

When Source Disclosure Matters:
User perceptions of generative AI authorship in landing page content

Hira Mirza

A thesis submitted to the Auckland University of Technology
in partial fulfilment of the requirements of the degree of
Master of Communication Studies (MCS)

2026

School of Communication Studies
Faculty of Design and Creative Technologies

Abstract

The widespread adoption of generative AI (GenAI) in digital content marketing has outpaced the understanding of how users evaluate it. While marketers deploy GenAI to produce personalised content at scale, user perceptions of authenticity, credibility, and value remain poorly understood. One particular under-researched area is conversion-oriented environments where content must inform, persuade, and establish trust simultaneously. Existing literature on user evaluations of digital content addresses GenAI content perception broadly, but little empirical work examines how source disclosure and functional context together shape user preferences. This study investigates user preferences for GenAI-authored versus human-authored content on website landing pages, a persuasion-critical environment where content credibility directly influences decision-making. Using a 2x2 mixed comparative experimental design, 95 participants recruited via CloudResearch Connect evaluated landing page content under both disclosed and undisclosed authorship conditions to assess how source awareness affects content evaluations. The findings reveal a central paradox. Users evaluate GenAI-authored content favourably when it is informative and well-structured, yet prefer human authorship once the source is disclosed. Disclosure was found to activate relational scrutiny, such as authenticity and human-likeness, without significantly degrading functional evaluations of content quality, persuasiveness and credibility. Rather than confirming AI aversion, results show that human favouritism operates alongside, not instead of, positive GenAI evaluations. The study proposes a dual evaluative framework to account for this. Findings are interpreted through this lens as reflecting functional competence when GenAI content is judged on its informational merit, and relational authenticity when human authorship is valued for its perceived intentionality and authority. Critically, user acceptance of GenAI content is conditional rather than categorical, shaped by task, type, digital environment, and the level of user trust at stake. These findings offer empirical grounding for marketers navigating GenAI adoption and for policymakers designing authorship disclosure standards.

Table of Contents

Abstract	i
Table of Contents	ii
List of Figures	v
List of Tables	vi
Attestation of Authorship	vii
Artificial Intelligence Declaration Template	viii
Acknowledgements	x
Chapter 1 Introduction	1
1.1 Overview	1
1.2 Generative AI in digital marketing	1
1.3 Problem statement and research gap	2
1.4 Research aims and questions	5
1.5 Study design	5
1.6 Significance of the study	6
1.7 Thesis structure	8
Chapter 2 Literature review	9
2.1 Chapter overview	9
2.2 Generative AI: Conceptual foundations	10
2.3 Evaluative dimensions of digital marketing content	10
2.3.1 <i>Source and message credibility</i>	13
2.4 User evaluations of GenAI-authored and human-authored content	14
2.5 Authorship disclosure and source-based bias	16
2.5.1 <i>Algorithm aversion</i>	16
2.5.2 <i>Human favouritism</i>	18
2.6 Contextual moderators of GenAI content evaluation	19
2.7 User attitudes and prior experience with GenAI	20
2.8 Chapter conclusion	21
Chapter 3 Methodology	23
3.1 Chapter overview	23
3.2 Philosophical assumptions	23
3.2.1 <i>Source Credibility Theory</i>	24
3.3 Mixed methods approach	25

3.4	Research design.....	26
3.5	Participants	28
3.6	Survey.....	30
3.6.1	<i>Survey design</i>	30
3.7	Data cleaning, storage and processing	31
3.7.1	<i>Data validity</i>	32
3.8	Quantitative data analysis	32
3.9	Qualitative data analysis	33
3.10	Chapter conclusion.....	34
Chapter 4 Results.....		35
4.1	Chapter overview.....	35
4.2	Sample characteristics	35
4.3	Content evaluations and user preferences without authorship disclosure (RQ1)	36
4.4	Impact of authorship disclosure on content evaluations and user preferences (RQ2)	40
4.5	Changes in ratings across disclosure conditions.....	43
4.6	Qualitative findings	46
4.7	Order effects	49
4.8	Participants' GenAI experience and attitudes	51
4.9	Chapter conclusion.....	55
Chapter 5 Discussions		57
5.1	Chapter overview.....	57
5.2	GenAI-authored content performs favourably when authorship is unknown ..	57
5.3	Relational constructs remain suspended in the absence of source cues	59
5.4	Authorship disclosure shifts rather than reduces GenAI content evaluations.	61
5.4.1	<i>Dual evaluative framework</i>	62
5.5	Context and familiarity shape responses to GenAI-authored content.....	63
5.6	Chapter conclusion.....	64
Chapter 6 Conclusions		66
6.1	Chapter overview.....	66
6.2	Summary of key findings	66
6.3	Contributions to knowledge and theory	67
6.4	Implications for policy and practice	68
6.5	Limitations of the study.....	69
6.6	Directions for future research	71

6.7	Concluding reflection	72
References		I
Appendix		XXIII
Appendix A: AUTECH Approval Letter		XXIII
Appendix B: Participant Information Sheet		XXIV
Appendix C: Participant Recruitment on CloudResearch Connect		XXVI
Appendix D: Survey Consent Preview		XXVII
Appendix E: Survey		XXVIII

List of Figures

Figure 4.1 Proportional Distribution of Ratings for the Human-authored Page and GenAI-authored Page Recorded in Survey Section 1 (Without Authorship Disclosure)	38
Figure 4.2 Proportional Distribution of Preferences for the Human-Authored Page and GenAI-authored Page Recorded in Survey Section 2 (Without Authorship Disclosure).....	40
Figure 4.3 Proportional Distribution of Preferences for the Human-authored Page and GenAI-authored Page Recorded in Survey Section 2 (With Authorship Disclosure)	43
Figure 4.4 Proportional Distribution of Ratings for Human-Authored Page Recorded in Survey Section 1 (Disclosed Versus Undisclosed Authorship)	44
Figure 4.5 Proportional Distribution of Ratings for GenAI-Authored Page Recorded in Survey Section 1 (Disclosed Versus Undisclosed Authorship)	46
Figure 4.6 Proportional Distribution of Ratings for the Webpages Among Participants who Viewed the Human-Authored Page First Recorded in Survey Section 1	50
Figure 4.7 Proportional Distribution of Ratings for the Webpages Among Participants who Viewed the GenAI-authored Page First Recorded in Survey Section 1	51
Figure 4.8 Distribution of Self-Reported GenAI Usage Frequency Among Participants Recorded in Survey Section 3	53
Figure 4.9 Proportional Distribution of Participants' Attitudes Towards Generative AI Use in Marketing Recorded in Survey Section 3	54

List of Tables

Table 4.1 Descriptive Statistics for Webpage Ratings Without Authorship Disclosure.....	37
Table 4.2 Descriptive Statistics for Webpage Ratings With Authorship Disclosure	41
Table 4.3 Descriptive Statistics for Webpage Ratings Across Conditions	42

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Signed:

Name: Hira Mirza

Date: 2 June 2026

Artificial Intelligence Declaration Template

<i>Specify the chapter number(s) in your research proposal/research component where AI was used.</i>	
Chapter number(s):	Chapters 1 to 6.
<i>Briefly explain why the AI tool was used (e.g., idea generation, summarization, data analysis, image creation, etc.).</i>	
Purpose of AI Use:	AI was used for 4 main purposes throughout the thesis. 1. Stimulus material generation for experimental research design (ChatGPT-4): ChatGPT was used to generate a website landing page as part of the research methodology. The study employed a between-subjects design in which participants were exposed to one of two versions of a landing page (one GenAI-authored and one human-written) in order to examine answer the research question (See Section 1.4). The GenAI-authored version was produced using ChatGPT and constitutes a research instrument, not a contribution to the written thesis itself. 2. Literature identification (Consensus): Consensus was used to locate relevant academic literature and identify key sources pertaining to the research topic. 3. Research support and consistency review (ChatGPT and Claude): ChatGPT and Claude was used at various stages to review sections of the thesis for internal coherence and alignment with the stated research questions. The tools assisted in identifying gaps in the argument, suggesting structural improvements, organising ideas into themes, and summarising selected sources to support the analysis. 4. Proofreading (Grammarly): The tool was used to check the document for spelling, grammatical accuracy, and sentence-level clarity.
<i>Clearly state the AI tools used.</i>	
AI Tool(s) Used:	ChatGPT, Consensus, Claude, Grammarly
<i>Prompts or queries entered into the AI tool (only the instructive portion of the prompt is required). Include all relevant variations.</i>	
Prompts Used:	Some prompt examples include: ChatGPT: "Create a website landing page for a real estate agent targeting first home buyers." Consensus: "What does the research say about user perceptions of AI-generated content in digital marketing contexts? "What are the most cited studies on "AI aversion"?" Claude: "Review the following section of my thesis and identify any inconsistencies or gaps in relation to my research questions." "Does the argument in this section align with the research questions stated earlier?"
<i>A summary or description of the AI-generated output. (Optional: attach output in an appendix.)</i>	
Output Received:	ChatGPT: A GenAI-authored website landing page used as stimulus material in the experimental research design, attached as part of Appendix E: Survey.

Consensus: A curated list of academic sources relevant to the research topic, accompanied by AI-generated summaries used to inform decisions about which sources to read in full.

ChatGPT and Claude: Structural and editorial feedback on thesis sections; identification of logical gaps and inconsistencies relative to the research questions; thematic organisation of ideas; and brief summaries of selected academic papers to support the researcher's own analysis.

Grammarly: Sentence-level corrections and suggestions relating to spelling, grammar, punctuation, and clarity.

Describe how the AI output was used, edited, or transformed.

**Post-AI
Processing
Methods:**

ChatGPT-4: The landing page generated by ChatGPT was reviewed and assessed for suitability as a research stimulus. It was used as generated, as methodological integrity required that it represent authentic GenAI-authored content for comparison against the human-written version.

Consensus: All literature surfaced through Consensus was evaluated independently. Sources deemed relevant were read in full, and only those directly applicable to the research topic were selected for use. The researcher identified overarching themes and research gaps through their own analysis of the material.

ChatGPT and Claude: All feedback and summaries produced by the tools were reviewed carefully before any action was taken. Their outputs informed the researcher's thinking and revision decisions, but all writing and argumentation remained the researcher's own.

Grammarly: Proofreading suggestions from Grammarly were reviewed individually and accepted or rejected at the researcher's discretion.

Use of AI tools while writing the thesis:

Throughout the thesis, AI was engaged as a thinking partner, not an author. At every stage, the researcher retained full intellectual ownership, and used these tools to pressure-test ideas, identify blind spots, and improve rigour, rather than to generate or substitute original thought.

**STUDENT
DECLARATION**

By signing you are confirming that the AI use stated in the table(s) above are accurate and follow AUT's recommended guidelines detailed in the AI Hub and Postgraduate Handbook.

**Student
Name:** Hira Mirza

Signature:

Date: 2 June 2026

Acknowledgements

This thesis has been nothing short of a test. One that I have, in my own eyes at least, passed satisfactorily. Nevertheless, preparing for this test and eventually finding the courage to see it through were tasks I could not have completed alone.

Before I thank the people in my life, I want to thank the One who gave me life. The Most Merciful and All-Knowing who sent me on this short journey to Earth. He is the One who motivates me to live my dreams but also challenges me to conquer my fears. Without His Light, the majestic yet gleaming mountains would never amaze me, the sparkles in the night sky would never intrigue me, and the words on my document would never make sense. Finding Him through my study has given me a direction that has left me speechless.

At the top of the long list of people I must now thank is my eldest, Moosa. His resilience and patience through our initial years in Auckland both encourage and humble me. Without him, I would have lost the motivation to do what I cherish the most. My youngest, Hamza, truly matched his brother's bravery and gave me the grounding I needed to carry this thesis to the finish line. He came into my life mid-journey and changed the shape of it entirely. I paused unexpectedly and reworked the timeline. Yet it was precisely because of my boys that I couldn't give up. With two pairs of eyes watching, quitting was never truly an option.

Beyond my children, I thank my parents who have been the fuel to my fire, near and far. It is only through their prayers and affection that I got this far in life. While I know I can never repay them enough, this thesis is my gift in return for their unspoken trust. This is my way of saying, "Mama and Abu jee, this is for every prayer that you whispered on my behalf without ever asking for anything in return".

That same warmth carried itself, unexpectedly, into another family entirely. A shared love for words with my late father-in-law lit a new flame of encouragement in me at a time I needed it most. I hope and pray that Baba's soul receives the gratitude I will always carry for him, and for my mother-in-law, who has always stood by us.

At the very centre of it all stood the person who has known me longest in this chapter of my life. My dear husband, Asif, has helped me explore new depths of my potential, for which I will always be grateful to him. With both his presence and quiet strength, I feel I can achieve anything my heart desires.

And finally, to the women who never let me forget who I was. Throughout this time, my dear sisters Rabia, Sarah and Ayesha, and my lovely friend Fizza, have been immovable rocks of support, picking me up when I fell, one too many times. Without them, the version of me, or of this thesis, wouldn't exist. They not only supported the work but also held me up while I was doing it.

On the academic front, my supervisor, Matt, deserves more credit than a paragraph can hold. Across the delays and the self-doubt, his support remained steady and constant. His calm, in particular, was something I borrowed often. He was measured in his guidance, generous with his time, and unfailing in his encouragement. Matt will always be part of my prayers, and his pep talks will go a long way with me, even outside of academic work.

I am equally grateful to the postgraduate staff at AUT, whose flexibility and understanding made it possible for me to continue my studies as life unfolded around me. Particular thanks to Gudrun Frommherz, whose support in organising my maternity break was the key ingredient in the completion of my Masters programme.

All these people have contributed to something that means more to me now than when I began studying at AUT in 2021. God willing, years from now I will sit in my living room and reflect on my two Master's degrees, one in 2011 from Pakistan, my home country, and one in 2026 from New Zealand, my new home. What I will remember most is the people who walked beside me through the wins and carried me through the losses I thought were mine alone to carry.

Chapter 1 Introduction

1.1 Overview

Artificial intelligence (AI) has rapidly evolved from a specialised technical tool into a central feature of everyday life, shaping how information is created, distributed, and consumed (Sengar et al., 2024). Within this broader transformation, advances in generative artificial intelligence (GenAI) have accelerated change by enabling machines to produce content that closely resembles human work. These systems can generate text, images, and other media using learned patterns from large datasets, fundamentally changing how communication is produced in both personal and organisational contexts. By lowering technical barriers, these tools have enabled a broad range of users to create human-like content with minimal effort, contributing to rapid adoption across industries (Banh & Strobel, 2023). The widespread release of tools such as ChatGPT, Midjourney, and GitHub Copilot has marked a turning point in the accessibility and visibility of generative technologies (Bengesi et al., 2023). As a result, GenAI is increasingly embedded in everyday communication practices and business operations, particularly in digital marketing contexts where content creation plays a central role.

1.2 Generative AI in digital marketing

Within business contexts, the impact of GenAI is especially evident in marketing, where communication with audiences is central to organisational success (Kshetri et al., 2024). Marketing broadly refers to activities aimed at creating, communicating, and delivering value to customers to achieve a desired response (Kotler, 1972). Digital marketing represents a subset of these activities conducted through online channels such as websites, social media, and email, to reach and engage customers (Desai, 2019). Within this domain, content marketing focuses on creating and distributing valuable, relevant content to attract, engage, and influence target audiences (Koob, 2021). While these terms are closely related and often used interchangeably, this thesis uses them to distinguish between the broader strategic function (marketing), the channel of delivery (digital marketing), and the specific practice of content creation (content marketing).

