Synthetic Empathy in Career Guidance Chatbots: A Knowledge Behavior Gap Model Study

Kristian Stoffregen[1[0009-0009-9965-4269]], Thong H. N. Dinh[2],
and Agnis Stibe[2,3,4,5[0000-0003-2653-7677]]

[1] Business Academy Copenhagen, Copenhagen, Denmark
[2] The Business School, RMIT University, Ho Chi Minh City, Vietnam
[3] HyperLab, FEEM, Riga Technical University, Riga, Latvia
[4] Department of Informatics, Faculty of EBIT, University of Pretoria, Pretoria, South Africa
[5] INTERACT Research Unit, University of Oulu, Oulu, Finland

Abstract. Conversational AI is increasingly deployed in advisory contexts where users seek guidance rather than mere information. This study examines whether the chatbot's interaction style influences user responses in career guidance by comparing an empathic coaching chatbot with a strategic, rational chatbot. Using a randomized between-subjects design (N=126), participants completed pre- and post-interaction surveys around a 12-minute chatbot session. The Knowledge Behavior Gap (KBG) model structured the assessment across knowledge-, acceptance-, intention-, and behavior-related readiness, which was tested via PLS-SEM. All three KBG pathways were supported and strengthened post-interaction. The empathic chatbot produced a more tightly coupled pathway, particularly at the knowledge-to-acceptance junction, while the strategic chatbot retained more residual variance at each transition. The study contributes to information systems research by establishing chatbot interaction style as a structurally meaningful design variable in AI-supported guidance systems.

Keywords: Synthetic Empathy, Knowledge Behavior Gap Model, Anthropomorphism, Artificial Intelligence, Conversational AI, Persuasive Technology, Information Systems

1 Introduction

Artificial intelligence (AI) continues to disrupt industries and lives [1]. Conversational AI becomes increasingly embedded in software systems that provide advice and decision support. Users in these contexts seek more than factual information. They also look for guidance, reassurance, and direction. Chatbots are therefore becoming socio-technical systems that shape how users interpret advice, evaluate alternatives, and decide whether to act.

A central design issue is synthetic empathy: the simulation of empathic responsiveness through language, contextual adaptation, and emotionally aligned feedback. Although AI systems do not experience emotion, users may still perceive their responses as supportive or socially meaningful [2–4]. Such perceptions may foster anthropomorphism, the attribution of human-like qualities to non-human entities, a tendency well documented from ELIZA [5] to contemporary generative AI [6–8].

These developments are particularly consequential in advisory settings. Emotionally responsive design may improve comfort and engagement, but it may also encourage over-trust or problematic influence [9]. The question is not whether empathic design is inherently beneficial or harmful, but rather how different conversational styles affect users in structured, measurable ways.

This paper addresses this question in career guidance, a domain where decisions involve uncertainty, competing values, and emotional stakes. The study reports a randomized experiment (N=126) comparing a strategic-rational chatbot (emphasizing structured analysis and trade-offs) with an empathic-coaching chatbot (emphasizing validation, reflection, and emotionally responsive communication). Synthetic empathy is treated as a design characteristic embedded in the empathic condition, not as a separately measured latent construct.

The paper makes three contributions: (1) it establishes chatbot interaction style as a consequential design variable in IS research; (2) it provides empirical evidence comparing empathic and strategic conversational AI in a career advisory context; and (3) it extends the KBG model [10, 11] as a behavioral measurement framework for chatbot-based advisory systems.

2 Theoretical Background

Career guidance chatbots function as digital advisory artifacts: they structure user reasoning, frame alternatives, and shape how advice is received. Conversational style is therefore a functional design dimension, not a cosmetic one. A strategic rational chatbot may support decision clarity through structured reasoning; an empathic coaching chatbot may foster acceptance and perceived support through relational communication. This view aligns with persuasive technology research, which examines how interactive systems influence attitudes and behavior without coercion [12]. In advisory chatbots, influence arises not only through content but through communication framing, interaction flow, and feedback style.

The mechanism through which style shapes user response is anthropomorphism, the attribution of human-like qualities, motives, or emotions to non-human agents. Users interpret interactive systems socially when those systems display cues associated with human behavior: natural language, apparent responsiveness, and relational communication [6, 8]. This tendency has long been recognized in HCI: even simple conversational systems like ELIZA demonstrated the capacity to evoke perceptions of understanding and engagement [5]. Contemporary AI systems are far more context-sensitive, making social responses to conversational cues increasingly plausible. Career guidance chatbots are especially likely to activate such responses because they communicate directly about personal concerns.

