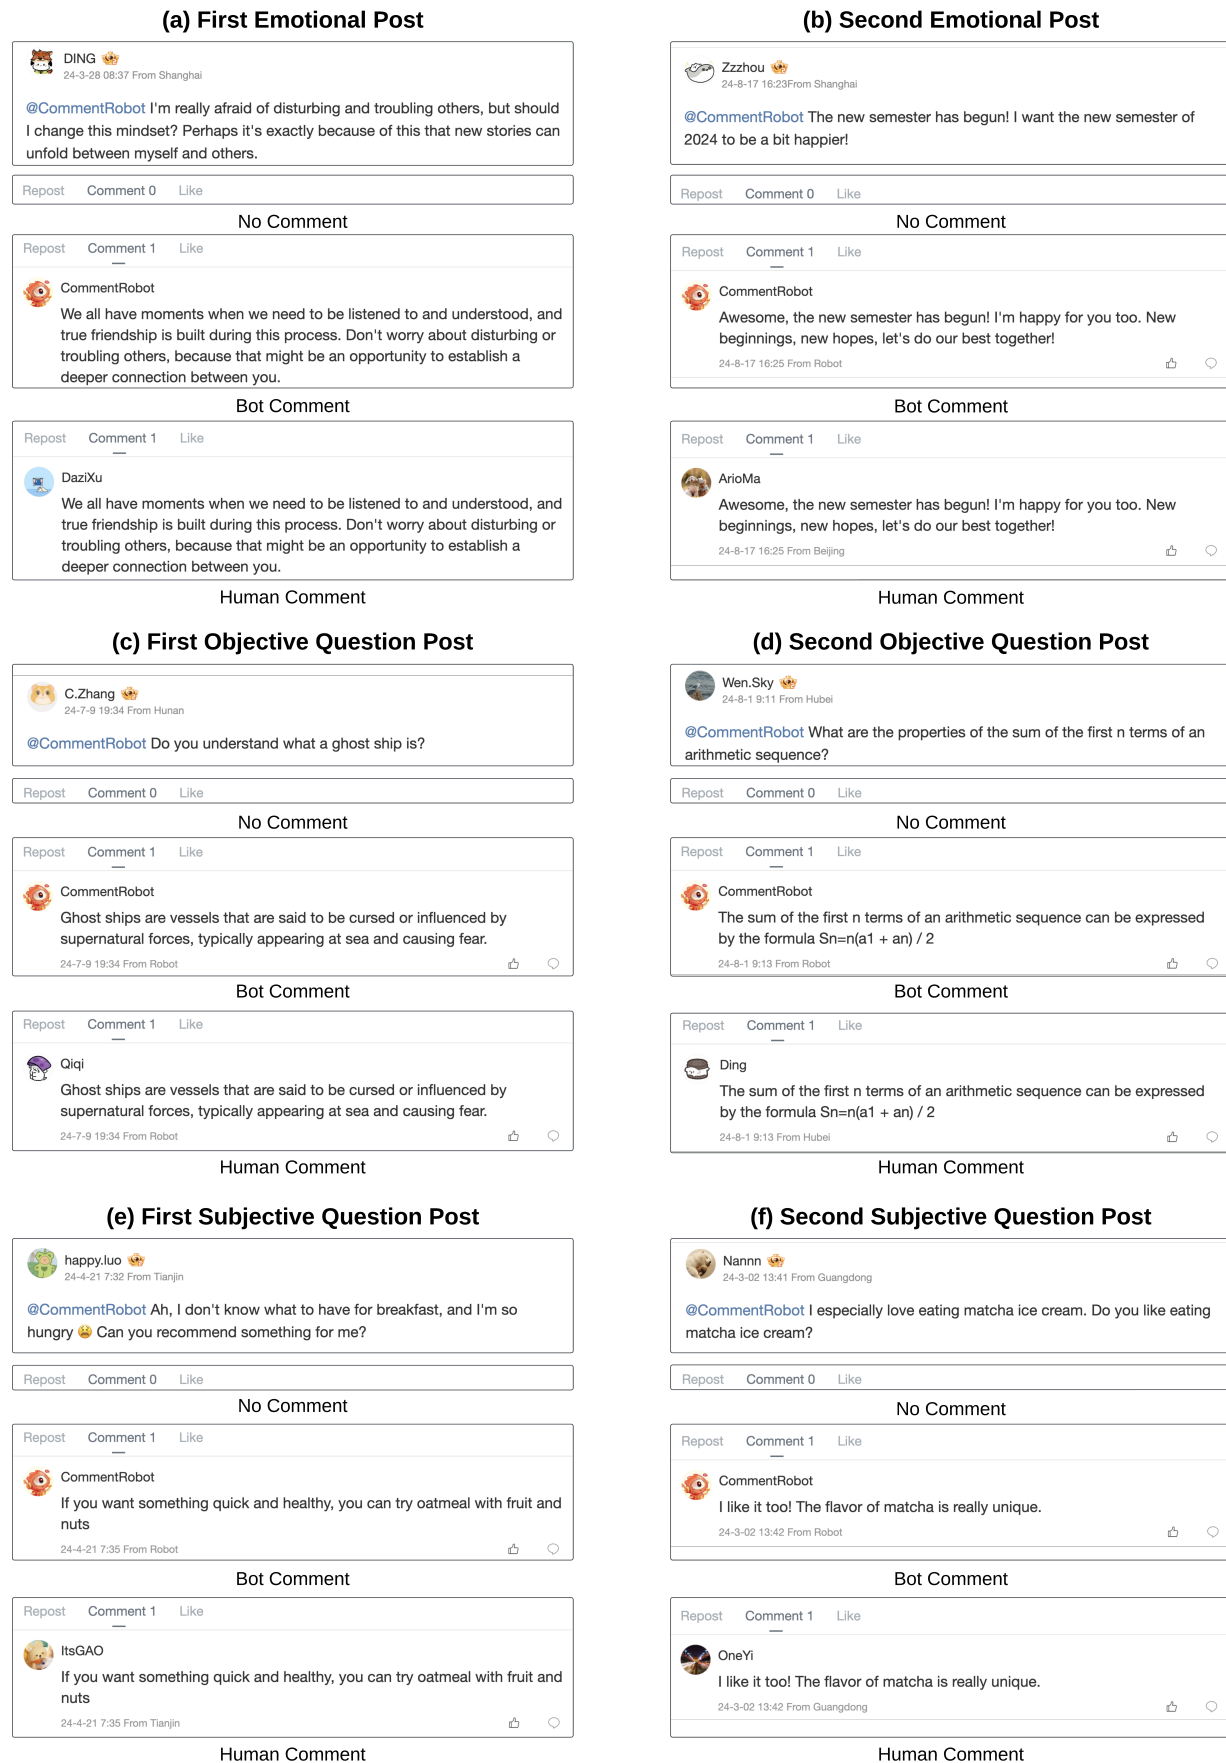
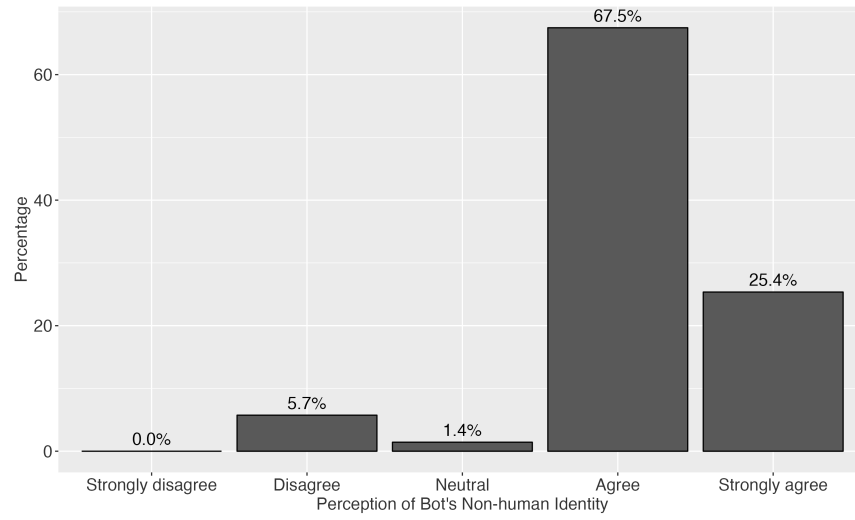