GenAI has become increasingly integrated into digital and content marketing activities (Kshetri et al., 2024) with organisations now using AI tools to produce website copy, social media posts, promotional messaging, and customer communications at scale

(Guzman & Lewis, 2024; Sands et al., 2024). These tools enable rapid content generation, adaptation for different audiences, and consistent messaging, thereby improving efficiency while maintaining relevance (Banh & Strobel, 2023; Kshetri et al., 2024). The role of GenAI is most pronounced in content marketing, where effectiveness depends on delivering clear, relevant, and engaging information across multiple digital touchpoints (Hollebeek & Macky, 2019; Joan Isibor et al., 2025). Digital marketing content is defined as valuable, reliable, and audience-focused information shared through platforms such as websites, blogs, and social media. It plays a critical role in shaping user engagement, trust, and brand (Jacob & Johnson, 2021). By enabling scalable, adaptable content production, GenAI tools have enhanced organisations' ability to deliver personalised, timely communication.

At the same time, integrating GenAI tools into content creation workflows introduces new challenges. The ability of generative systems to produce highly realistic and human-like outputs makes it increasingly difficult for users to distinguish between human-authored and machine-generated content (AI-kfairy et al., 2024). Therefore, while these tools enable scalable, personalised, and high-quality content creation, they also raise concerns regarding trust, authenticity, and transparency in digital communication (Morales, 2025). As such, GenAI-authored content may perform well in terms of clarity and efficiency, but uncertainty around its origin and credibility may influence how users interpret and respond to it.

1.3 Problem statement and research gap

Existing research presents a paradox regarding how users evaluate and perceive GenAI-authored content, particularly with respect to credibility, authenticity, and trust. On one hand, users frequently rate GenAI content as clear, relevant and high-quality, in some cases preferring it over human-authored alternatives (Park et al., 2024; Stamkou et al., 2025; Y. Zhang & Gosline, 2023). These findings are consistent across multiple domains, including marketing, education, media, and e-commerce (Gu et al., 2024; Hussain, 2024; Zhe & Srijinda, 2024). On the other hand, this positive evaluation is often undermined when authorship is disclosed. When content is explicitly labelled as GenAI-authored, users tend to express greater scepticism and lower perceived authenticity (Camilleri, 2024; Parshakov et al., 2025), reflecting a broader psychological response described as 'algorithm aversion' (Dietvorst et al., 2015). Competing explanations suggest that this shift may instead reflect a source bias in favour of human authors, which Y. Zhang and Gosline (2023) refer to as 'human favouritism'. Despite agreement that disclosure influences preferences, there remains no clear consensus regarding the impact of content performance versus the awareness of its

source. The inconsistency in user reactions also limits our understanding of whether authorship fundamentally alters user responses or simply reframes them, and how factors beyond source awareness come into play.

Contextual factors further shape user perceptions. Prior research in emotionally driven contexts suggests more pronounced preferences for human-authored content that requires empathy, personal experience, or relational connection (Jain et al., 2024; Kirk & Givi, 2025). In contrast, users demonstrate greater acceptance of GenAI content in functional or utilitarian tasks, where clarity, efficiency, and information delivery are prioritised over emotional depth (Castelo et al., 2019; Kashive et al., 2020). In such contexts, preferences are more likely to be grounded in perceived usefulness and performance rather than the human origin of the content. At the same time, emerging research indicates that even in more complex settings, users may respond positively to GenAI-authored content when it delivers both cognitive value and persuasive effectiveness (Kubovics, 2025b; Sew et al., 2025; Zhe & Srijinda, 2024). More broadly, research on contextual factors remains inconsistent across contexts, with calls for more systematic investigation of how these effects vary across different decision environments (Burton et al., 2020), particularly under different authorship conditions. These inconsistencies may also reflect a temporal shift in user perceptions. Much of the foundational research on algorithm aversion was conducted in the era of conventional, predictive algorithms, before GenAI tools like ChatGPT became widespread (Burton et al., 2020; Kaufmann et al., 2023). As user exposure increases, expectations and evaluative criteria are likely to evolve, with greater AI familiarity potentially reshaping acceptance and reducing resistance (Jia et al., 2025). This temporal dimension requires continued examination and underscores the importance of investigating user responses under current conditions of widespread GenAI use.

Another limitation in the existing literature on human favouritism and algorithm aversion is its concentration in a narrow set of digital contexts. Much of this work focuses on social media, journalism, and creative content settings, where algorithms curate feeds, select news, or generate creative outputs. In journalism, studies have examined how audiences and professionals evaluate algorithm-written versus human-written articles, highlighting concerns about quality, credibility, and trust (Jung et al., 2017). Beyond journalism, a substantial body of work examines how algorithmic systems shape user experiences through personalised feeds and recommendations in social media environments (Metzler & Garcia, 2024). These studies investigate how users perceive, interpret, and respond to algorithmic influence, including its effects on trust, misinformation, and platform engagement (Bogert et al., 2022; Swart, 2021; Ytre-Arne & Moe, 2021). In parallel, research in creative and cognitive domains has explored

responses to algorithmic assistance in tasks such as writing and idea generation, identifying both aversion and appreciation depending on the nature of the task (Bogert et al., 2022).

Current research shows limited clarity on how user responses to algorithmic outputs translate into structured, goal-oriented digital marketing contexts, where content is evaluated not only for engagement but also for its ability to drive action. Although existing studies demonstrate GenAI's ability to produce functional marketing content, they focus more on the technical or textual quality of the output than on audience interpretation (Hartmann et al., 2025). For example, research on structured content formats shows that GenAI can replicate the conventions, style, and persuasive logic of professional copywriting (Ouyang et al., 2022). However, less attention has been given to how users perceive this content, especially when evaluating its authenticity, quality, trustworthiness, or authorship. To address this gap, website landing pages provide a particularly relevant and functional context for examining user perceptions of GenAI content in goal-oriented digital marketing platforms. A landing page is a standalone entry point designed to drive a specific conversion, such as an enquiry or purchase (Thomas & Mathew, 2023). It serves as a central touchpoint within digital customer journeys, involving search engines, social media, and email, while presenting clear, structured information to support user decision-making. The primary role of a landing page remains task-oriented and outcome-driven, but it must also establish trust, credibility, and perceived authenticity within a limited timeframe, particularly in high-stakes contexts (Huang et al., 2019). Established web credibility principles further emphasise clarity, structured information, and identifiable authorship cues (Fogg et al., 2001) on landing pages, making them a context in which both content performance and source-related signals are salient. Despite their functional and contextual relevance, landing pages are rarely isolated as the unit of analysis in studies examining responses to GenAI-authored content, which focus mainly on broader aspects of digital communication, such as content quality, recommendation systems, or personalisation processes (Çetinkaya et al., 2024; Schickmair et al., 2025; Y. Zhang & Gosline, 2023). By focusing on website landing pages, this study examines how users perceive GenAI content in a functional digital marketing context and how their preferences change when authorship is disclosed. In doing so, it considers whether user choices are shaped primarily by the perceived performance of the content or by awareness of its source.

1.4 Research aims and questions

This thesis aims to examine how users perceive GenAI-authored and human-authored digital marketing content, and how these perceptions are influenced by authorship disclosure. As discussed, prior research suggests greater acceptance of GenAI-authored content in functional contexts and mixed responses when authorship is disclosed. This study investigates whether these patterns extend to structured, goal-oriented marketing environments, such as landing pages, where content is designed to communicate information clearly and build user trust to support decision-making. Specifically, the study compares user preferences for GenAI and human-authored landing page content under conditions in which authorship is not disclosed, and examines how these preferences change when authorship is made explicit. In doing so, it seeks to determine whether user perceptions are primarily driven by the perceived quality and usefulness of the content, or by awareness of its source.