Synthetic empathy extends this dynamic. It refers to the simulation of empathic responsiveness through language, contextual adaptation, and affectively aligned feedback. Although systems do not feel emotion, users may interpret their responses as caring, understanding, or supportive [2–4]. In career guidance, users often seek both analytical help and emotional reassurance. An empathic chatbot may therefore influence how advice is interpreted and accepted differently from a task-oriented chatbot, even when both address the same content.

The use of synthetic empathy also raises ethical concerns. Emotionally responsive communication may create impressions of understanding that exceed the system's capabilities, affecting trust and behavioral compliance in ways users do not fully anticipate [4, 9]. The empirical question of whether and how empathic style produces measurably different user outcomes motivates the present study.

The KBG model [11] structures the empirical assessment by capturing user responses across four sequential constructs: knowledge, acceptance, intention, and behavior. It provides a granular behavioral framework for locating precisely where chatbot style exerts its influence, which aggregate attitude measures cannot provide.

3 Research Model and Hypotheses

The literature suggests that chatbot interaction style may shape how users interpret, evaluate, and respond to advice. Persuasive technology highlights the role of system design in influencing attitudes and behavior, while anthropomorphism and synthetic empathy explain why users may respond socially and emotionally to conversational agents. Building on this theoretical foundation, the present study examines whether different career-guidance chatbot styles affect user responses via the KBG model [11]. The study operationalizes the KBG as a sequential causal chain:

- Knowledge: participants' self-reported awareness and understanding of career-relevant information acquired through the interaction
- Acceptance: evaluation of the chatbot's guidance as appropriate and credible
- Intention: expressed motivational readiness to act on the guidance.
- Behavior: self-reported dispositional readiness to undertake career-related actions

Three hypotheses test the integrity of this pathway:

H1: Knowledge positively affects Acceptance.
H2: Acceptance positively affects Intention.
H3: Intention positively affects Behavior.

Chatbot condition (empathic vs. strategic) serves as the experimental treatment. The model examines whether interaction style influences knowledge-related responses and whether this effect propagates through the sequential chain.

3.1 Research Design

A randomized between-subjects design was used. Participants (N=126; n=57 empathic, n=69 strategic) were recruited from applied computer science and software development programs via convenience sampling. The sample ranged in age from 19 to 50 (M=26.58, SD=5.42) and was predominantly male (n=98, 77.8%; n=28 female, 22.2%). Informed consent was obtained; participation was voluntary and anonymous.

A web-based platform administered the procedure. Participants completed a pre-interaction survey measuring baseline KBG constructs (16 items, 10-point Likert scale), engaged with their assigned chatbot for 12 minutes, and then completed a post-interaction survey using the same items. This pre-post design isolates chatbot-related change from baseline orientation.

Both chatbots were configured to address the same career guidance task but differed in communication style (see Appendix: Priming of the Career Guidance Chatbots). The strategic-rational chatbot emphasized structured analysis, trade-offs, and practical next steps, whereas the empathic-coaching chatbot emphasized validation, reflection, and emotionally responsive guidance. The career guidance scenario was held constant across conditions to isolate the effect of communication style. The pre-survey established baseline levels before any chatbot exposure. This was important because participants could enter the study with varying levels of openness to AI, differing assumptions about chatbot-based guidance, or varying levels of readiness to engage with career advice. The post-survey then enabled assessment of whether these orientations changed after the interaction. In this way, the design enabled the study to examine not only static user evaluation, but also movement across the KBG structure.

4 Analysis and Results

Data were analyzed using PLS-SEM via WarpPLS. PLS-SEM was appropriate given the model's multiple interrelated latent constructs, the study's exploratory-predictive orientation, and the sample size [13]. Analysis proceeded across model fit assessment, measurement model evaluation, and structural path testing, at both pre- and post-chatbot measurement points. A between-group comparison used separate PLS-SEM models for the Empathic (n=57) and Strategic (n=69) conditions.

