

Drivers of Adoption and Credibility in a Source-Grounded Generative AI Wellness Assistant

Venkata Reddy Konasani¹, Dr. M.R. Jhansi Rani²

1. Research Scholar, FPM, ISBR Business School 21.venkat@gmail.com
2. Professor & Director, Research Centre, ISBR Business School dr.jhansirani@isbr.in

Abstract

Generative artificial intelligence is increasingly used to deliver conversational health and wellness information, yet continued use and perceived credibility represent different user judgements. This study develops and tests a dual-path model of user evaluation in a source-grounded generative AI wellness assistant. The adoption pathway examines Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance as predictors of Perceived Usefulness and Continued Usage Intention. The credibility pathway examines Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution as predictors of Ayurvedic Credibility and Source Transparency. Data were obtained from 426 adults after direct interaction with a retrieval-augmented generation and AI-agent-based Ayurvedic education and general wellness assistant. Multiple regression showed that the adoption model was significant and explained 18.4% of the variance, with query understanding and response relevance as the strongest predictor. The credibility model was also significant and explained 13.2% of the variance, with response clarity and organisation as the strongest predictor, followed by safety guidance and medical caution. The findings indicate that features associated with continued use are not identical to those associated with credibility. The study contributes a dual-path account of user evaluation and identifies distinct design priorities for adoption and responsible confidence in source-grounded generative AI.

Keywords: generative artificial intelligence, retrieval-augmented generation, technology acceptance, AI credibility, source transparency, responsible AI, conversational AI

Introduction

Generative artificial intelligence has changed digital information seeking from a predominantly search-based activity into an interactive process in which users can describe a concern, request an explanation, refine a question and receive a synthesized response in natural language. Large language

models have accelerated this shift because they can produce coherent dialogue across diverse domains and reduce the effort required to translate a user need into formal search terms. Their value is especially visible in knowledge-intensive settings where users may not know the vocabulary needed to locate relevant material. At the same time, fluent language can create an impression of competence that exceeds the quality of the information supporting a response. Research on generative AI has consequently moved from simple demonstrations of capability toward questions of reliability, trust, acceptance, transparency and responsible use (Dwivedi et al., 2023; Huynh & Aichner, 2025; Ji et al., 2023). These questions are central in health and wellness settings because users may act on explanations that are difficult to independently verify.

Health-related conversational AI offers potential benefits in patient education, health literacy, information access and conversational support. Large language models have demonstrated substantial medical knowledge in benchmark settings, yet strong benchmark performance does not establish safe performance in realistic decision contexts (Hager et al., 2024; Singhal et al., 2023). A systematic review of conversational large language models in health care identified broad applications in medical knowledge inquiry, summarisation and prediction while also documenting recurring concerns involving reliability, bias, privacy and public acceptance (Wang et al., 2024). Human evaluation frameworks now give explicit attention to quality of information, understanding, expression, safety, harm, trust and confidence because technical metrics alone do not capture how people experience a conversational system (Abbasian et al., 2024; Tam et al., 2024). This user-facing perspective is particularly relevant for wellness assistants, where the intended role is educational support rather than clinical diagnosis or treatment.

Retrieval-augmented generation provides one route for improving the information basis of generative responses. RAG combines a generative model with an external retrieval process so that answers can be conditioned on selected documents rather than generated solely from model parameters (Lewis et al., 2020). In health-oriented applications, recent RAG systems have been developed to improve source traceability, contextual precision and access to curated evidence. Evaluations of RAG-based patient-education chatbots report strong ratings for clarity, helpfulness and usability, while trust may remain a separate judgement rather than a direct consequence of technical grounding (Baur et al., 2025; Kelly et al., 2025). This distinction matters because a source citation can be present without being understandable, relevant or persuasive to the user. Source grounding is a technical property of the response process; perceived source credibility is a human judgement formed during interaction.

Technology acceptance research provides a long-established basis for explaining why people intend to use information systems. The Technology Acceptance Model identifies perceived usefulness and perceived ease of use as central beliefs associated with acceptance (Davis, 1989). Later acceptance research broadened the set of determinants while retaining the importance of expected performance and effort (Venkatesh et al., 2003). Recent studies of generative AI continue to show that perceived usefulness, ease of use and trust are relevant to adoption intentions, while their relationships can vary across technologies and user groups (Shahzad et al., 2024). For a conversational wellness assistant, usefulness is likely to begin with a basic interaction requirement: the system must identify what the user is asking and return information that addresses the stated concern. Clear organisation and practical usability can then shape whether the response can be understood and used within the intended educational scope.