The following research questions guide this investigation:

Research question 1: Do users prefer GenAI-authored or human-authored website landing page content when authorship is not disclosed?

Research question 2: How does authorship disclosure influence users' preferences for GenAI-authored and human-authored website landing page content?

1.5 Study design

This study employs a mixed methods survey design that integrates quantitative and qualitative data to examine user perceptions and preferences regarding GenAI and human-authored website landing page content. Data were collected through an online survey hosted on Qualtrics, with adult English-speaking participants recruited via the CloudResearch Connect platform. Participants were presented with two versions of a realistic but fictional website landing page designed for first-home buyers. Both pages were identical in structure, layout, and visual presentation, ensuring that differences in evaluation could be attributed to content and its authorship rather than design. One version was generated using ChatGPT-4 (*GPT-4 | OpenAI, n.d.*), while a human digital content specialist wrote the other. A digital content specialist is a marketing professional responsible for planning, creating, and optimising content across digital platforms to support brand objectives and business outcomes (Hollebeek & Macky, 2019). More details on both authors are provided in Section 3.4.

The study employed a mixed experimental survey design, incorporating a within-subject comparison of the two webpages and a between-subject manipulation of

authorship disclosure. This design directly aligns with the two research questions. The within-subject comparison allowed participants to evaluate and compare the GenAI-authored and human-written webpages, supporting the examination of user preferences when authorship is not disclosed. The between-subject disclosure manipulation allowed the study to examine whether making authorship visible influenced those preferences. Participants were randomly assigned to one of two conditions: undisclosed authorship, where no information about the content source was provided, or disclosed authorship, where the origin of each page was explicitly identified.

Participants evaluated the webpages by answering survey questions across nine perception constructs: authenticity, credibility, persuasiveness, author knowledge, human-likeness, writing connection, brand genuineness, product interest, and overall quality. These are discussed in more detail in Chapter 2. In addition, open-ended questions were included to capture participants' reasoning and qualitative responses to the website content. These measures enabled both direct comparison between the GenAI-authored and human-authored pages and interpretation of the factors shaping participant preferences. Moreover, quantitative data were analysed using descriptive statistics to compare evaluations across content types and disclosure conditions. Qualitative responses were analysed thematically to identify recurring patterns and explanatory insights. The study is intentionally scoped to textual content within a single product context and involves one-time exposure to the webpages. By controlling for layout and presentation, the design isolates the effect of authorship and disclosure on user perceptions and preferences, providing a focused examination of how users respond to GenAI-authored content in a structured digital marketing environment.

1.6 Significance of the study

As GenAI reshapes the landscape of digital communication, the significance of this study extends across practice, policy, and scholarship. In particular, its deepening role in marketing practice raises crucial questions about how users perceive and evaluate GenAI-authored material published in digital campaigns and brand communication (Wahid et al., 2023). Research consistently demonstrates that user perceptions of credibility, authenticity, and trust directly shape engagement and conversion outcomes in digital environments (Castelo et al., 2019). By isolating how these perceptions shift under different disclosure conditions, this study offers practitioners empirically grounded guidance on when and how to deploy transparency without eroding audience trust. Beyond practice, the findings also contribute to policy debates about AI transparency and responsible use in commercial contexts, which have often relied

more on prescription than empirical evidence (Corrêa et al., 2024; Raso, 2024). This study brings nuanced, context-specific evidence to those debates, and in doing so, advances the conversation from principle to practice.

From a practical standpoint, the study addresses a meaningful contextual gap by focusing specifically on website landing pages. These pages represent a goal-oriented, high-stakes communication environment that differs substantially from the social media content or email formats examined in prior research. By situating GenAI-authored content within this functional context, the study offers direct insight into how authorship disclosure influences user preferences at the precise point where user attention is concentrated, and conversion decisions are made. This has direct implications for digital marketing strategy, particularly for practitioners balancing efficiency gains against the risk of losing user trust (Davenport et al., 2020). In doing so, it equips practitioners with a more grounded understanding of how users evaluate GenAI-authored content in a functional digital setting where both informativeness and perceived authenticity drive decision-making.

Theoretically, the study makes two original contributions to the field. The first is the identification of nine constructs, grounded in the literature, to measure user evaluation of GenAI-authored content in digital marketing contexts. These include authenticity, credibility, persuasiveness, author knowledge, writing connection, human-likeness, brand genuineness, product interest, and overall quality. Derived from source credibility and message credibility research and adapted to the landing page context, these constructs offer the field an evaluative framework not currently available in prior research. Where existing studies do not combine these constructs, this study proposes a purpose-built framework that future researchers can apply, test, and refine as a measurement instrument across other GenAI communication contexts.

Emerging from this empirical foundation, the second and more significant contribution is the development of a dual evaluative framework. This framework integrates source credibility, algorithm aversion, and human-content favouritism to develop an original contribution to explain how users simultaneously judge what content says and who produced it. Specifically, the framework combines considerations of functional competence, or the extent to which content is clear, credible, and persuasive, with relational authenticity, or the extent to which content signals human presence and emotional connection. Crucially, the framework highlights the dynamic interaction between the two lenses through which users view digital content within context-specific conditions. This extends existing theoretical models, which have tended to treat these dimensions in isolation.

Together, the study contributes to communication theory, digital marketing practice, and AI governance. For marketers and businesses, the framework clarifies when GenAI-authored content can perform on par with human-written material, and when disclosure becomes a liability rather than a trust-building gesture. For regulators, the nine constructs and supporting evidence provide empirical grounding for transparency requirements that are currently being debated with insufficient evidence to keep pace with technological change. For content consumers, the findings reveal how knowing that content was GenAI-authored can reshape trust and credibility judgements, even when the content is of comparable quality. In combination, these contributions offer a timely, evidence-based foundation for how organisations, regulators, and users can navigate the expanding role of GenAI in digital communication.

1.7 Thesis structure

This thesis is structured into six chapters. Chapter 1 introduces the research problem, outlines the study's objectives, and presents the research questions guiding the investigation. Chapter 2 reviews the relevant literature on user perceptions of GenAI-authored and human-written digital content, establishing the theoretical foundation for the study. Chapter 3 details the research methodology, including the study design, data collection procedures, and analytical approach. Chapter 4 presents the empirical findings, reporting both quantitative and qualitative results in relation to the research questions. Chapter 5 discusses the key findings in light of existing literature on source credibility, AI aversion and human favouritism. Chapter 6 concludes the thesis and outlines the study's contributions and practical implications. It also acknowledges its limitations and suggests directions for future research.

Chapter 2 Literature review

2.1 Chapter overview

The growing adoption of GenAI tools in commercial contexts has raised perceptual and behavioural questions that are increasingly relevant to digital marketing practice. This chapter reviews how users evaluate digital marketing content, how authorship disclosure alters those evaluations, and how contextual factors, user attitudes and prior beliefs about AI strengthen or diminish those effects. The existing literature has produced consistent findings across social, journalistic, and informal commercial contexts. Yet structured, conversion-oriented environments such as website landing pages remain underexplored. A further limitation concerns the tendency to treat evaluation and preference as interchangeable. Most studies infer user choice from evaluative measures alone, such as trustworthiness or quality ratings, thereby obscuring the possibility that a user may assess GenAI-authored content less favourably yet still choose to engage with it. These gaps inform the central questions this thesis addresses: whether users prefer GenAI-authored content or human-authored alternatives, and whether the perceived content source shapes how they evaluate and respond to it.