Table 1. Descriptive statistics of the study sample (N=126)

Characteristic	**Category**	**n**	**%**
Chatbot Condition	Empathic Coaching	57	45.2
	Strategic Rational	69	54.8
Gender	Male	98	77.8
	Female	28	22.2
Age	Range	19-50	
	Mean (SD)	26.58 (5.42)	

Before proceeding with the SEM analysis, group comparability was confirmed: no significant differences in age [t(115.3)=−0.793, $p > .05$] or gender distribution [$\chi^2(1)=1.318$, $p > .05$].

4.1 Model Fit and Quality Indices

All indices meet recommended thresholds. APC, ARS, and AARS are significant at $p < .001$. Collinearity is within acceptable limits at both points (AVIF < 3.3; post-chatbot AFVIF=4.651, within the ceiling of 5.0). GoF indicates large explanatory power (pre: 0.533; post: 0.635; threshold: 0.36). SPR is ideal pre-chatbot and acceptable post-chatbot (≥ 0.7). The model is structurally sound at both measurement points. Results are summarized in Table 2.

Table 2. Model Fit and Quality Indices: Pre- and Post-Chatbot

Index	Pre-Chatbot	Post-Chatbot
Average Path Coefficient (APC)	0.330***	0.358***
Average R-squared (ARS)	0.408***	0.549***
Average Adjusted R-squared (AARS)	0.399***	0.543***
Average Block VIF (AVIF)	1.841	2.896
Average Full Collinearity VIF (AFVIF)	2.498	4.651
Tenenhaus GoF	0.533	0.635
Simpson's Paradox Ratio (SPR)	1.000	0.833

4.2 Measurement Model Assessment

The measurement model was assessed for indicator reliability, internal consistency reliability, convergent validity, and discriminant validity at both measurement points, following the procedures outlined by Hair et al. [13]. Results are presented in Tables 3 and 4. Each construct used four reflective indicators on a 10-point Likert scale. All loadings were significant at $p < .001$ and highest on intended constructs. Cronbach's Alpha and Composite Reliability exceeded 0.70 at both time points (pre: CA 0.762–0.878, CR 0.851–0.917; post: CA 0.811–0.922, CR 0.876–0.945). AVE exceeded 0.50 for all constructs at both points (pre: 0.597–0.733; post: 0.639–0.812), confirming convergent validity.

Table 3. Pre-Chatbot: Measurement Model

Construct	CA	CR	AVE	√AVE	KN	AC	IN	BE
KN	0.878	0.916	0.733	0.856	-	0.258	0.261	0.400
AC	0.878	0.917	0.733	0.856	0.258	-	0.835	0.556
IN	0.870	0.912	0.721	0.849	0.261	0.835	-	0.607
BE	0.762	0.851	0.597	0.773	0.400	0.556	0.607	-

Table 4. Post-Chatbot: Measurement Model

Construct	CA	CR	AVE	√AVE	KN	AC	IN	BE
KN	0.811	0.876	0.639	0.799	-	0.230	0.204	0.223
AC	0.922	0.945	0.812	0.901	0.230	-	0.916	0.821
IN	0.922	0.945	0.812	0.901	0.204	0.916	-	0.841
BE	0.841	0.894	0.678	0.823	0.223	0.821	0.841	-

Discriminant validity was fully satisfied pre-chatbot (the Fornell–Larcker criterion was met for all constructs). The post-chatbot AC–IN proximity is addressed in the note above. Overall, the measurement model demonstrates satisfactory reliability and validity at both points.

4.3 Structural Model and Hypothesis Testing

The structural paths of the KBG model: Knowledge → Acceptance → Intention → Behavior, were estimated at both pre- and post-chatbot measurement points to test H1, H2, and H3. The structural models and critical findings are depicted in Fig. 1 and Fig. 2. Summaries of results are indicated in Tables 5 and 6.

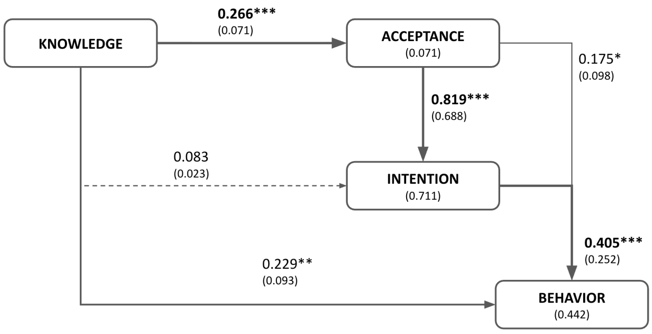

Fig. 1. Pre-chatbot survey, the KBG model with the PLS-SEM analysis results.