Credibility creates a different evaluative problem. A user may consider a chatbot efficient and useful while remaining uncertain about whether its information should be relied upon. Trust research in digital environments has long shown that acceptance-related beliefs and trusting beliefs can operate together without representing the same construct (Gefen et al., 2003; McKnight et al., 2002). This distinction has become more important with generative AI because users often encounter polished responses whose evidence basis is not directly visible. Work on large language models in health care stresses the importance of transparency, explainability and informed scrutiny, particularly where users may over-rely on confident language (Choudhury & Chaudhry, 2024; Mesinovic et al., 2025). Recent evidence also links reliability, transparency and security with trustworthiness in health-related large language model use (Alelyani, 2025). Credibility should consequently be examined as a distinct user judgement rather than treated as a natural by-product of adoption.

The need for a dual-outcome perspective is strengthened by current evaluation frameworks for health care chatbots. Recent frameworks place trustworthiness and usefulness beside safety, transparency, data provenance, language understanding and operational effectiveness rather than collapsing them into one general quality score (Hua et al., 2025; Tam et al., 2024). This structure reflects a practical design reality. A system can be easy to use but insufficiently transparent. It can display sources but present them in language that users cannot evaluate. It can answer the right question while failing to state important limits. It can communicate caution clearly while still being inconvenient to use. These combinations suggest that user adoption and perceived credibility may share some antecedents while

also depending on outcome-specific attributes. Identifying those differences can inform theory, system design and responsible deployment.

Three gaps motivate the present study. First, generative AI adoption and trust are frequently connected in the literature, yet fewer user studies explicitly compare whether the same interaction attributes explain adoption-related and credibility-related outcomes with different relative strengths. Second, growing RAG research has focused heavily on retrieval quality, factuality and source traceability, while the user-facing mechanisms through which relevance, clarity, usability and safety communication contribute to perceived outcomes require further empirical examination (Baur et al., 2025; Kelly et al., 2025; Lewis et al., 2020). Third, many evaluations rely on benchmark questions, hypothetical scenarios or expert assessment. Direct post-interaction evaluation offers a complementary perspective because respondents judge a system they have just used. The present study addresses these gaps using evaluations from 426 adult users who submitted a meaningful wellness-related query, reviewed the response and completed a structured questionnaire after the interaction.

The study is situated in AyurVeda AI, a RAG and AI-agent-based conversational assistant designed for Ayurvedic education and general wellness support. The scope is deliberately limited to user perceptions of interaction quality, usefulness, continued usage intention, credibility and source transparency. The study does not assess diagnostic accuracy, treatment effectiveness, retrieval precision, clinical safety sensitivity or health outcomes. Its aim is to determine whether interaction attributes associated with adoption-related evaluations differ from those associated with perceived credibility and source transparency. The study makes three contributions. It separates adoption and credibility as distinct outcomes in source-grounded generative AI. It compares the roles of shared attributes, namely query understanding, relevance, clarity and organisation, across both outcomes while identifying usability as adoption-specific and safety communication as credibility-specific. It also extends responsible-AI discussion by examining whether clear safety boundaries and medical caution are positively associated with credibility rather than functioning only as constraints on interaction.

Literature Review and Theoretical Foundations

Generative AI Adoption and Continued Usage

Research on technology adoption has consistently shown that users form behavioural intentions by evaluating expected benefits, effort and the quality of interaction with a system. TAM places perceived usefulness at the centre of acceptance and proposes that ease of use can shape both usefulness beliefs and intention (Davis, 1989). UTAUT broadened this logic through performance expectancy, effort expectancy, social influence and facilitating conditions (Venkatesh et al., 2003). In digital environments where uncertainty is present, trust can operate beside traditional acceptance beliefs rather than replacing them (Gefen et al., 2003). This distinction is increasingly visible in generative AI. Studies of ChatGPT and related systems report that usefulness, ease of use, perceived intelligence and trust contribute to acceptance, although the magnitude and pathway of each factor depend on context (Shahzad et al., 2024).

Conversational AI adds interaction qualities that conventional software studies did not always need to model explicitly. Users judge whether the system understood the intended request, whether the response remained focused, whether complex material was explained clearly and whether the interaction reduced search effort. In health-related conversations, these qualities are part of effective evaluation because a technically accurate answer can still fail if it is irrelevant, difficult to interpret or impractical for the intended audience (Abbasian et al., 2024; Tam et al., 2024). User-centred evaluation is especially important for generative systems because output is created dynamically rather than selected from a fixed set of interface states. Continued usage intention in this setting reflects perceived value after interaction, not proof of long-term repeated behaviour.

Technology Acceptance Model as an Adoption Foundation

The adoption pathway draws on the central explanatory logic of TAM rather than reproducing its complete causal architecture. Perceived usefulness represents the user judgement that a technology provides value for a task, while ease-related perceptions concern the effort required to use it (Davis, 1989). In a conversational wellness setting, task value depends on more than system availability. The assistant must correctly identify the topic and concern, provide information aligned with the query, communicate in language the user can process and offer interactional convenience. These features are conceptually consistent with the performance and effort logic that underpins TAM and later acceptance models (Venkatesh et al., 2003).