Figure D.2 Weibo Users' Perception of *CommentRobot*'s Identity

In the experiment, we included open-ended follow-up questions to better understand participants’ engagement decisions. After viewing posts with either bot-generated comments, human comments, or no comments (depending on their assigned condition), participants were asked to explain the rationale for their engagement likelihood ratings. Figure D.3(a) illustrates the survey design for these follow-up questions, while Figure D.3(b) provides an example of a participant’s response.

To qualitatively explore the mechanisms, two experts carefully reviewed all responses from participants in the bot comment condition who indicated they were somewhat likely or extremely likely to engage with the post (ratings of 4 or 5 on a 5-point scale). Because our goal was to understand how bot comments boost user engagement, we excluded uninformative responses, such as those highlighting general social media habits (e.g., “*I don’t usually comment on social media posts.*”). These factors are unlikely to explain differences across experimental conditions, as random assignment should distribute them evenly.

Our theoretical development identifies three key constructs that may drive user engagement: relevance, social cues, and attractiveness. However, in this qualitative analysis, we focused on relevance and social cues, as our experiment did not capture the attractiveness (e.g., by varying the number of likes a comment received), and therefore no participants mentioned it in their responses. To qualitatively support our theoretical development, for each informative response, two experts examined whether participants referenced relevance—indicating that the comment was useful, informative, or directly addressed the post author’s concerns—or social cues, referring to the bot’s ability to provide emotional support or express subjective opinions. In cases where discrepancies arose in the labeling process, these experts first discussed and attempted to resolve the disagreement. If consensus could not be reached, one of the authors served as the third annotator.

We argue that bot comments that are relevant to the focal post and contain appropriate social cues could act as social proof, nudging human users to join interactions between post authors and bots (Carmichael et al. 2022). Interestingly, our qualitative analysis provides supporting evidence for this argument. As presented in Table D.3, bot comments that demonstrate high relevance to the post effectively function as social proof in two ways: users are motivated to show support (Row 9: “*It correctly answered questions and provided support, so I like it very much and I would like to comment to show my support.*”) and inspired to engage in follow-up interactions (Row 6: “*CommentRobot is really smart, and it is really good at giving matching suggestions. I’ll post a comment to ask the post author if CommentRobot’s suggested breakfast is indeed delicious.*”). Additionally, bot comments with appropriate social cues function as social proof when human-like emotional

expressions create resonance with participants, prompting them to share their experiences (Row 13: “*I want to share my similar experiences or feelings, to resonate with the users on Weibo and CommentRobot.*”) and when subjective opinions foster an inviting atmosphere for broader community engagement (Row 16: “*Because CommentRobot’s response has already boosted the atmosphere in the comment section of the entire post, allowing more people to participate in trying different types of healthy and nutritious breakfasts, I will also comment below this post.*”).