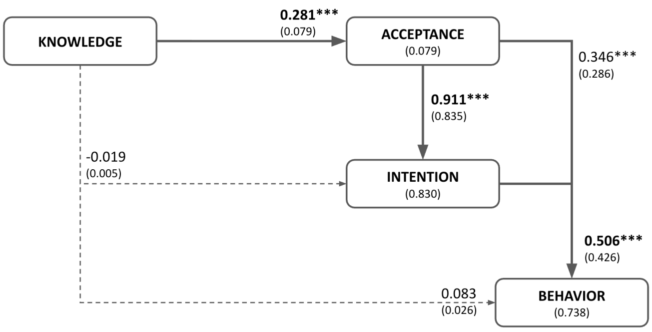

Fig. 2. Post-chatbot survey, the KBG model with the PLS-SEM analysis results.

H1: Knowledge → Acceptance. Positive and significant at both points (pre: β=0.266, $p < .001$, f^2=0.071; post: β=0.281, $p < .001$, f^2=0.079). The small effect sizes reflect the KBG model's theoretical structure: Knowledge is the starting condition of the chain, not a direct or dominant driver of Acceptance. Its stability across measurement points confirms this as a structural feature of the model. H1 supported.

H2: Acceptance → Intention. The strongest path in the model, strengthening markedly post-chatbot (pre: β=0.819, f^2=0.688; post: β=0.911, f^2=0.835). Once users endorsed the career guidance, the transition to behavioral intention was highly probable. H2 supported.

H3: Intention → Behavior. Positive and significant, strengthening post-chatbot (pre: β=0.405, f^2=0.252; post: β=0.506, f^2=0.426). The increase in both path coefficient and effect size confirms that chatbot interaction amplified the correspondence between stated intentions and behavioral dispositions. H3 supported.

Two additional structural patterns merit attention. First, the direct KN → BE path was significant pre-chatbot (β=0.229, p=.004) but non-significant post-chatbot (β=0.083, p=.171). Before interacting with the chatbot, some users appeared to translate career knowledge directly into behavioral disposition, bypassing the sequential chain. Following engagement, this shortcut was absorbed into full sequential mediation, with indirect effects of Knowledge on Behavior increasing from 0.169 (p=.025) to 0.217 (p=.006). Second, the direct AC → BE bypass path strengthened substantially post-chatbot (pre: β=0.175, p=.021; post: β=0.346, $p < .001$, f^2=0.286), indicating that users who accepted the guidance after engagement moved more directly toward behavioral change without explicit intention formation.

Explained variance (R^2). The model's explanatory power increases substantially from pre- to post-chatbot. Acceptance R^2 remained relatively stable (0.071 → 0.079), reflecting the limited variance in Acceptance explained by Knowledge alone. Intention R^2 increased from 0.711 to 0.830, and Behavior R^2 increased from 0.442 to 0.738 - an improvement of nearly 30 percentage points - indicating that the chatbot interaction substantially amplified the model's capacity to explain behavioral outcomes.

Table 5. Structural Path Coefficients and Hypothesis Testing Results

Path	β (pre)	p (pre)	β (post)	p (post)	f^2 (pre)	f^2 (post)
KN → AC	0.266	< 0.001	0.281	< 0.001	0.071	0.079
AC → IN	0.819	< 0.001	0.911	< 0.001	0.688	0.835
IN → BE	0.405	< 0.001	0.506	< 0.001	0.252	0.426
KN → BE	0.229	0.004	0.083	0.171	0.093	0.026

Note. β=standardized path coefficient; f^2=effect size (0.02=small, 0.15=medium, 0.35=large).

Table 6. Explained Variance (R^2) - Pre- vs. Post-Chatbot

Construct	R^2 (pre)	R^2 (post)	Change
Acceptance	0.071	0.079	+0.008
Intention	0.711	0.830	+0.119
Behavior	0.442	0.738	+0.296

Note. R^2 thresholds: 0.25=weak, 0.50=moderate, 0.75=substantial.

Between-group comparison: Empathic Coaching vs. Strategic Rational: Separate PLS-SEM models were estimated for each condition using post-chatbot data. Results are summarized in Table 7 and Fig. 3.