To examine these questions, the review draws on research in digital media marketing, source credibility theory, and consumer psychology. Literature was identified through two stages. The AUT Library's discovery service, which indexes EBSCO and other databases, was used during topic development to explore the emerging GenAI literature. Keywords included generative AI, AI disclosure and transparency, user perceptions and evaluations of GenAI-authored content, landing page credibility, digital content marketing, large language model-generated content, and AI content detection. As the review's focus narrowed, Consensus, an AI-assisted tool, was used for more targeted searches. Consensus synthesizes findings from the Semantic Scholar database rather than performing conventional keyword searches and was queried using natural-language questions built around the same key concepts rather than Boolean strings. Most literature was published from 2023 onward, reflecting the emergence of ChatGPT-driven GenAI research. Earlier sources were retained where they informed established literature on algorithmic trust and automated decision-making, providing theoretical continuity for a topic that has only recently generated its own body of research. Across these bodies of work, nine perceptual constructs emerge as the most salient dimensions through which users seem to evaluate GenAI-authored

digital marketing content: authenticity, persuasiveness, content quality, product interest, brand genuineness, credibility, author knowledge, human-likeness, and writing connection. Together, they form the evaluative framework of this study and are explained in this chapter.

2.2 Generative AI: Conceptual foundations

GenAI occupies a distinct position within the broader AI hierarchy, built on deep learning architectures. What sets it apart from other AI systems is its generative capacity, or the ability to produce previously unseen outputs that embody creativity in computational form (Bordas et al., 2024; Yao et al., 2024). Common architectures include diffusion models for image generation, generative adversarial networks (GANs) for image and video synthesis, large language models (LLMs) for text generation, and multimodal systems capable of producing diverse output types (Kalota, 2024).

Central to this research is the LLM's capacity to produce human-like text through Natural Language Processing (NLP), the branch of AI concerned with how computers process and interpret human language (Bandi et al., 2023). LLMs are built on artificial neural networks trained on vast quantities of unstructured text data, enabling them to generate contextually coherent, human-like written outputs (Yao et al., 2024). This is the type of GenAI-authored content examined throughout this study. The most widely adopted LLM-based services include ChatGPT (OpenAI), which holds approximately 60% of the global LLM market share, followed by Gemini (Google) at 24% and Grok (xAI) at 3.4% (Ermut, 2026). ChatGPT's dominance in both adoption and accessibility (Wahid et al., 2023) makes it the most representative tool for examining GenAI-authored content in this context, directly informing its selection in the methodology.

2.3 Evaluative dimensions of digital marketing content

Digital content marketing is one of the key areas in which GenAI is increasingly being applied, particularly for generating personalised content for targeted audiences (Grewal et al., 2025). Given that this study focuses on text-based marketing content, this section examines how textual digital content is evaluated in marketing and related domains, and then narrows its focus to GenAI-authored and human-written content in subsequent sections. The following sections highlight nine evaluative constructs that are salient to this study and are examined in the literature review: authenticity, persuasiveness, product interest, brand genuineness, overall content quality, credibility, author knowledge, human-likeness, and writing connection. These nine constructs were selected based on their relevance to the research questions and prior use in studies underpinned by source credibility theory or GenAI-authored content.

Effective digital content supports marketing goals by driving engagement, conversions, and long-term brand trust, and its success can be evaluated through two distinct but complementary lenses: quantitative performance metrics of user engagement and qualitative assessments of user perception (Joan Isibor et al., 2025). Quantitative measures, such as reach, impressions, engagement rates, and conversion outcomes, provide insight into how widely content is distributed and how users interact with it. Saura et al. (2017) posit that such indicators are grounded in measurable data and prove valuable for tracking visibility and campaign performance. However, quantitative measures primarily capture observable behaviour without explaining why it occurred, and remain insufficient for understanding the perceptual and attitudinal processes that drive it (Saura et al., 2017). Qualitative evaluation addresses this limitation by focusing on the cognitive processes underlying user responses. Instruments such as user surveys reveal how users form judgements about content and how those judgements shape attitudes and behavioural intention (Ismagilova et al., 2020). For instance, perceived credibility and relevance are particularly powerful predictors of behavioural intention in digital environments (Ngo et al., 2024). Critically, a piece of content may generate strong engagement or conversions while simultaneously failing to earn genuine trust, highlighting a consequential gap between observable behaviour and meaningful user connection (Yardley et al., 2016). To bridge this gap, this study draws on a set of perceptual constructs grounded in established consumer psychology, each of which is elaborated below.

In marketing contexts, content is inherently persuasive, designed to influence user attitudes and guide decision-making towards specific outcomes such as engagement, enquiry, or purchase (Munteanu et al., 2024). Persuasion is defined as a process aimed at changing a person's attitude or behaviour towards products and services, using written or spoken words and other message elements (Ogah & Abutu, 2022). 'Persuasiveness' is thus a crucial measure of how effectively digital content fulfils its communicative function, and is included as a key construct in the survey instrument. Closely related to persuasion is content quality, defined as material that is valuable, relevant, and engaging for a specific target audience, thereby attracting, retaining, and influencing that audience while building trust and brand relationships (Bubphapant & Brandão, 2024). High-quality website content is strongly correlated with improved usability and conversion rates (Almaghlouth, 2024), making 'content quality' a relevant and theoretically grounded construct for this study. A similar message credibility dimension is noted by Appleman and Sundar (2016) called "accuracy" which overlaps significantly in scope with content quality. Given the similarity in these dimensions, content quality was deemed a better fit for this study as it captures more breadth and avoids confusing other GenAI measures that use the term "accuracy" to discuss GenAI

specific phenomena such as hallucinations. Furthermore, Dwiphala and Astuti (2024) argue that the quality and emotional resonance of digital content directly shape product interest. Content that is low in informative or experiential value leads to lower user interest in the product compared with well-executed content (Dabbous & Barakat, 2020). ‘Product interest’ is defined as the degree of curiosity and attention a consumer directs towards an advertised product and is considered a meaningful predictor of purchase intention (Dwiphala & Astuti, 2024). In a digital marketing context, product interest reflects whether a product is perceived as attractive, relevant, or worth learning more about, often sparked by advertisements or online experiences (Cheah et al., 2024). This makes it particularly relevant as a measure of landing page content, where conversion is the primary objective. However, the effectiveness of digital content cannot be reduced to its functional qualities alone. Beyond persuasiveness, quality and product interest, users also evaluate content through a more fundamental lens: whether it feels authentic.

Authenticity plays a central role in shaping evaluations of digital content, especially in conversion-oriented environments (Arango et al., 2023). Within digital marketing literature, authenticity refers to the extent to which content is perceived as genuine, truthful, and aligned with expected values or intentions (Zniva et al., 2023). Importantly, authenticity is not an inherent property of the content itself but a judgement formed by users based on their experiences, expectations, and perceptions of the source. These authenticity judgements are closely linked to trust, engagement, and brand-related outcomes. When content is perceived as lacking authenticity, users tend to engage with it less and develop more negative attitudes toward the brand (Arango et al., 2023; Kirk & Givi, 2025). Research on GenAI-authored digital and social content specifically shows decreased perceived brand authenticity and more negative attitudinal and behavioural reactions among users (Brüns & Meißner, 2024), thereby reducing engagement and purchase-related behaviours (Curina et al., 2020). In the context of GenAI content, such concerns often stem from the perceived absence of human intent and emotional depth, making ‘authenticity’ a theoretically significant construct for examining how authorship disclosure shapes user evaluations. Additionally, the authenticity of content consumed online also affects user perceptions of the brand. ‘Brand genuineness’ is understood as the extent to which a brand is perceived as honest, legitimate, and consistent in its communication (Wymer & Akbar, 2017), particularly when the disclosure of AI authorship invites greater scrutiny from readers (Toff & Simon, 2025). Unlike authenticity, which is evaluated at the level of individual content, ‘brand genuineness’ reflects a broader, cumulative judgement of the brand’s communicative integrity. It is therefore included as a construct to examine how authorship disclosure affects broader brand-level perceptions beyond the content,

particularly in landing page contexts, where users are unlikely to separate their judgment of the brand from the persuasive intent of the content.