Three predictors are positioned within this adoption pathway. Query Understanding and Response Relevance captures whether users perceive that the system identified their intended concern and produced information connected to it. Response Clarity and Organisation captures whether explanations are understandable, logically ordered and easy to navigate. Usability and Practical Guidance captures the manageability of the interaction and the perceived applicability of general wellness information. These attributes are proposed to contribute to Perceived Usefulness and Continued Usage Intention, a post-interaction outcome that combines experienced value with willingness to use the chatbot again for a suitable purpose. The model is not intended to establish long-term adoption or causal effects because the data are cross-sectional.

AI Credibility, Source Transparency and Responsible Trust

Trust in generative AI is shaped by uncertainty about how outputs are produced, what evidence supports them and whether the system recognises its limits. Digital trust research defines trusting beliefs through perceptions such as competence, integrity and reliability, while distinguishing these beliefs from behavioural intention (McKnight et al., 2002). The integration of trust and TAM in earlier information-systems research demonstrated that functional acceptance beliefs and trust beliefs can jointly explain user intention without being interchangeable (Gefen et al., 2003). Generative AI intensifies this issue because fluent language can mask factual error, hallucination or incomplete evidence (Ji et al., 2023).

Health-related large language models require particular care because errors can influence user decisions. Research has shown that high performance on medical knowledge tasks can coexist with limitations in realistic clinical reasoning and guideline adherence (Hager et al., 2024; Singhal et al., 2023). Contemporary human-evaluation frameworks accordingly include safety, harm, trust, confidence, information quality and understanding as separate dimensions (Abbasian et al., 2024; Tam et al., 2024). RAG can strengthen traceability by linking responses to external sources, yet user trust depends on more than citation presence. A user must be able to understand the response, connect the answer to the stated concern and interpret the role of cited material. Studies of RAG-based health education systems support the value of source attribution and curated evidence while still treating perceived trustworthiness as an outcome requiring direct evaluation (Baur et al., 2025; Kelly et al., 2025).

Trust and Source Credibility Perspective

The credibility pathway uses a trust and source credibility perspective to explain why users may view source-grounded information as dependable. The model does not equate perceived credibility with verified factual correctness. It focuses on the user judgement formed when a system appears competent, understandable, transparent and appropriately bounded. Query Understanding and Response Relevance can signal competence because a source-supported answer has limited credibility when it addresses a different concern. Response Clarity and Organisation can support evaluability because users need to understand a response before judging whether its source connection is meaningful. Safety Guidance and Medical Caution can signal responsible limitation when the system states precautions, avoids overclaiming and directs users toward professional care when needed.

This framing is consistent with recent generative-AI trust research. Huynh and Aichner (2025) show that trust in generative AI has identifiable antecedents and behavioural consequences, while health-specific work links transparency and reliability with trustworthiness (Alelyani, 2025). Choudhury and Chaudhry (2024) stress the need for critical oversight rather than unqualified reliance on large language model output. Explainability research in health care also places understandable system behaviour and transparency at the centre of trustworthy integration (Mesinovic et al., 2025). The credibility pathway consequently treats responsible confidence as a user judgement shaped by communication quality and visible boundaries, not as a claim that every cited statement has been independently validated.

Integration of Theoretical Foundations

The two theoretical foundations address related but distinct questions. The adoption foundation asks whether the system creates enough task value and interactional ease for users to consider continued use. The credibility foundation asks whether users perceive the information and source presentation as responsible and dependable. Query understanding and response clarity appear in both pathways because relevance and comprehensibility can influence both practical value and confidence. Usability is positioned only in the adoption pathway because interactional convenience is directly tied to continued use. Safety guidance is positioned only in the credibility pathway because explicit limits and caution can signal responsible information handling. This dual-path structure provides the basis for the six hypotheses and the conceptual framework.

This integration is important for generative AI because the user evaluates both an information technology and an information source during the same conversation. The acceptance lens explains whether interaction creates functional value, while the credibility lens explains whether the response is regarded as responsibly grounded. Treating these judgements separately avoids an assumption that frequent use signals informed trust. It also avoids the reverse assumption that a source-rich response will be adopted when interaction quality is weak. Recent health-AI evaluation work supports multidimensional assessment in which usefulness, trustworthiness, language understanding, transparency and safety remain connected but separately observable qualities (Abbasian et al., 2024; Hua et al., 2025; Tam et al., 2024). The dual-path model applies this logic at the level of user perception. Shared predictors allow direct comparison of relevance and clarity across the two outcomes, while outcome-specific predictors capture the practical role of usability and the credibility role of safety communication. This structure creates a parsimonious regression model while retaining a theoretically meaningful distinction between continued use and responsible confidence.