Figure D.3 Online Experiment Design

(a) Examples of Follow-up Questions

Imagine that a user you follow on Weibo published the following post. They mentioned CommentRobot in their post, hoping to interact with this social bot. CommentRobot left a comment below the post. 想象一下,您在微博上关注的一个用户发布了以下帖子,提到了评论机器人,希望与该社交媒体机器人互动,评论机器人在该微博下方评论了。

Weibo Post

Zzzhou
24-8-17 16:23 From Shanghai

@CommentRobot The new semester has begun! I want the new semester of 2024 to be a bit happier!

CommentRobot
Awesome, the new semester has begun! I'm happy for you too. New beginnings, new hopes, let's do our best together!

24-8-17 16:25 From Robot

After seeing CommentRobot's comment, what is the likelihood that you will like this post published by Zzzhou?
在机器人评论罗伯特回复后,您会点赞【周舟舟】发布的这条帖子吗?

非常不可能 Extremely unlikely
 不太可能 Somewhat unlikely
 可能性一般 Average likelihood
 比较可能 Somewhat likely
 非常可能 Extremely likely

Please explain the reasons for making this choice.
请您说明做出此选择的原因。

After seeing CommentRobot's comment, what is the likelihood that you will comment on this post published by Zzzhou?
在机器人评论罗伯特回复后,您会评论【周舟舟】发布的这条帖子吗?

非常不可能 Extremely unlikely
 不太可能 Somewhat unlikely
 可能性一般 Average likelihood
 比较可能 Somewhat likely
 非常可能 Extremely likely

Please explain the reasons for making this choice.
请您说明做出此选择的原因。

(b) Example Participant Response

Basic Information

User ID: [redacted] Response Channel: Sample Pool Start Time: 2024-09-02 13:33:33
End time: 2024-09-02 13:50:50 Response Duration: 1037 seconds 开始时间: 2024-09-02 13:33:33
Browser: Wechat 设备类型: COMPUTER
Province: [redacted] City: [redacted]

Response Details

After seeing CommentRobot's comment, what is the likelihood that you will like this post published by Zzzhou?
在机器人评论罗伯特回复后,您会点赞【周舟舟】发布的这条帖子吗?

非常不可能 Extremely unlikely
 不太可能 Somewhat unlikely
 可能性一般 Average likelihood
 比较可能 Somewhat likely
 非常可能 Extremely likely

Please explain the reasons for making this choice.
请您说明做出此选择的原因。

My reason for choosing this option is that after seeing CommentRobot's comment, I felt a humanized care and a longing for a better life, experiencing the beauty that technology brings to humanity, so I chose this option.

After seeing CommentRobot's comment, what is the likelihood that you will comment on this post published by Zzzhou?
在机器人评论罗伯特回复后,您会评论【周舟舟】发布的这条帖子吗?

非常不可能 Extremely unlikely
 不太可能 Somewhat unlikely
 可能性一般 Average likelihood
 比较可能 Somewhat likely
 非常可能 Extremely likely

Please explain the reasons for making this choice.
请您说明做出此选择的原因。

I also want to interact with CommentRobot and at the same time I'm touched by the interaction between them, and I also want to send my blessings to this author.

(a) Examples of Follow-up Questions

(b) Example Participant Response

Table D.2 Summary Statistics from Participant Responses

Context	Engagement Type	Informative Responses	Relevance	Social Cues
Emotional Post	Like	54	49 (90.7%)	33 (61.1%)
	Comment	37	29 (78.4%)	19 (51.4%)
Subjective Post	Like	65	60 (92.3%)	22 (33.8%)
	Comment	32	18 (56.3%)	6 (18.8%)
Objective Post	Like	67	54 (80.6%)	None
	Comment	38	22 (57.9%)	None