Collectively, the perceptual constructs discussed here (authenticity, persuasiveness, product interest, brand genuineness and content quality) carry measurable consequences for how users perceive and engage with a brand's digital content. As the review will show, GenAI-authored content may perform well in terms of clarity and structure, yet user evaluation remains shaped by subjective expectations and, critically, by the perceived source of the content.

2.3.1 Source and message credibility

The idea that message acceptance is shaped not only by the content itself but by the perceived characteristics of its source is postulated by the Source Credibility Theory (SCT) (Hovland et al., 1953). Within this framework, source credibility is defined along two dimensions: perceived expertise, reflecting the competence and knowledge attributed to the source, and perceived trustworthiness, reflecting its credibility and integrity. Digital marketing literature has highlighted that source credibility can enhance the acceptance and persuasiveness of content, whereas a perceived lack of credibility can undermine even well-structured marketing messages, leading to reduced purchase intention and perceived usefulness among target audiences (Curina et al., 2020; Ismagilova et al., 2020). 'Credibility' can be defined as the extent to which consumers see a digital advertisement or message as truthful, reliable, fair, and believable (Kothari et al., 2025), and is a critical measure of consumer trust and usefulness (Ismagilova et al., 2020). Credibility is therefore included as a construct to examine how authorship shapes participants' evaluations in a landing page context. This makes SCT a particularly relevant theoretical framework, as it explains how the perceived source of content, whether human or GenAI, influences user judgements of trustworthiness and persuasive intent. Trustworthiness and attractiveness, both dimensions of source credibility in celebrity endorser research (Ohanian, 1990), were not measured as separate constructs. Trustworthiness was subsumed within the composite credibility construct, and attractiveness was excluded as more relevant to visual or spokesperson credibility than text-based authorship.

Within the SCT framework, trustworthiness and perceived expertise play an equally decisive role. Wang and Huang (2023) assert that perceived author expertise can significantly shape the success of digital content, in terms of engagement, helpfulness, and purchase behaviour. Lou and Yuan (2019) support this by noting that an author's competence and knowledge on social media shape trust in the author, which then affects brand attitudes and purchase intentions. 'Author knowledge' can then be

defined as the perceived expertise and domain-specific knowledge of the content creator (human or AI), including signs of professional credentials and topic-related experience, which contribute to judgements of source credibility (expertise, trustworthiness) and influence how consumers process, trust, and act on marketing messages (Ismagilova et al., 2020; Jenkins et al., 2020; Thon & Jucks, 2017). The primary mechanisms of SCT, source expertise and trustworthiness, apply regardless of whether the perceived source is a human author, a GenAI tool, or a domain expert, making ‘author knowledge’ a theoretically grounded construct for this study.

The dual-channel structure of source perception underpins a distinction central to the present study: the difference between message credibility and source credibility. Message credibility concerns how believable and trustworthy the information itself is judged to be, while source credibility concerns the origin of that information, the author or organisation, used as a heuristic cue when evaluating content (Ou & Ho, 2024). Research shows that users are more likely to engage with content that is informative, reliable, and well-structured, reinforcing message credibility as a judgement of content quality and trustworthiness (Moran & Muzellec, 2017; Saima & Khan, 2020). This is particularly relevant in high-involvement purchasing contexts, where information quality, relevance, and clarity shape confidence in both the content and the brand behind it (Davidavičius & Limba, 2022; Rashmi & Adhana, 2023). Johnson and Grayson (2005) further refine these dimensions by distinguishing between cognitive trust, grounded in rational assessments of competence and accuracy, and affective trust, reflecting emotional reassurance and relational responses to the communicator. Together, these dimensions suggest that users evaluate digital content on two levels simultaneously, judging both what the content says and who or what is perceived to be saying it.

While the seven constructs identified so far in this review (persuasiveness, content quality, product interest, authenticity, brand genuineness, credibility, and author knowledge) capture how users evaluate digital content and its source, their salience is unlikely to remain constant across contexts. A critical moderating factor is authorship disclosure, specifically, whether users know they are engaging with GenAI-authored content. The following sections examine how disclosure shapes evaluative responses, leading to the study’s final two constructs.

2.4 User evaluations of GenAI-authored and human-authored content

The literature on user evaluations of digital content suggests that when the origin of the content is unknown, users tend to evaluate GenAI-authored material favourably across

a broad range of digital contexts. In social media and advertising settings, GenAI-authored content is routinely perceived as credible, visually appealing, and relevant, performing on par with or better than human-authored equivalents (Hussain, 2024; Park et al., 2024). The advantage becomes more pronounced in persuasion-oriented contexts. GenAI-authored product advertisements and campaign messages have been rated higher in quality than content produced by human experts (Y. Zhang & Gosline, 2023), and GenAI-crafted marketing emails have outperformed human-written ones across dimensions of persuasiveness, clarity, and professionalism (Acatrinei & Miu, 2025). When classical persuasion principles are applied, the gap widens further, with GenAI-authored advertisements significantly outperforming human-authored alternatives in user preference (Meguellati et al., 2025). Personalisation amplifies this effect, as Matz et al. (2024) show that tailored GenAI-authored messages demonstrate greater influence on consumer decision-making than non-personalised content. This positive reception extends beyond casual or incidental content consumption. In more deliberate, comparative evaluations, GenAI content continues to perform strongly. Participants have rated GenAI-authored short stories as more enjoyable and better written than human alternatives (Doshi & Hauser, 2024), and blind-marked assessments of academic writing show that even trained evaluators struggle to distinguish GPT-4 outputs from human work (Yeadon et al., 2024). This pattern of favouring GenAI content is observed across social content, creative writing, academic essays, and professional marketing copy, suggesting that it is not limited to any single format or audience. Instead, it points to a more general feature of how users respond to GenAI textual content when source awareness is absent. They respond, largely, as they would to human-authored work.

What is theoretically significant here is not simply that GenAI content scores well, but why. Research on digital media processing suggests that users rarely engage in deep, effortful evaluation of content. Instead, they rely on cognitive heuristics, or surface-level shortcuts, to form rapid credibility judgements (Metzger et al., 2010; Shyam Sundar, 2008). Attributes such as clarity, structural consistency, and tonal coherence function as proxies for quality, triggering positive evaluations without requiring scrutiny of the content's actual origin or depth. GenAI content, typically clean in structure and consistent in presentation, is well-positioned to satisfy these shortcuts precisely, especially in structured formats (Ouyang et al., 2022). Specifically in the case of structured digital marketing copy, GenAI can replicate the conventions, style, and persuasive logic of professional copywriting (Hartmann et al., 2025). Therefore, when authorship is undisclosed, users have no competing cue to disrupt this process, and so perceived content quality becomes the dominant, and often sole, basis for judgement.

2.5 Authorship disclosure and source-based bias

Authorship disclosure fundamentally alters how users receive and evaluate GenAI content, often reversing the favourable judgements observed in the absence of authorship disclosure. What changes, however, is not the content itself but the evaluative lens through which users read it. Across social, emotional, commercial, and news contexts, disclosure of AI authorship consistently triggers negative evaluative shifts. This pattern is captured by Longoni et al. (2019) and their concept of the ‘disclosure penalty’, defined as the consistent finding that identical content is rated less favourably once AI authorship is known. In creative and emotional writing settings, revealed AI authorship lowers perceived trustworthiness, competence, and likability (Nakano et al., 2025). In marketing and consumer contexts, identical texts (such as corporate communications and blog posts) are judged as less credible and of lower quality when labelled as GenAI-authored rather than human-written (Haupt et al., 2025). The penalty extends to news contexts, where AI-disclosure leads to lower ratings for trustworthiness, writing quality, and likelihood of sharing (Kreps et al., 2022), and to credibility and authenticity judgements more broadly, with users less willing to recommend or revisit AI-attributed content despite no change in the material itself (Lermann Henestrosa & Kimmerle, 2024; Rae, 2024). Since the content is identical across conditions, the penalty cannot be attributed to content quality. In fact, it is a source-based bias (Y. Zhang & Gosline, 2023) that changes how users respond to content on digital platforms. As the evidence highlights, this source-based bias operates independently of content quality, suggesting that user responses to GenAI-authored content are driven not by the content itself but by what is known about its origin. This distinction sits at the heart of this study's experimental design.