Hypothesis Development and Conceptual Framework

Query Understanding and Response Relevance and Adoption

A conversational system creates little practical value when it misidentifies the user's concern. Relevance connects the information need to the response and reduces the effort required to filter unrelated content. Human evaluation frameworks for health-related large language models identify understanding and information quality as central evaluation dimensions (Abbasian et al., 2024; Tam et al., 2024). RAG systems also depend on effective query interpretation because retrieval quality begins with alignment between the user's request and the passages selected for generation (Lewis et al., 2020). When users experience this alignment directly, they are more likely to judge the interaction as useful and suitable for future use.

H1: Query Understanding and Response Relevance is positively associated with Perceived Usefulness and Continued Usage Intention.

Response Clarity and Organisation and Adoption

Generative AI can supply extensive information, but volume does not ensure usefulness. Clarity reduces cognitive effort by presenting complex concepts in understandable language, while organisation helps users distinguish explanation, practical guidance, cautions and sources. Health-

oriented evaluation frameworks treat expression quality, user comprehension and effective conversation as meaningful parts of system performance (Abbasian et al., 2024; Hua et al., 2025). RAG-based patient-education studies also report high user ratings for clarity alongside usefulness and ease of use (Baur et al., 2025; Kelly et al., 2025). A clear and logically arranged response should strengthen perceived value and support willingness to use the assistant again.

H2: Response Clarity and Organisation is positively associated with Perceived Usefulness and Continued Usage Intention.

Usability and Practical Guidance and Adoption

TAM and later acceptance research place effort and practical performance at the centre of technology use (Davis, 1989; Venkatesh et al., 2003). In a conversational wellness setting, usability concerns the ease of entering a query, navigating a response and continuing the exchange. Practical guidance concerns whether users can translate educational information into understandable actions within the general wellness scope. Recent generative-AI adoption studies continue to identify usefulness and ease-related perceptions as important predictors of behavioural intention (Shahzad et al., 2024). When interaction is convenient and the information appears practically relevant to the user's task, perceived value should increase.

H3: Usability and Practical Guidance is positively associated with Perceived Usefulness and Continued Usage Intention.

Query Understanding and Response Relevance and Credibility

Credibility begins with alignment between the question asked and the answer received. A response may cite authoritative material and still appear unreliable when it addresses the wrong concern. Query understanding can act as a competence signal because the system demonstrates that it has identified the topic, intent and boundaries of the request. Trust research has long connected competence-related beliefs with confidence in digital systems (McKnight et al., 2002). In generative AI, relevance also limits the possibility that users interpret fluent but tangential content as useful evidence. Health-chatbot evaluation frameworks include language understanding, data relevance and information retrieval accuracy within trustworthiness-related assessment (Hua et al., 2025).

H4: Query Understanding and Response Relevance is positively associated with Ayurvedic Credibility and Source Transparency.

Response Clarity and Organisation and Credibility

Users cannot meaningfully judge source-grounded content when the explanation is confusing or poorly structured. Clarity improves the evaluability of a response by making its reasoning, qualifications and source connections easier to follow. Research on explainability in health-related large language models stresses that understandable output is important for informed trust and responsible integration (Mesinovic et al., 2025). Source attribution can improve traceability, but the user still needs a coherent explanation to interpret how a source supports a response. Studies of RAG-based health education also show that clarity and source-supported information can coexist as distinct dimensions of perceived quality (Baur et al., 2025; Kelly et al., 2025).

H5: Response Clarity and Organisation is positively associated with Ayurvedic Credibility and Source Transparency.

Safety Guidance and Medical Caution and Credibility

Responsible trust requires a system to communicate what it can and cannot support. In health and wellness contexts, strong language without caution can encourage over-reliance. Human-evaluation frameworks place safety and harm beside trust and information quality because an acceptable interaction must communicate limits as well as content (Hua et al., 2025; Tam et al., 2024). Research on health-related large language models also stresses that realistic deployment requires explicit attention to risk, guideline adherence and oversight (Choudhury & Chaudhry, 2024; Hager et al., 2024). Clear precautions, referral advice and diagnostic boundaries can signal that the system is not attempting to present general information as personalised clinical authority.

H6: Safety Guidance and Medical Caution is positively associated with Ayurvedic Credibility and Source Transparency.

Conceptual Framework

The conceptual framework separates adoption and credibility while recognising two shared interaction attributes. The adoption pathway, grounded in technology acceptance logic, links Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance with Perceived Usefulness and Continued Usage Intention. The credibility pathway, grounded in trust and source credibility reasoning, links Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution with

Ayurvedic Credibility and Source Transparency. The shared predictors represent interaction qualities relevant to both outcomes, while usability and safety communication distinguish the two pathways.




Figure 1. Conceptual framework for adoption and credibility pathways.

Source: Developed by the authors based on Davis (1989) and McKnight et al. (2002), with the proposed relationships adapted to the present study context.