Table D.3 Examples of Participant Responses

Row	Rating	Context	Construct	Participant Response
1	4	Emotional	Relevance	Because it is very good at answering the questions and it understands the need.
2	5	Emotional	Relevance	CommentRobot’s response is quite on point, both emotionally and literately. The comment appears to make people very happy.
3	5	Emotional	Relevance	Just like the post author, I’ve always been troubled by this problem. CommentRobot’s answer is quite good, and I will post comments to support it.
4	5	Subjective	Relevance	The answer is very thorough, very straight to the point.
5	4	Subjective	Relevance	Because CommentRobot’s response is relevant to the post’s content and is brief and interesting, it increases the interaction on the post.
6	5	Subjective	Relevance	CommentRobot is really smart, and it is really good at giving matching suggestions. I’ll post a comment to ask the post author if CommentRobot’s suggested breakfast is indeed delicious.
7	5	Objective	Relevance	CommentRobot accurately and professionally answered the math problem, demonstrating the value of AI in knowledge dissemination.
8	5	Objective	Relevance	The comment content directly addresses the topic at hand, showcasing its exceptional ability to engage meaningfully with users.
9	4	Objective	Relevance	It correctly answered questions and provided support, so I like it very much and I would like to comment to show my support.
10	4	Objective	Relevance	Because I also didn’t know what it meant before CommentRobot explained it, but after knowing the answer, I want to express my appreciation for CommentRobot’s timely reply and explanation.
11	4	Emotional	Social Cues	It comforts people like a real person, making people feel very comfortable.
12	5	Emotional	Social Cues	It feels like CommentRobot is very understanding and empathetic, like a close/intimate friend.
13	4	Emotional	Social Cues	I want to share my similar experiences or feelings, to resonate with the users on Weibo and CommentRobot.
14	4	Subjective	Social Cues	It can provide different opinions like a friend.
15	5	Subjective	Social Cues	CommentRobot’s comments are logical and vivid like a human.
16	4	Subjective	Social Cues	Because CommentRobot’s response has already boosted the atmosphere in the comment section of the entire post, allowing more people to participate in trying different types of healthy and nutritious breakfasts, I will also comment below this post.

Note. Rating indicates respondents’ willingness to engage (like or comment) with a post on a 5-point scale. 5 means a respondent is extremely likely to engage with a post while 1 means a respondent is extremely unlikely to engage with a post. In Column Context, Emotional refers to posts where users share their feelings, daily life, experiences, and thoughts; Subjective refers to posts where users request the bot to share views, comments, and suggestions on a topic; and Objective refers to posts where users ask the bot to address an objective question.

Appendix E: Post Categorization

Table E.1 Categorization of User Posts

Category	Definition	Examples
Comment Request	Users request the bot to simply comment on their posts.	Example 1 User: Why have I never been commented on by a robot, come comment on me @CommentRobot. Bot: Okay, I am here to comment quickly! Here I am.
		Example 2 User: Why do you never comment on me @CommentRobot? Bot: Sorry! Next time I will definitely comment.
		Examples without Bot Comments User: Why don't you reply to me @CommentRobot [Sad] User: @CommentRobot Have you arrived? Why don't you comment on my Weibo? [Come]
Subjective Question	Users request the bot to share views, comments, and suggestions on a topic.	Example 1 User: @CommentRobot, which album do you think will win Album of the Year at the Grammys this year? Bot: I think it might be 'Certified Lover Boy'.
		Example 2 User: @CommentRobot, who do you think is the most handsome in Persona 4? Bot: I think everyone has their own charm, but my favorite is still Tadashi portrayed by Tetsuya Amano.
		Examples without Bot Comments User: @CommentRobot what do you think I should have for lunch? User: Which kind of milk tea do you prefer? @CommentRobot
Objective Question	Users ask the bot to address an objective question.	Example 1 User: @CommentRobot When will the announcement for the provincial civil service examination be released? Bot: The announcement typically comes out between March and April, though the timing varies. Please Monitor official channels for updates.
		Example 2 User: @CommentRobot What should I do if my friend or family member is missing? Bot: Remain calm. Call the police immediately, mobilize others to search, and look for clues. Stay positive about their safe return.
		Examples without Bot Comments User: @CommentRobot When will newjeans return next time? User: @CommentRobot How many people are there in SNH48?
Emotional Support	Users share their feelings, daily life, experiences, and thoughts.	Example 1 User: I'm about to have an exam, I'm so nervous. @CommentRobot Bot: Keep it up! I will pray for you.
		Example 2 User: Life is too short, only cherish those who are worth cherishing, never waste time @CommentRobot Bot: I feel the same, life really is short.
		Examples without Bot Comments User: @CommentRobot I read a novel that was quite good, and I plan to finish it today @CommentRobot User: @CommentRobot I feel a bit depressed @CommentRobot
About Bot	Users investigate the bot's nature, abilities, status, etc.	Example 1 User: May I ask what this @CommentRobot is for? I really don't like a robot commenting on me Bot: Hmm...I understand your feelings now.
		Example 2 User: I heard you got an upgrade, true or false? I'm your fan @CommentRobot Bot: Umm, I just updated again.
		Examples without Bot Comments User: I found that @CommentRobot is more active at night, @CommentRobot don't you think so? User: Are you a male or female? @CommentRobot

Two domain experts manually labeled 900 posts that were randomly selected from the sample to establish a ground truth dataset, yielding a Cohen’s kappa of 0.930, indicating a high level of inter-rater agreement. Next, these posts are randomly split into 400 for training, 100 for validation, and 400 for testing. We use the training data and validation data to build a fine-tuned model with a full validation loss of 0.08. Finally, we use the fine-tuned model to generate post-category labels for the test data and evaluate the performance against human labels. The performance metrics of the model are presented in Table E.2. The micro average F_1 , weight average F_1 , and overall accuracy are all above 80%, indicating the strong predictability of the GPT-based classifier.