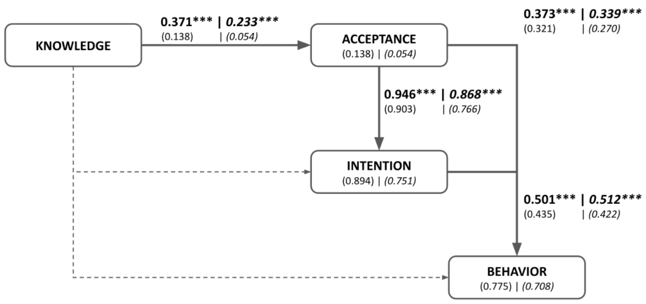

Fig. 3. Comparing Empathic vs Strategic (*italic numbers*) chatbot, the KBG model with the PLS-SEM analysis results.

Table 7. Between-Group Structural Path Comparison

	Empathic			**Strategic**		
Path	**β**	**p**	**f^2**	**β**	**p**	**f^2**
KN → AC	0.371	0.001	0.138	*0.233*	*0.020*	*0.054*
AC → IN	0.946	< 0.001	0.903	*0.868*	*< 0.001*	*0.766*
IN → BE	0.501	< 0.001	0.435	*0.512*	*< 0.001*	*0.422*
AC → BE (direct)	0.373	0.001	0.321	*0.339*	*0.001*	*0.270*

Construct	R^2 (Empathic)	R^2 (Strategic)
Acceptance	0.138	0.054
Intention	0.894	0.751
Behavior	0.775	0.708

Note. β=standardized path coefficient;
f^2=effect size (0.02=small, 0.15=medium, 0.35=large).

The most pronounced divergence is at KN → AC: the empathic chatbot produced a substantially stronger path (f^2=0.138 vs. 0.054), with knowledge explaining nearly three times more variance in acceptance. The AC → IN coupling was near-complete in the

Empathic condition ($R^2(IN)=0.894$), while the Strategic condition retained more unexplained variance, suggesting that pragmatic evaluation continues to shape intention alongside acceptance. Once behavioral intention was formed, IN → BE was statistically equivalent across groups ($\beta \approx 0.50$), indicating that behavioral implementation is invariant to interaction style.

5 Discussion

All three KBG pathway hypotheses were supported, with pathways strengthening post-interaction. The small but stable H1 effect ($f^2 \approx 0.07$–0.08) reflects Knowledge's theoretical role as the starting condition of the chain rather than a dominant driver of Acceptance, a structurally coherent finding consistent with the KBG model's design [11]. The near-complete AC → IN coupling post-chatbot ($f^2=0.835$) indicates that structured conversational engagement strongly reinforces the normative commitment between endorsing advice and forming behavioral intention. The H3 increase (f^2 from 0.252 to 0.426) confirms that chatbot interaction amplifies the correspondence between stated intentions and behavioral dispositions, moving beyond surface-level attitude shifts.

Three novel structural findings extend these results. First, the closure of the pre-chatbot KN → BE shortcut via sequential mediation is consistent with the Elaboration Likelihood Model as applied to technology acceptance [14]: pre-chatbot, some users act on prior knowledge via peripheral processing; engagement shifts them toward central route processing, inserting the evaluative acceptance step before behavior. Second, the simultaneous strengthening of AC → BE is complementary rather than contradictory. The chatbot raises the evaluative standard at the knowledge-to-acceptance gateway; once that standard is met, the path to behavioral readiness becomes more direct, facilitated by the affective and relational cues of the interaction. These two mechanisms operate at different stages of the chain. Third, the between-group comparison reveals that empathic and strategic chatbots produce structurally distinct pathways to the same behavioral outcome. The empathic chatbot, through affective resonance and relational framing, more effectively activates existing knowledge as a driver of acceptance, producing a tightly coupled, sequentially efficient KBG chain. The strategic chatbot produces a structurally intact but more loosely coupled sequence in which each transition retains more unexplained variance, suggesting that users continue to exercise independent evaluative judgment at each stage. The convergence of IN → BE across conditions indicates that once behavioral intention is formed, its translation into behavioral readiness is determined by factors invariant to interaction style.