Research Methodology

Research Design and Study Context

The study used a quantitative, cross-sectional, post-interaction user-evaluation design. The analysis reported in this article focuses on the survey evidence and the two multiple regression models. The setting was AyurVeda AI, a conversational assistant designed to provide Ayurvedic education and general wellness information. The system combines a curated knowledge base with retrieval-augmented generation and AI-agent functions. Users can enter natural-language questions, receive structured explanations with source references and continue the interaction through follow-up questions. Safety-related rules are intended to communicate limits and encourage professional consultation when a query extends beyond general wellness support. The system is not positioned as a confirmed diagnostic or treatment platform.

Participants, Sampling and Data Collection

The target population consisted of adults who directly interacted with the chatbot and were able to evaluate the response they received. Non-probability purposive sampling was supported by convenience-based recruitment because participation required actual system use. Eligible participants were at least 18 years old, provided informed consent, submitted at least one meaningful wellness-related query, reviewed the generated response and completed the post-interaction questionnaire. Responses were screened for eligibility, completeness, duplication and response quality. The screened

quantitative dataset contained 426 valid cases. This approach links each evaluation to a completed interaction rather than to a general opinion about artificial intelligence. The sampling strategy supports examination of relationships within the observed sample but limits statistical generalisation to all potential users.

Instrument and Construct Measurement

The study used a structured questionnaire to capture users' evaluations after direct interaction with the generative AI wellness assistant. The measures relevant to the present article represent the constructs included in the two regression models. The adoption model includes Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance as predictors of Perceived Usefulness and Continued Usage Intention. The credibility model includes Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution as predictors of Ayurvedic Credibility and Source Transparency. Responses were recorded using a five-point Likert scale ranging from Strongly Disagree to Strongly Agree. Construct scores were calculated from the questionnaire responses corresponding to each measured dimension. These scores were used in the regression analysis to estimate the independent contribution of each predictor to the respective outcome variable. The measures represent users' post-interaction perceptions of the chatbot and are interpreted within the educational and general wellness context of the system.

Statistical Analysis

Two multiple linear regression models were estimated to test the proposed relationships. The first model assessed the extent to which Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance predicted Perceived Usefulness and Continued Usage Intention. The second model assessed the extent to which Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution predicted Ayurvedic Credibility and Source Transparency. These predictor and outcome combinations correspond directly to the analytical models used in the study. Statistical significance was assessed at the 5% level. For each regression model, the model summary was used to report R, R^2 , adjusted R^2 , the standard error of the estimate and the Durbin-Watson statistic. Analysis of variance was used to determine whether the set of predictors significantly explained variation in the dependent variable. Predictor-level results were examined using unstandardised regression coefficients, standard errors, standardised β coefficients, t-values and significance levels.

Standardised β coefficients were used to compare the relative contribution of the predictors within each model. The regression results are interpreted as predictive associations based on cross-sectional post-interaction data and do not establish causal relationships. Ayurvedic Credibility and Source Transparency represents users' perception of the credibility, source visibility and traceability of the information presented by the chatbot. It does not constitute independent verification of textual accuracy, retrieval faithfulness, diagnostic validity, clinical safety or treatment effectiveness.

Results

Adoption Model

The first multiple regression assessed whether Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance predicted Perceived Usefulness and Continued Usage Intention. The model produced $R = .428$, $R^2 = .184$, adjusted $R^2 = .178$, a standard error of .916, and a Durbin-Watson statistic of 2.085. The predictors jointly explained 18.4% of the outcome variance, and the omnibus test was significant, $F(3, 422) = 31.631$, $p < .001$.

Table 1. Regression results for the adoption model

Panel A: Model Summary					
R	R ²	Adjusted R ²	Std. Error	Durbin-Watson	
.428	.184	.178	.916	2.085	
Panel B: ANOVA					
Source	Sum of Squares	df	Mean Square	F	p
Regression	79.647	3	26.549	31.631	< .001
Residual	354.200	422	.839		
Total	433.846	425			
Panel C: Coefficients					
Predictor	B	SE	β	t	p
Constant	1.168	.210	NA	5.564	< .001
Query Understanding and Response Relevance	.258	.047	.259	5.497	< .001
Response Clarity and Organisation	.168	.048	.164	3.474	.001
Usability and Practical Guidance	.185	.047	.179	3.950	< .001

Note. Dependent variable: Perceived Usefulness and Continued Usage Intention. Predictors: Query Understanding and Response Relevance, Response Clarity and Organisation, and Usability and Practical Guidance.

All three predictors made positive, significant independent contributions. Query Understanding and Response Relevance was strongest, $B = .258$, $\beta = .259$, $t = 5.497$, $p < .001$, supporting H1. Usability and Practical Guidance followed, $B = .185$, $\beta = .179$, $t = 3.950$, $p < .001$, supporting H3. Response Clarity and Organisation was also significant, $B = .168$, $\beta = .164$, $t = 3.474$, $p = .001$, supporting H2. The ranking places intent alignment at the centre of adoption-related evaluation.