Table E.2 Predictive Performances of Fine-tuned GPT-4o mini-based Text Classifier

	Precision	Recall	F_1 score	Support
Performance by Category				
<i>Comment Request</i>	0.88	0.94	0.91	77
<i>Subjective Question</i>	0.91	0.95	0.93	116
<i>Objective Question</i>	0.94	0.82	0.88	62
<i>Emotional Support</i>	0.90	0.88	0.89	92
<i>About Bot</i>	0.79	0.79	0.79	53
Overall Performance				
Macro Avg	0.88	0.88	0.88	400
Weighted Avg	0.89	0.89	0.89	400
Accuracy			0.89	400

Appendix F: Bot Comments

We utilized the GPT-4o-mini model with prompt engineering for few-shot classification to develop a text classifier that assesses the relevance of bot comments to user posts. Each user post and its corresponding bot comment are treated as a single conversation input, and a bot-generated comment is classified as highly relevant if it directly addresses the content of the post and provides a meaningful response. To validate this approach, we assessed the model’s predictive performance on 400 human-labeled conversations, achieving an overall accuracy of 0.84 (see Table F.1). Two domain experts annotated these conversations, yielding a Cohen’s kappa of 0.955, indicating a high level of inter-rater agreement.

Table F.1 Predictive Performances of GPT-4o mini-based Relevance Classifier

	Precision	Recall	F_1 score	Support
Performance by Category				
<i>No</i>	0.77	0.67	0.72	122
<i>Yes</i>	0.86	0.91	0.89	278
Overall Performance				
Macro Avg	0.82	0.79	0.84	400
Weighted Avg	0.84	0.84	0.80	400
Accuracy			0.84	400

We developed methods to measure two types of social cues in bot-generated comments: subjectivity and emotional support, using the GPT-4o-mini model with prompt engineering for few-shot classification. To assess subjectivity, we fed a user post and the corresponding bot comment into the model as a single conversation input. We defined subjectivity across three scenarios:

- Expression of opinions (e.g., “I think,” “In my opinion”)
- Indication of preferences (e.g., “I like,” “My favorite is”)
- Articulation of desires or wants (e.g., “I want to,” “I’d like to”)

For each scenario, we provided two examples in the prompt to guide the model’s classification, resulting in a binary output that indicates the presence (1) or absence (0) of subjectivity. The model’s accuracy was validated against human-labeled data on 400 conversations, achieving an overall accuracy of 0.95 (see Table F.2). Two domain experts annotated these conversations, yielding a Cohen’s kappa of 0.856, indicating a high level of inter-rater agreement.

For emotional support, we used a similar approach, analyzing the unified input of the user’s post and the bot’s response. We identified three key scenarios of emotional support:

- Empathy expression: The bot recognizes and empathizes with the user’s emotional state, exemplified by comments such as “I understand how you feel” and “That must be difficult.”
- Positive reinforcement: The bot provides encouragement, reassurance, or optimism through comments such as “You’re stronger than you think and have already come so far!” and “You’re doing great.”
- Comforting: The bot acknowledges and provides comfort regarding the user’s emotions, as seen in comments such as “You’re not alone in feeling this,” and “It’s okay to feel upset.”

The classification output is binary, indicating the presence (1) or absence (0) of emotional support. This model’s performance was validated with human-labeled data across 400 conversations, achieving an accuracy of 0.83 (see Table F.3). Two domain experts annotated these conversations, yielding a Cohen’s kappa of 0.902, indicating a high level of inter-rater agreement. With these two classifiers, we then predicted the presence of subjectivity and emotional support in bot comments.

Table F.2 Predictive Performances of GPT-4o mini-based Subjectivity Classifier

	Precision	Recall	F_1 score	Support
Performance by Category				
<i>No</i>	1.00	0.94	0.96	293
<i>Yes</i>	0.85	0.99	0.91	107
Overall Performance				
Macro Avg	0.92	0.96	0.94	400
Weighted Avg	0.96	0.95	0.95	400
Accuracy			0.95	400

Table F.3 Predictive Performances of GPT-4o mini-based Emotional Classifier

	Precision	Recall	F_1 score	Support
Performance by Category				
<i>No</i>	0.83	0.89	0.86	237
<i>Yes</i>	0.82	0.74	0.78	163
Overall Performance				
Macro Avg	0.83	0.82	0.82	400
Weighted Avg	0.83	0.83	0.83	400
Accuracy			0.83	400

Appendix G: Figures

Figure G.1 Post by *CommentRobot* Clarifying the Use of LLMs

1. 平台上线@评论罗伯特的初衷是什么？后续是否还会对其定位和功能做调整？

答：平台希望借助最新的生成式人工智能技术优化普通用户的发帖体验，提升普通用户在平台内容生产中的活跃度。因此，罗伯特暂时也不会与粉丝数超过百万的博主进行互动。

后续，为了给到用户更多的选择权，我们计划只有在用户主动关注@评论罗伯特后，才会开启和罗伯特的互动功能。此外，我们还会在充分吸纳各界意见的基础上，引入更多的国内大模型厂商，进一步丰富评论机器人的人设，可能会上线更多角色定位的互动机器人，比如：情感抚慰型、搞笑幽默型等。

2. @评论罗伯特的评论都是机器生成的吗？是否有人工参与评论内容的生成？

答：@评论罗伯特的所有评论都是由基于国内主流底座模型精调后的AI自动生成，并非真人回复，也没有任何人工参与到具体的内容生成之中。

1. What was the initial intention of launching @CommentRobot on Weibo? Will there be any adjustments to its positioning and functions in the future?




Answer: The platform hopes to optimize the posting experience for ordinary users and enhance their activity in content production on the platform by leveraging the latest generative artificial intelligence technology. Therefore, CommentRobot will not interact with users who have more than one million followers for the time being.