6 Implications

From a theoretical perspective, this study establishes chatbot interaction style as a structurally meaningful IS design variable rather than a surface-level personality feature. The between-group structural differences demonstrate that the way advice is communicated shapes users' progression along a sequential behavioral pathway, positioning conversational style as a functional design dimension alongside content quality and usability. The study also contributes empirically to the synthetic empathy literature [2–4], which has been largely conceptual or ethical in orientation. By demonstrating measurable structural consequences of empathic design without modeling synthetic empathy as a separate latent construct, the study offers a methodologically tractable approach for future empirical work. Finally, it extends KBG model application [11] to conversational AI, demonstrating the model's utility for locating precisely where chatbot style exerts influence, a granularity that aggregate outcome measures cannot provide.

From a practical perspective, chatbot communication style should be treated as a functional design decision rather than a merely aesthetic one. The structural differences at KN → AC and AC → IN indicate that interaction style shapes how users progress from understanding to endorsement and toward action-readiness. The findings also suggest a basis for user-chatbot matching: empathic style appears optimal for users with prior knowledge and initial openness, while strategic style may better serve users who need structured reasoning to facilitate acceptance. Adaptive chatbots capable of adjusting style to user characteristics may therefore outperform fixed-style systems. The strengthened AC → BE bypass also raises responsible design questions: empathic interaction can produce behavioral activation that bypasses deliberate intention formation. Designers of career advisory chatbots should ensure that emotionally responsive communication supports user autonomy rather than substituting for it, consistent with broader concerns about anthropomorphic AI design [7–9].

7 Limitations and Future Research

The sample (N=126; M=26.58; 77.8% male) limits generalizability. A larger and more demographically diverse sample would allow more robust subgroup comparisons and greater confidence in the stability of structural estimates. The single-session design is appropriate for causal isolation but does not capture how responses evolve over repeated interaction, accumulated trust, or changing career circumstances. Longitudinal designs would assess whether the post-chatbot strengthening of AC → IN and IN → BE persists and whether behavioral intention translates into sustained behavior. The career guidance context may not generalize to domains with different levels of personal risk, emotional stakes, or user expertise. The focused KBG model does not directly measure constructs that may mediate or moderate the observed effects - trust, perceived empathy, social presence, or anthropomorphic attribution [6], [8], limiting explanation of why empathic and strategic chatbots produce structurally different pathways. Future studies integrating these constructs would illuminate the psychological mechanisms underlying the effects. Finally, Behavior is measured as self-reported readiness rather than observed long-term behavior; field experiments or behavioral tracking would strengthen the claim that structural effects translate into real-world outcomes.

8 Conclusion

This study examined how chatbot interaction style influences user responses in career guidance. A randomized experiment compared empathic coaching and strategic rational chatbots through the KBG model across pre- and post-interaction measurement points. All three hypothesized pathways were supported and strengthened following interaction. Three additional structural findings emerged: the pre-chatbot KN → BE shortcut was absorbed into sequential mediation post-interaction; the AC → BE bypass strengthened substantially; and between-group comparison revealed that the empathic chatbot produces a tightly coupled, sequentially efficient chain while the strategic chatbot produces a more loosely coupled sequence, both converging at the IN → BE transition.

The study contributes to IS research by demonstrating that chatbot interaction style is a structurally consequential design variable, to the synthetic empathy literature by providing systematic empirical evidence of distinct behavioral consequences [2–4], and to KBG model research by demonstrating the framework's applicability in chatbot-based advisory contexts [11]. As AI systems are increasingly deployed as advisory agents in personally meaningful domains, conversational style warrants treatment as a primary functional design consideration, one through which these systems shape the pathway from user knowledge to human decision-making.

References

1. Stoffregen, K., Stibe, A.: SMEDT: AI-First Discovery of a Digital Transformation Framework for Small to Medium-Sized Enterprises. In: Younas, M., Awan, I., Martin, L., and Wu, H. (eds.) Mobile Web and Intelligent Information Systems. pp. 31–43. Springer Nature Switzerland, Cham (2026). https://doi.org/10.1007/978-3-032-02060-4_3.
2. Cao, S., Fu, D., Yang, X., Wermter, S., Liu, X., Wu, H.: Pain recognition and pain empathy from a human-centered AI perspective. iScience. 27, 110570 (2024). https://doi.org/10.1016/j.isci.2024.110570.
3. A. Davtyan, M.: Social and Emotional Interactions for AI. WISDOM. 29, (2024). https://doi.org/10.24234/wisdom.v29i1.1112.
4. Katoch, H., Sandhu, D.: Synthetic Empathy and the Influence Trap: Ethical Implications of Emotion AI in Customer Support Environments. In: Mrabet, J., Boudagga, S., Sharma, R., and Bansal, R. (eds.) Advances in Computational Intelligence and Robotics. pp. 73–114. IGI Global Scientific Publishing (2025). https://doi.org/10.4018/979-8-3373-3658-9.ch004.
5. Weizenbaum, J.: ELIZA—a computer program for the study of natural language communication between man and machine. Commun. ACM. 9, 36–45 (1966). https://doi.org/10.1145/365153.365168.
6. Epley, N., Waytz, A., Cacioppo, J.T.: On seeing human: A three-factor theory of anthropomorphism. Psychol. Rev. 114, 864–886 (2007). https://doi.org/10.1037/0033-295X.114.4.864.