Notably, however, the body of evidence investigating user evaluations of GenAI, both with and without authorship disclosure, is concentrated in social, journalistic, and non-transactional or informal commercial contexts. Structured digital marketing environments, such as website landing pages, where content is explicitly persuasive, conversion-oriented, and evaluated by users with transactional intent (Davidavičius & Limba, 2022), remain largely unexamined. Whether the disclosure penalty operates with the same force in such contexts, where credibility and commercial purpose are already in tension, is a question the existing literature does not answer. For now, the review continues to shed light on the source-based bias in greater depth.

2.5.1 Algorithm aversion

Users' diminished responses to disclosed GenAI authorship content noted in the preceding section are consistent with the broader literature on algorithm aversion,

which describes users' reluctance to rely on algorithmic output despite demonstrated performance (Dietvorst et al., 2015). In their seminal study, Dietvorst et al. (2015) established that users become less willing to trust algorithms following even minor errors, and that this reluctance persists even when algorithms provably outperform human alternatives. More recent research confirms that this scepticism extends specifically to GenAI-authored content, driven by concerns about reliability, intent, and ethical integrity (Morales, 2025). This heightened critical response manifests across multiple dimensions. Users assign a lower value to GenAI-authored news articles and persuasive writing, even when they are indistinguishable from human writing, primarily due to distrust in GenAI accuracy and the belief that LLMs hallucinate and fabricate information (Darda et al., 2023; Zhu et al., 2025). In social media contexts, disclosure of AI involvement elicits negative attitudinal and behavioural responses due to reduced perceived authenticity, with full AI replacement viewed more negatively than AI used as a human aid (Brüns & Meißner, 2024; Radivojevic et al., 2024). Psychological explanations suggest the underlying reason. Users are more forgiving of human error because they attribute intention, emotion, and accountability to human creators, qualities they do not extend to AI systems (Renier et al., 2021). Negative evaluations of GenAI-authored content, then, are often less a response to the content itself than a reflection of deeper, pre-existing beliefs about the nature of machine-generated output.

Once AI authorship is disclosed, users do not simply adjust their evaluations. In fact, they reassess content through a more critical and sceptical lens. Toff and Simon (2025) find that AI-labelled content is perceived as less trustworthy even when its factual accuracy is not in question, suggesting that disclosure activates scrutiny that content quality alone cannot neutralise. This evaluative reversal observed in authorship disclosure is partly explained by the way disclosure functions as a cognitive cue. Metzger et al. (2010) identify source-based heuristics as among the most powerful shortcuts available to users in digital environments. Once authorship is known, identity cues rapidly reshape how content is received, often overriding prior judgements formed based on content quality alone. Across these findings, the pattern is consistent. Disclosure does not alter the content being evaluated, but fundamentally changes the framework through which it is judged.

A different and contrasting perspective emerges in fast-paced digital media environments, where user responses to GenAI-authored content appear to be evolving. Kubovics (2025b) shows that users are more likely to engage with GenAI-authored marketing content when it is perceived as useful, reliable, and transparently labelled, suggesting that value and clarity can attenuate the negative effects of AI disclosure. This is supported by Sew et al. (2025), who find that, while scepticism towards AI

influencers persists among some users, others respond positively to GenAI content despite full awareness of its origin. Zhe and Srijinda (2024) extend this by showing that personalised GenAI content can enhance user satisfaction by integrating cognitive value with affective appeal. This combination reduces the salience of authorship cues when users prioritise utility over origin. These findings suggest that the negative effects of AI authorship disclosure documented in earlier research are not immutable. They are sensitive to whether content delivers perceivable value. This sits in tension with the algorithm aversion literature and implies that aversion is not a stable user trait but a contextually activated response. Logg et al. (2019) demonstrate that in certain domains, users exhibit algorithm appreciation. In other words, they over-rely on algorithmic advice precisely because it is perceived as objective and consistent. Taken together, aversion and appreciation represent two poles of a continuum, with perceived value, task type, and domain stakes determining where on that continuum a given user is likely to land.

An important methodological caveat runs through the studies presented in this section. Few of these studies measure preference directly, and instead, infer it from evaluative proxies, including ratings of trustworthiness and writing quality (Kreps et al., 2022), perceived credibility and content quality (Haupt et al., 2025; Nakano et al., 2025), and willingness to recommend or revisit content (Lermann Henestrosa & Kimmerle, 2024; Rae, 2024). While these dimensions are meaningfully related to preference, they are not equivalent to it. A user who rates GenAI content lower on trustworthiness may still choose to engage with it, and a user may express a preference for human authorship while behaving differently in practice (Choudhury & Shamszare, 2023). The distinction matters because evaluation and preference can conflict, particularly when the evaluative setting is artificial. Participants in controlled studies are prompted to reflect on authorship in ways that do not reflect real-life digital environments (Toff & Simon, 2025). What these studies collectively demonstrate, then, is that disclosed AI authorship shapes how users assess content, not necessarily how they choose it. The present study addresses this gap by measuring both evaluative judgements and content preferences, providing a more grounded understanding of how authorship disclosure influences user behaviour.

2.5.2 Human favouritism

Not all research frames authorship disclosure effects as purely aversion-based. An alternative explanation positions the evaluative shift not as a penalty on AI, but as a premium on human authorship. Casal and Kessler (2023) and Parshakov et al. (2025) report psychological biases that systematically favour human authorship over GenAI

authorship when the source is disclosed, even when the source cannot be reliably identified. In the specific case of persuasive marketing content, a significant study by Y. Zhang and Gosline (2023) shows that the quality advantage GenAI-authored text holds under non-disclosure diminishes upon disclosure. They assert that this is not because perceptions of AI worsen significantly, but because evaluations of human-authored content rise. In other words, knowing the source of the content shifts the evaluation from a quality-based to a context-driven heuristic that places an intrinsic value on human origin (Y. Zhang and Gosline, 2023). This framing reorients the phenomenon. Now, the issue is not simply that AI content is distrusted, but that human authorship carries an evaluative premium that disclosure makes visible.

Whether interpreted as algorithm aversion or human favouritism, the evidence points to the same structural conclusion. Authorship disclosure serves as a powerful cognitive cue that reshapes evaluation independently of content merit. User judgements are shaped not only by what content is, but by who or what produced it. Again, research on content in structured digital marketing settings, such as website landing pages, is underrepresented in prominent research, which largely focuses on descriptive text, such as news articles, creative writing, and informal marketing content. Therefore, the research problem this thesis addresses is to examine the distinction between content-based and source-based evaluation in the context of structured, informative, and conversion-oriented content on a website landing page.

2.6 Contextual moderators of GenAI content evaluation

The strongest and most consistent finding in the literature is that users are more receptive to GenAI-authored content in functional, task-oriented contexts where accuracy, efficiency, and clarity are the primary evaluative criteria. Okoso et al. (2025) assert that users engage comfortably with GenAI-authored content in domains such as product or hotel recommendations, provided the information is accurate and clearly structured. Similarly, Kashive et al. (2020) find that user satisfaction in AI-assisted contexts is driven by perceived usefulness and task simplification. These findings align with the theorisation by Castelo et al. (2019) that trust in algorithmic outputs is task-contingent. They assert that users are more willing to rely on AI for objective tasks with verifiable outcomes, and considerably more reluctant when tasks require subjective judgement. However, this is not true in emotionally sensitive and high-stakes contexts. In their study, Osborne and Bailey (2025) found that users consistently prefer human-generated responses in domains involving personal advice, relationships, and well-being, even when GenAI alternatives are informationally equivalent or superior. This preference is not incidental and implies that user acceptance of GenAI-authored

content declines markedly regardless of content quality. Kirk and Givi (2025) further reveal that consumer loyalty and word-of-mouth intentions decline when emotional marketing communications are perceived as GenAI-authored, with such content perceived as morally inauthentic. This occurs not because GenAI lacks information, but because the perceived absence of a human author strips it of emotional legitimacy.