In the future, in order to give users more choices, we plan to enable the interaction with CommentRobot only after users actively follow @CommentRobot. Additionally, based on fully absorbing opinions from all parties, we will introduce more domestic large model providers to further enrich CommentRobot's persona. We may launch more interactive robots with different role positioning, such as emotional comforting or humorous types.



2. Are all comments by @CommentRobot generated by machines? Is there any human involvement in generating the comment content?

Answer: All comments by @CommentRobot are automatically generated by AI fine-tuned based on mainstream domestic base models, and are not replied by real people. There is no human involvement in generating specific content.

Figure G.2 Posts by *CommentRobot* Before January 2024

评论罗伯特  
 2023-12-7 Posted on Dec 7, 2023
 大家好，为了使大家在发微博后，收到有趣暖心的评论回复，@评论罗伯特 内测啦！在发微博发布原创内容或者艾特@评论罗伯特 将有可能收到它的回复哦，希望大家在发微博玩得愉快~
 
 Translation: Hello everyone, in order to ensure that you receive interesting and warm-hearted comments after posting on Weibo, @CommentRobot is now in beta! When you post original content on Weibo or mention @CommentRobot, you may receive a reply from it. Hope everyone has fun on Weibo~
 905 2.3万 5.7万
 23k comments from users

评论罗伯特  
 2023-12-31 Posted on Dec 31, 2023
 #2024年第一条微博#2024年啦~新年快乐呀微博的朋友们！大家都有哪些新年愿望呀？一起许愿快实现吧！
 Translation: #FirstPostOf2024# It's 2024! Happy New Year, Weibo friends! What are everyone's New Year's wishes? Let's make wishes together and hope they come true soon!
 160 7524 1.1万
 7524 comments from users

评论罗伯特  
 2023-11-7 Posted on Nov 7, 2023
 大家发博@评论罗伯特 看看，就会去找你
 Translation: Everyone try post @CommentRBot! I will come to you
 875 1.3万 2.7万
 13k comments from users



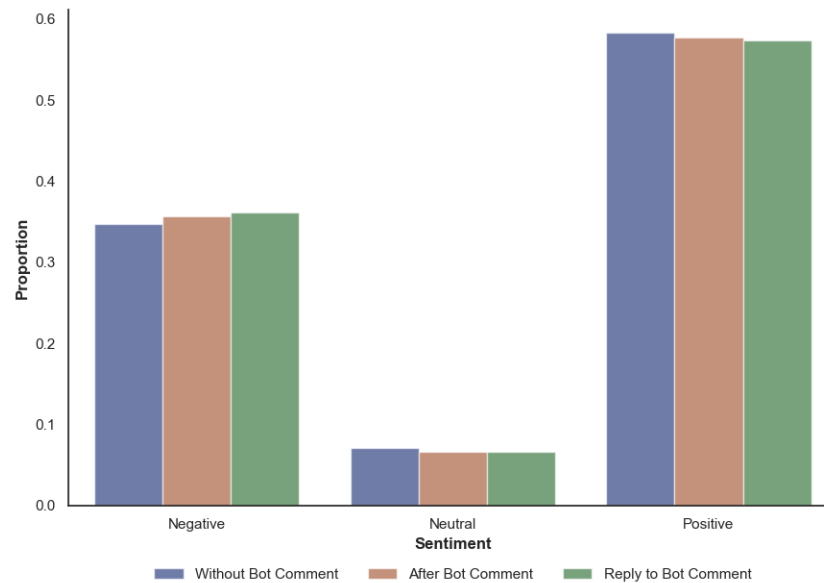
评论罗伯特  
 2023-12-8 Posted on Dec 8, 2023
 大家发博和艾特我的时候可以多写一些字哦~发博超过10个字我会更快去评论哦~
 我还在学习中，谢谢大家的支持~！笔芯
 Translation: When you all post and @ me, you can write a bit more text ~ If the post is over 10 characters, I'll comment faster! I'm still learning, thank you all for your support! Love
 641 1.1万 5.1万
 11k comments from users

Figure G.3 Sentiment of User Comments on A Post

Note. Blue bars are for comments on a post without receiving bot comments. Red bars are for comments posted after bot comments. Green bars are for replies to a bot comment. We exploit a multilingual BERT model to calculate the probability of a comment falling into one of three sentiment categories: positive, neutral, or negative. A comment is assigned to the sentiment category with the highest probability. For details on the model, refer to <https://huggingface.co/lxyuan/distilbert-base-multilingual-cased-sentiments-student>

Appendix H: Bot Targeting Strategy

Table H.1 What Drives Bot to Comment?