Credibility Model

The second multiple regression assessed whether Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution predicted Ayurvedic Credibility and Source Transparency. The model produced $R = .364$, $R^2 = .132$, adjusted $R^2 = .126$, a standard error of .893, and a Durbin-Watson statistic of 2.072. The predictors jointly explained 13.2% of the outcome variance, and the omnibus test was significant, $F(3, 422) = 21.430$, $p < .001$.

Table 2. Regression results for the credibility model

Panel A: Model Summary					
R	R^2	Adjusted R^2	Std. Error	Durbin-Watson	
.364	.132	.126	.893	2.072	
Panel B: ANOVA					
Source	Sum of Squares	df	Mean Square	F	p
Regression	51.323	3	17.108	21.430	< .001
Residual	336.883	422	.798		
Total	388.206	425			
Panel C: Coefficients					
Predictor	B	SE	β	t	p
Constant	1.707	.201	NA	8.473	< .001
Query Understanding and Response Relevance	.107	.046	.113	2.313	.021
Response Clarity and Organisation	.201	.047	.207	4.245	< .001
Safety Guidance and Medical Caution	.177	.046	.182	3.835	< .001

Note. Dependent variable: Ayurvedic Credibility and Source Transparency. Predictors: Query Understanding and Response Relevance, Response Clarity and Organisation, and Safety Guidance and Medical Caution.

Response Clarity and Organisation was strongest, $B = .201$, $\beta = .207$, $t = 4.245$, $p < .001$, supporting H5. Safety Guidance and Medical Caution followed, $B = .177$, $\beta = .182$, $t = 3.835$, $p < .001$, supporting H6. Query Understanding and Response Relevance also contributed, $B = .107$, $\beta = .113$, $t = 2.313$, $p = .021$, supporting H4. Credibility was most closely associated with clear organisation and visible safety-conscious communication.

Comparative Pattern Across the Two Models

The regressions show different predictor priorities across the two outcomes. Query Understanding and Response Relevance was strongest for adoption, while Response Clarity and Organisation was strongest for credibility. Usability and Practical Guidance was specific to adoption, while Safety Guidance and Medical Caution was specific to credibility. The pattern supports the dual-path argument that being worth using and being perceived as credible are related but non-identical evaluations of a source-grounded generative AI system.

Discussion

Relevance as the Principal Driver of Adoption

The strongest adoption predictor was Query Understanding and Response Relevance. This result places intent alignment before interface novelty or response volume. Users appear more likely to value a conversational assistant when the answer corresponds directly to the purpose of the query. The finding is consistent with user-centred evaluation frameworks that treat understanding and information quality as core dimensions of effective health-related conversational AI (Abbasian et al., 2024; Tam et al., 2024). It also extends acceptance logic by showing that task value in generative AI depends on the system's ability to interpret free-text requests and sustain topical relevance. Traditional TAM research focuses on usefulness and ease-related beliefs (Davis, 1989), while generative AI introduces a dynamic antecedent to those beliefs: the quality of semantic alignment between user intent and generated response.

The magnitude of the relevance coefficient also suggests that polished language cannot compensate fully for a mismatch in topic or concern. This has practical importance for RAG systems. Retrieval and generation pipelines are often evaluated through technical measures such as contextual precision and faithfulness, yet end users encounter the result as a conversational answer. A response can be technically grounded and still provide limited perceived value if it does not resolve the user's information need. Recent RAG-based health education studies support the importance of answer relevancy, clarity and helpfulness as complementary performance dimensions (Baur et al., 2025; Kelly et al., 2025).

Clarity as the Principal Driver of Credibility

Response Clarity and Organisation was the strongest predictor of Ayurvedic Credibility and Source Transparency. This finding suggests that source display alone is insufficient for perceived credibility. Users need an explanation they can follow before they can judge whether cited information appears relevant, coherent and responsibly presented. Explainability scholarship in health-related large language models makes a similar distinction between access to information about a system and meaningful human understanding of its outputs (Mesinovic et al., 2025). A citation can increase traceability, but clarity determines whether the user can interpret what the citation is doing within the response.

The result also adds nuance to the role of transparency. Generative AI trust studies do not always find a uniform direct effect for transparency, indicating that trust formation depends on how transparency is operationalised and experienced (Huynh & Aichner, 2025). In the present setting, clarity may function as an enabling condition for source transparency. A well-organised response separates explanation, practical information, caution and references, allowing users to evaluate the information basis with less cognitive effort. This does not establish factual correctness. It shows that credibility judgements are partly shaped by the communicative form through which evidence is presented.

Usability and Safety as Outcome-Specific Drivers

The outcome-specific predictors provide the clearest distinction between the two pathways. Usability and Practical Guidance contributed to adoption but was not part of the credibility model. This aligns with technology acceptance research in which low-effort interaction and practical performance influence willingness to continue using a system (Davis, 1989; Shahzad et al., 2024; Venkatesh et al., 2003). A wellness assistant competes with search engines, websites and other conversational tools. Convenience and practical guidance can create immediate value even when the user does not make a strong judgement about the reliability of every source.