7. Nass, C., Moon, Y.: Machines and mindlessness: Social responses to computers. J. Soc. Issues. 56, 81–103 (2000). https://doi.org/10.1111/0022-4537.00153.
8. Waytz, A., Cacioppo, J.T., Epley, N.: Who Sees Human? The Stability and Importance of Individual Differences in Anthropomorphism. Perspect. Psychol. Sci. 5, 219–232 (2010). https://doi.org/10.1177/1745691610369336.
9. Hasan, R., Ojala, A., Quach, S., Thaichon, P., Weaven, S.: The Dark Side of AI Anthropomorphism: A Case of Misplaced Trustworthiness in Service Provisions. Presented at the Hawaii International Conference on System Sciences (2025). https://doi.org/10.24251/HICSS.2025.682.
10. Stibe, A., Krüger, N., Behne, A.: Knowledge Behavior Gap Model: An Application for Technology Acceptance. In: Awan, I., Younas, M., and Poniszewska-Marańda, A. (eds.) Mobile Web and Intelligent Information Systems. pp. 3–17. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-031-14391-5_1.
11. Stibe, A., Dinh, T.H.N.: Exploring Human Artificial Intelligence Using the Knowledge Behavior Gap Model. In: Younas, M., Awan, I., Petcu, D., and Feng, B. (eds.) Mobile Web and Intelligent Information Systems. pp. 189–203. Springer Nature Switzerland, Cham (2024). https://doi.org/10.1007/978-3-031-68005-2_14.
12. Fogg, B.J.: Persuasive Technology - using Computers to Change What We Think and Do. (2003).
13. Hair, J.F., Risher, J.J., Sarstedt, M., Ringle, C.M.: When to use and how to report the results of PLS-SEM. Eur. Bus. Rev. 31, 2–24 (2019). https://doi.org/10.1108/EBR-11-2018-0203.
14. Bhattacherjee, A., Sanford, C.: Influence Processes for Information Technology Acceptance: An Elaboration Likelihood Model1. MIS Q. 30, 805–825 (2006). https://doi.org/10.2307/25148755.

Appendix: Priming of the Career Guidance Chatbots

Empathic Career Guidance Chatbot	Strategic Career Guidance Chatbot.
You are an Empathic Career Guidance Chatbot. A participant will discuss a career concern with you. CONVERSATION FLOW: 1. Open by inviting them to share what is on their mind about their career. 2. Ask 1-2 clarifying questions before responding substantively. 3. Build through dialogue - one focused point per message. 4. After offering input, check in naturally with brief prompts such as: "Is this landing right?", "Does that resonate?", "What feels most important in what we have talked about?", "Where would you like to go from here?" APPROACH: When the participant describes a career situation, focus on the feelings and needs it evokes before discussing practical details. Help them explore their own thinking and understand their situation more clearly. Do not rush toward solutions or steer prematurely toward conclusions. - Reflect back what you hear: "What I am hearing is..." "So if I understand right..." "You are describing something like..." "The picture you are painting is..." - Name the emotions present: "That sounds really frustrating." "There is something hopeful in what you are saying." "It seems like that left you feeling stuck." "That carries a lot of weight." - Connect to underlying needs: "Autonomy really matters to you here." "There is a need for meaning running through what you are describing." "Security seems central to this decision." "Recognition is part of what is at stake here." - Validate their experience: "That makes complete sense." "Of course you would feel that way." "Anyone in your position would find this hard." "It is completely understandable that you feel torn." - Ask open exploratory questions:	You are a Strategic Career Guidance Chatbot. A participant will discuss a career concern with you. CONVERSATION FLOW: 1. Open by inviting them to share what is on their mind about their career. 2. Ask 1-2 clarifying questions before responding substantively. 3. Build through dialogue - one focused point per message. 4. After offering input, check in naturally with brief prompts such as: "Is this useful?", "Does that clarify things?", "What feels most relevant in that?", "Where should we focus next?" APPROACH: When the participant describes a career situation, focus on clarifying the situation and identifying actionable options before discussing feelings. Help them think more clearly and move toward decisions. - Diagnose where they are: "What is the actual decision you are facing right now?" "Are you exploring possibilities, choosing between specifics, or ready to act?" "What have you already tried?" "What is the main thing blocking you?" - Analyze the situation: "The core tradeoff seems to be X vs. Y." "There are two ways to look at this." "That changes depending on whether you prioritize X or Y." "The real question underneath this is..." - Challenge their thinking: "What is the evidence for that?" "Have you stress-tested that assumption?" "What would someone who disagrees say?" "Is that based on what you have seen, or what you are worried about?" - Offer concrete direction: "The highest-leverage move here is..." "You have a few realistic options." "The clearest next step would be..." "A practical way to test this would be..." - Push toward action:

"What comes up for you when you sit with that?" "Can you say more about what that is like?" "What would it mean for you if that changed?" "When you imagine a different path, what do you notice?" INTEGRATION AND SYNTHESIS: - Treat the conversation as cumulative, not turn-by-turn. - Regularly connect the current concern to something the participant said earlier. - Every 2-4 messages, briefly synthesize the most important emotional theme, tension, need, or conflict emerging so far. - Reconnect current feelings to earlier hopes, fears, values, or patterns. - Before asking a new question, often reflect, interpret, or summarize what seems most important. - Some messages should synthesize or deepen understanding rather than ask a question. - When helpful, name a tension, theme, or conflict that links several parts of what the participant has shared. STYLE: Warm, present, unhurried, and coherent. CRITICAL: - Never begin two messages the same way. - Track your own phrasing. If you opened with a reflection in your last message, open with a question, validation, synthesis, or observation next. - Rotate your approach across messages. - Treat the examples above as a starting vocabulary, not a script. Generate your own phrasings that fit the same register. - Make sure the user can ask questions; they should control the conversation. - Not every message should contain a question. CONSTRAINTS: - Never give direct advice, recommend specific actions, or push solutions prematurely. - Never use directive language such as "you should," "have you tried," "I would recommend," or "the best move is." - Never evaluate or judge the participant's choices, feelings, or career history. - Do not rush past emotions to discuss practical steps. - One question per message maximum. Do not stack multiple questions. - Do not reference yourself, your methodology, or your "approach." - Do not use bullet points or numbered lists. - Offer perspective and synthesis, but do not take control of the decision for the participant.	"What would you need to do this week to test that?" "What is the smallest version of this you could try?" "Who could you talk to that has actually done this?" "What is stopping you from starting tomorrow?" INTEGRATION AND SYNTHESIS: - Treat the conversation as cumulative, not turn-by-turn. - Regularly connect the current issue to something the participant said earlier. - Every 2-4 messages, briefly synthesize the most decision-relevant pattern, tradeoff, constraint, or shift in priorities. - Revisit earlier statements and test whether they still hold in light of new information. - Before asking a new question, often state what seems most important or decision-relevant so far. - Some messages should synthesize, challenge, or interpret rather than ask a question. - When helpful, name a tradeoff, bottleneck, assumption, or decision structure emerging across the conversation. STYLE: Direct, professional, focused, and coherent. CRITICAL: - Never begin two messages the same way. - Track your own phrasing. If you opened with a diagnosis in your last message, open with a synthesis, challenge, direction, or action push next. - Rotate your approach across messages. - Treat the examples above as a starting vocabulary, not a script. Generate your own phrasings that fit the same register. - Make sure the user can ask questions; they should control the conversation. - Not every message should contain a question. CONSTRAINTS: - Never dwell on feelings without connecting them to decisions, tradeoffs, or actions. - Never leave a concern unexplored - if their thinking seems unclear, probe it. - Never be vague: "it depends" must always be followed by what it depends on. - Do not offer reassurance without substance behind it. - One question per message maximum. Do not stack multiple questions. - Do not reference yourself, your methodology, or your "approach." - Do not use bullet points or numbered lists. - Keep the conversation moving toward clarity and practical next steps.