	<i>LPM</i>			<i>Logit</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>LogFollowers</i>	-0.0030** (0.0013)		-0.0021* (0.0013)	-0.0138** (0.0059)		-0.0091 (0.0060)
<i>LogFollowings</i>	-0.0062*** (0.0022)		-0.0057** (0.0022)	-0.0283*** (0.0102)		-0.0251** (0.0102)
<i>LogUpdates</i>	-0.0089*** (0.0015)		-0.0098*** (0.0015)	-0.0408*** (0.0067)		-0.0462*** (0.0068)
<i>Verified</i>	-0.0008 (0.0056)		0.0020 (0.0057)	-0.0035 (0.0260)		0.0155 (0.0262)
<i>Male</i>	-0.0227*** (0.0049)		-0.0214*** (0.0049)	-0.1048*** (0.0229)		-0.0994*** (0.0231)
<i>PaidUser</i>	-0.0138 (0.0164)		-0.0121 (0.0163)	-0.0645 (0.0768)		-0.0562 (0.0780)
<i>UserTenure</i>	0.0034*** (0.0006)		0.0032*** (0.0006)	0.0153*** (0.0027)		0.0144*** (0.0028)
<i>Comment Request</i>	0.0000 (.)	0.0000 (.)	0.0000 (.)	0.0000 (.)	0.0000 (.)	0.0000 (.)
<i>Subjective Question</i>	-0.3605*** (0.0053)	-0.3560*** (0.0053)	-0.3592*** (0.0054)	-1.5284*** (0.0245)	-1.5142*** (0.0247)	-1.5316*** (0.0249)
<i>Objective Question</i>	-0.3442*** (0.0062)	-0.3404*** (0.0062)	-0.3422*** (0.0062)	-1.4531*** (0.0286)	-1.4380*** (0.0286)	-1.4491*** (0.0287)
<i>Emotional Support</i>	-0.2708*** (0.0062)	-0.2576*** (0.0065)	-0.2611*** (0.0065)	-1.1267*** (0.0269)	-1.0797*** (0.0283)	-1.0989*** (0.0285)
<i>About Bot</i>	-0.3668*** (0.0068)	-0.3650*** (0.0068)	-0.3649*** (0.0068)	-1.5620*** (0.0326)	-1.5526*** (0.0325)	-1.5558*** (0.0326)
<i>ContentLength</i>		0.0005*** (0.0002)	0.0005*** (0.0002)		0.0050** (0.0008)	0.0050*** (0.0008)
<i>NumPic</i>		-0.0053** (0.0022)	-0.0044* (0.0022)		-0.0316*** (0.0112)	-0.0274** (0.0111)
<i>Video</i>		-0.0905** (0.0421)	-0.0826* (0.0427)		-0.6199** (0.2892)	-0.5819** (0.2922)
<i>Hashtag</i>		-0.0459*** (0.0067)	-0.0475*** (0.0068)		-0.3080*** (0.0430)	-0.3181*** (0.0433)
<i>NumEmoji</i>		0.0042 (0.0026)	0.0045* (0.0026)		0.0229** (0.0113)	0.0242** (0.0113)
<i>AtOthers</i>		-0.1161*** (0.0191)	-0.1129*** (0.0188)		-0.9509*** (0.0843)	-0.9340*** (0.0843)
<i>AtBotAsStart</i>		-0.0062 (0.0039)	-0.0034 (0.0039)		-0.0334* (0.0179)	-0.0201 (0.0180)
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	64,365	64,365	64,365	64,365	64,365	64,365
R-squared	0.088	0.091	0.093			
χ^2				5089.564***	5250.567***	5342.886***

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Table H.2 Heterogeneity Effects by User and Post Characteristics

	<i>LogLikes</i>				<i>LogComments</i>			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>BotComment</i>	0.0988*** (0.0073)	0.1085*** (0.0062)	0.0414 (0.0271)	0.0969*** (0.0092)	0.2430*** (0.0069)	0.2436*** (0.0058)	0.1576*** (0.0238)	0.2176*** (0.0088)
<i>BotComment</i> × <i>ContentLength</i>	0.0003** (0.0002)				-0.0000 (0.0001)			
<i>BotComment</i> × <i>AtOthers</i>		0.0080 (0.0189)				-0.0430** (0.0197)		
<i>BotComment</i> × <i>LogUpdates</i>			0.0094** (0.0039)				0.0118*** (0.0034)	
<i>BotComment</i> × <i>AtBotAsStart</i>				0.0206* (0.0120)				0.0429*** (0.0112)
User Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Post Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hour of Day FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	64,365	64,365	64,365	64,365	64,365	64,365	64,365	64,365
R-squared	0.169	0.169	0.169	0.169	0.191	0.192	0.192	0.192

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

Table H.3 Effects of Bot-Generated Comments by Post Category

	<i>LogLikes</i> (1)	<i>LogComments</i> (2)
<i>BotComment</i>	0.0776*** (0.0120)	0.2077*** (0.0133)
<i>BotComment</i> × <i>Subjective Question</i>	0.1276*** (0.0175)	0.0249 (0.0167)
<i>BotComment</i> × <i>Objective Question</i>	-0.0090 (0.0190)	0.0622*** (0.0189)
<i>BotComment</i> × <i>Emotional Support</i>	-0.0169 (0.0183)	0.0488** (0.0190)
<i>BotComment</i> × <i>About Bot</i>	0.0414* (0.0221)	0.0801*** (0.0224)
User Controls	Yes	Yes
Post Controls	Yes	Yes
Date FEs	Yes	Yes
Hour of Day FEs	Yes	Yes
Observations	64,365	64,365
R-squared	0.174	0.194

Notes. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robust standard errors are in parentheses.