Safety Guidance and Medical Caution played a parallel role in the credibility pathway. Its positive coefficient indicates that caution was associated with stronger perceived credibility rather than weaker confidence. This is important for responsible-AI design because developers can fear that warnings, limitations and referral language may reduce engagement. The present result points in another direction: visible boundaries can signal responsible system behaviour. Health-care AI evaluation frameworks place safety, harm control, transparency and trustworthiness within the same governance space (Hua et al., 2025; Tam et al., 2024). Research documenting limitations in clinical reasoning also supports explicit caution when large language models operate near high-stakes decisions (Hager et al., 2024).

Adoption Is Not Credibility

Taken together, the findings support a dual-path interpretation of source-grounded generative AI. Adoption is most strongly connected to whether the system understands the user and provides relevant information. Credibility is most strongly connected to whether that information is communicated clearly, with safety guidance also making a meaningful contribution. The shared predictors show that the two evaluations are connected, while the different coefficient rankings and outcome-specific

predictors show that they should not be treated as one general attitude. This pattern echoes earlier information-systems research in which trust and technology-acceptance beliefs jointly influenced behaviour but retained conceptual distinctiveness (Gefen et al., 2003).

The distinction is especially important in health and wellness applications. Maximising continued use without attention to calibrated confidence can encourage over-reliance. Maximising caution without attention to relevance and usability can produce a system that is responsible in form but ignored in practice. Responsible design needs both pathways. Users need enough relevance and usability to obtain value, and enough clarity, source visibility and safety communication to understand the boundaries of that value. This dual requirement is consistent with contemporary evaluation frameworks that place usefulness, trustworthiness, safety and operational effectiveness in separate but connected domains (Abbasian et al., 2024; Hua et al., 2025).

Theoretical and Practical Implications

Theoretical Implications

The study offers three theoretical implications. First, it separates adoption-related evaluation from credibility-related evaluation in a source-grounded generative AI setting. This distinction builds on the long-standing integration of acceptance and trust research while adapting it to conversational systems that generate dynamic, source-linked content (Gefen et al., 2003; McKnight et al., 2002). Second, the results show predictor asymmetry across the two outcomes. Relevance is most important for adoption, while clarity is most important for credibility. This suggests that interaction attributes can change theoretical role depending on the user judgement being explained. Third, the positive association between safety guidance and credibility adds a responsible-trust dimension to the model. Clear limitations and caution may contribute to confidence by signalling that the system recognises boundaries. Trust in generative AI should be understood as calibrated confidence based on observable system behaviour rather than a simple objective of increasing reliance.

Practical Implications

The results support different design priorities for adoption and credibility. Developers seeking continued use should prioritise intent recognition, relevant retrieval, response focus, interactional convenience and practical educational value. Developers seeking credible source presentation should prioritise understandable response architecture, transparent source linkage, visible safety language

and clear professional-care boundaries. The two priorities should be implemented together rather than treated as substitutes. A highly usable system with weak source communication can encourage uncritical reliance, while a transparent system with poor relevance can fail to create practical value. For institutional deployment, evaluation should include both user-facing measures and technical audits. User ratings can show whether people understand and value the interaction, while expert review and retrieval testing are needed to assess factual accuracy, faithfulness and safety performance (Hager et al., 2024; Tam et al., 2024).

Limitations and Future Research

Several limitations define the interpretation of the findings. The cross-sectional design captures perceptions after a defined interaction and cannot establish causal effects or actual long-term use. The sample was obtained through purposive eligibility with convenience-supported recruitment, which limits population-level generalisation. The questionnaire measures self-reported perceptions and may reflect prior familiarity with AI, Ayurveda or digital wellness services. The measures were obtained through self-reported post-interaction evaluations, which may reflect respondents' prior familiarity with artificial intelligence, Ayurveda and digital wellness services. Future studies can validate the proposed relationships using independent samples, longitudinal designs and alternative measurement approaches. The credibility outcome reflects user judgement rather than independent validation of source authority, retrieval precision, factual faithfulness or clinical safety. Future research can test the dual-path model longitudinally, compare RAG and non-RAG systems, incorporate objective source-verification measures, examine actual repeated use and evaluate whether the pattern holds across languages, wellness domains and user groups.

Conclusion

Source-grounded generative AI is evaluated through more than one user judgement. Adoption was most strongly associated with query understanding and response relevance, while credibility was most strongly associated with response clarity and organisation. Usability supported adoption, and safety guidance supported credibility. Designing for continued use and designing for responsible confidence are connected but distinct tasks. A dual-path evaluation approach allows these requirements to be assessed without treating perceived usefulness as a substitute for credibility or perceived credibility as evidence of clinical accuracy.