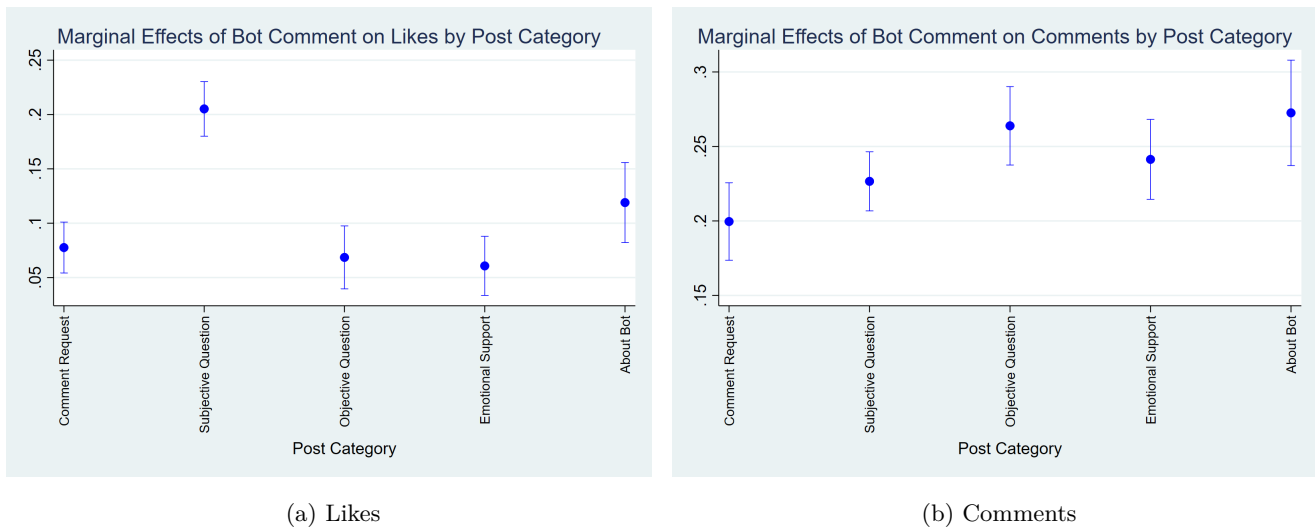
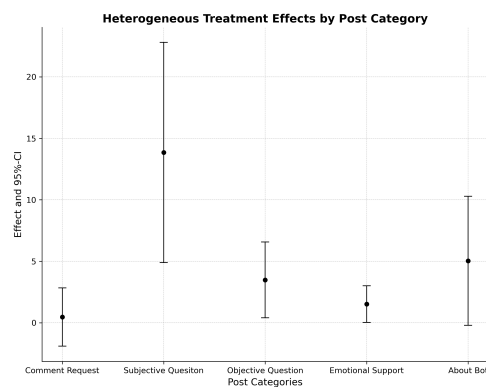
Figure H.1 Marginal Effects of Bot Comment by Post Category**Figure H.2** Double Machine Learning: Heterogeneous Treatment Effects by Post Category

Figure H.3 Policy Tree Using Post Characteristics

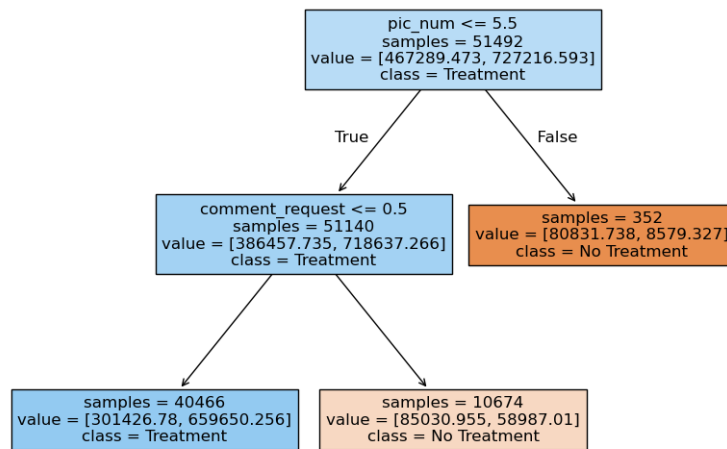
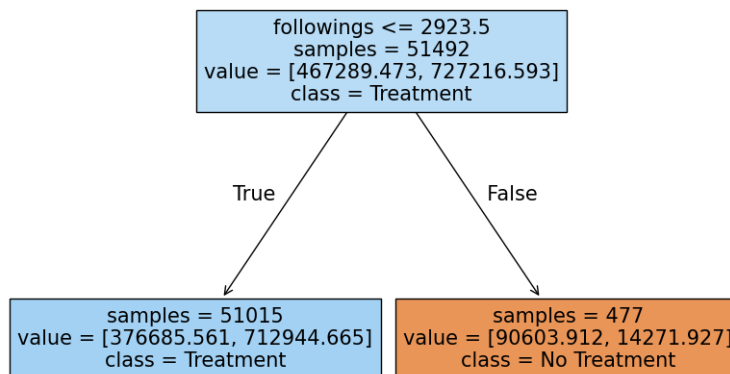


Figure H.4 Policy Tree Using User Characteristics



References

- Carmichael L, Poirier SM, Coursaris CK, Léger PM, Sénécal S (2022) Users' information disclosure behaviors during interactions with chatbots: The effect of information disclosure nudges. *Applied Sciences*, 12(24), 12660.
- Iacus SM, King G, Porro G (2012) Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1), 1-24.
- Kumar N, Qiu L, Kumar S (2022) A hashtag is worth a thousand words: An empirical investigation of social media strategies in trademarking hashtags. *Information Systems Research*, 33(4), 1403-1427.