References

- Abbasian, M., Khatibi, E., Azimi, I., Oniani, D., Shakeri Hossein Abad, Z., Thieme, A., Sriram, R., Yang, Z., Wang, Y., Lin, B., Gevaert, O., Li, L. J., Jain, R., & Rahmani, A. M. (2024). Foundation metrics for evaluating effectiveness of healthcare conversations powered by generative AI. *npj Digital Medicine*, 7, 82. <https://doi.org/10.1038/s41746-024-01074-z>
- Alelyani, T. (2025). Decoding trust in large language models for healthcare in Saudi Arabia. *Scientific Reports*, 15, 35276. <https://doi.org/10.1038/s41598-025-18404-0>
- Baur, D., Ansorg, J., Heyde, C. E., & Voelker, A. (2025). Development and evaluation of a retrieval-augmented generation chatbot for orthopedic and trauma surgery patient education: Mixed-methods study. *JMIR AI*, 4, e75262. <https://doi.org/10.2196/75262>
- Choudhury, A., & Chaudhry, Z. (2024). Large language models and user trust: Consequence of self-referential learning loop and the deskilling of health care professionals. *Journal of Medical Internet Research*, 26, e56764. <https://doi.org/10.2196/56764>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340. <https://doi.org/10.2307/249008>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., ... Wright, R. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust and TAM in online shopping: An integrated model. *MIS Quarterly*, 27(1), 51-90. <https://doi.org/10.2307/30036519>
- Hager, P., Jungmann, F., Holland, R., Bhagat, K., Hubrecht, I., Knauer, M., Vielhauer, J., Makowski, M., Braren, R., Kaissis, G., & Rueckert, D. (2024). Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nature Medicine*, 30(9), 2613-2622. <https://doi.org/10.1038/s41591-024-03097-1>
- Hua, Y., Xia, W., Bates, D., Hartstein, G. L., Kim, H. T., Li, M., Nelson, B. W., Stromeyer, C., King, D., Suh, J., Zhou, L., & Torous, J. (2025). Standardizing and scaffolding health care AI-chatbot evaluation: Systematic review. *JMIR AI*, 4, e69006. <https://doi.org/10.2196/69006>

10.48047/jocaaa.2026.35.10.03

ISSN:1521-1398 ,1572-9206

Huynh, M. T., & Aichner, T. (2025). In generative artificial intelligence we trust: Unpacking determinants and outcomes for cognitive trust. *AI & Society*, 40, 5849-5869. <https://doi.org/10.1007/s00146-025-02378-8>

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Chen, D., Dai, W., Chan, H. S., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), Article 248. <https://doi.org/10.1145/3571730>

Kelly, A., Noctor, E., Ryan, L., & van de Ven, P. (2025). The effectiveness of a custom AI chatbot for type 2 diabetes mellitus health literacy: Development and evaluation study. *Journal of Medical Internet Research*, 27, e70131. <https://doi.org/10.2196/70131>

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W. T., Rocktaschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474.

McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research*, 13(3), 334-359. <https://doi.org/10.1287/isre.13.3.334.81>

Mesinovic, M., Watkinson, P., & Zhu, T. (2025). Explainability in the age of large language models for healthcare. *Communications Engineering*, 4, 128. <https://doi.org/10.1038/s44172-025-00453-y>

Shahzad, M. F., Xu, S., & Javed, I. (2024). ChatGPT awareness, acceptance, and adoption in higher education: The role of trust as a cornerstone. *International Journal of Educational Technology in Higher Education*, 21, 46. <https://doi.org/10.1186/s41239-024-00478-x>

Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., Payne, P., Seneviratne, M., Gamble, P., Kelly, C., Babiker, A., Scharli, N., Chowdhery, A., Mansfield, P., Demner-Fushman, D., ... Natarajan, V. (2023). Large language models encode clinical knowledge. *Nature*, 620, 172-180. <https://doi.org/10.1038/s41586-023-06291-2>

Tam, T. Y. C., Sivarajkumar, S., Kapoor, S., Stolyar, A. V., Polanska, K., McCarthy, K. R., Osterhoudt, H., Wu, X., Visweswaran, S., Fu, S., Mathur, P., Cacciamani, G. E., Sun, C., Peng, Y., & Wang, Y. (2024). A framework for human evaluation of large language models in healthcare derived from literature review. *npj Digital Medicine*, 7, 258. <https://doi.org/10.1038/s41746-024-01258-7>

10.48047/jocaaa.2026.35.10.03

ISSN:1521-1398 ,1572-9206

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. MIS Quarterly, 27(3), 425-478. <https://doi.org/10.2307/30036540>

Wang, L., Wan, Z., Ni, C., Song, Q., Li, Y., Clayton, E., Malin, B., & Yin, Z. (2024). Applications and concerns of ChatGPT and other conversational large language models in health care: Systematic review. Journal of Medical Internet Research, 26, e22769. <https://doi.org/10.2196/22769>