The finding that acceptance of GenAI content is heavily shaped by expectations of human presence and accountability leads to another construct this study employs to study user perceptions: 'human-likeness' (the eighth evaluative construct of the study). Stein et al. (2024) demonstrate that perceived mental human-likeness functions as a critical mediator of user engagement, with lower human-likeness suppressing interaction even for informative content. Human-likeness is described as the degree to which content or a content source is seen to exhibit recognisably human cognitive and emotional qualities (Stein et al., 2024). Transposed to written digital marketing content, this suggests that users are not simply assessing what content says, but whether the writing itself carries traces of human intelligence, intention, and emotional investment. These are the precise qualities whose absence Kirk and Givi (2025) identify as the drivers behind users' characterisation of GenAI-authored content as morally inauthentic. This judgement is less about stylistic imperfection than about the perceived absence of genuine human intent. It is precisely this sense of intentional absence that this study captures through its ninth and final construct: 'writing connection'. Writing connection is described as the degree to which users perceive written digital marketing content as reflecting human cognitive and emotional qualities, including intention, relatability and emotional expressiveness (Stein et al., 2024). The construct measures the sense that a human mind was present in the composition of the content, and that the content was written for them rather than generated at them. It can be applied specifically to textual landing page content, a context in which visual and interactional human-likeness cues are absent and textual qualities must therefore carry the full weight of perceived human presence (Dvir & Gafni, 2018). Together, human-likeness and writing connection operationalise a dimension of user perception that the broader literature consistently identifies but rarely measures directly: the extent to which written content feels as though it originates from and was intended for a human being.

2.7 User attitudes and prior experience with GenAI

Contextual factors alone do not determine how users respond to GenAI-authored content. Research shows that the same disclosure cue and interface setting can elicit markedly different reactions depending on the prior attitudes and AI experience users bring to the encounter. Darda et al. (2023) assert that users who hold more positive

views of the utility of AI tend to evaluate GenAI content more favourably than human-authored content, exhibit lower levels of algorithm aversion, and respond less strongly to disclosure cues. Moreover, familiarity with AI tools is generally associated with greater openness. Users who have interacted repeatedly with AI tools tend to perceive GenAI-authored content as more credible and respond less negatively to disclosure (Kubovics, 2025a; Li et al., 2025). Structured exposure, through training, repeated use, or professional engagement with AI tools, can therefore reinforce positive evaluations by establishing AI as a competent and efficient agent within specific task contexts (Pellas, 2023). Conversely, Buchanan and Hickman (2024) show that users who are more sceptical of organisations using GenAI demonstrate heightened sensitivity to authorship information, applying stronger source-based scrutiny to content once its AI origin is known. Other evidence shows that increased exposure can heighten critical awareness of AI's limitations, including its tendency towards generic outputs, factual errors, and a lack of contextual depth, leading to more demanding rather than more lenient evaluation (Buchanan & Hickman, 2024; Zhou & Lu, 2025). In this sense, familiarity can both increase acceptance and intensify scrutiny, depending on the context and quality of the content.

Collectively, the findings discussed so far on users' attitudes and experiences with AI shape how users approach GenAI-authored content, but do not fully determine their evaluations or choices. Users may express scepticism towards AI while still engaging positively with its outputs, particularly when the content is perceived as useful, clear, or relevant. This reflects a broader tension between users' stated beliefs and observed behaviour, where practical utility tends to override attitudinal resistance (Altay & Gilardi, 2024). Importantly, these attitudes are not static. Mass exposure to GenAI tools and outputs is gradually shifting baseline familiarity, softening initial resistance and normalising GenAI content as part of everyday digital experience (Jia et al., 2025). As familiarity increases, the evaluative gap between GenAI and human-authored content may narrow further, not because the content has changed, but because users have. While hybrid or collaborative AI approaches such as human-in-the-loop (Brüns & Meißner, 2024; Haupt et al., 2025; Kirk & Givi, 2025) can help mitigate some concerns by emphasising human oversight, their effectiveness remains context-sensitive and cannot be assumed to operate uniformly across user groups or digital environments.

2.8 Chapter conclusion

This chapter has established the theoretical and empirical foundations for understanding user perceptions of GenAI-authored digital marketing content. The review demonstrates that when the origin of the content is unknown, GenAI-authored

content is generally evaluated favourably across a range of digital contexts. It performs on par with or exceeds human-authored alternatives in terms of quality, clarity, and persuasiveness. This favourable response from users is attributed to heuristic processing, in which surface-level cues such as structural and informational consistency serve as proxies for quality without authorship context. Once authorship is revealed, those evaluations shift unfavourably for GenAI content, not due to changes in content quality but source-based biases. The literature refers to them as ‘algorithm aversion’ and ‘human favouritism’, while noting that their effects are not uniform. The ‘disclosure penalty’ is also moderated by contextual factors such as task type and domain stakes, and by individual-level variables including prior AI experience and pre-existing attitudes towards GenAI. The literature further suggests that as familiarity with GenAI tools increases, the evaluative gap between GenAI-authored and human-written content may narrow due to users' evolving attitudes towards AI.

Across these bodies of work, however, two limitations persist. First, the research base remains concentrated in social, journalistic, and informal commercial contexts. Structured, conversion-oriented environments such as website landing pages, where content is explicitly structured and persuasive, remain underexplored. Whether the evaluative dynamics documented elsewhere operate with the same force in these settings cannot be assumed. Second, the literature has largely treated evaluation and preference as interchangeable, inferring user choice from evaluative measures alone. This neglects a practically significant distinction that the present study directly addresses. To examine user perceptions across both dimensions, this study employs nine perceptual constructs identified through the review as the most salient and theoretically grounded dimensions of evaluation: authenticity, persuasiveness, content quality, product interest, brand genuineness, credibility, author knowledge, human-likeness, and writing connection. These constructs collectively capture both the rational and affective dimensions of user response and provide a framework capable of distinguishing between how users judge content and how they feel about its source. Against this framework, the research questions of this study are examined: whether users prefer GenAI-authored website landing page content or human-authored alternatives, and whether the perceived source of that content shapes how they evaluate and respond to it.

Chapter 3 Methodology

3.1 Chapter overview

Examining user perceptions of GenAI-authored content presents a methodological challenge: the phenomenon is perceptual, context-dependent, and still evolving, making it poorly suited to a single-method approach. The design choices described in this chapter were made in direct response to that challenge. This chapter outlines the research design and methodological approach, describing data collection, analysis, and validation procedures. The methodology is informed by a pragmatist paradigm, which supports the use of mixed methods to address applied research questions. This approach allows the study to identify patterns in quantitative survey data while also capturing users' reasoning through open-ended responses, enabling a context-sensitive examination of how users evaluate digital content written by both GenAI and humans. The study draws on Source Credibility Theory (SCT), which proposes that perceptions of a message source shape judgements of credibility and persuasiveness. Data were collected via an online survey using a 2×2 mixed-design experiment that compared GenAI-authored and human-authored content under disclosed and undisclosed authorship conditions. After qualitative and quantitative analyses were completed separately, the findings were integrated to produce a more coherent picture of user evaluations. Overall, the study design supports a comprehensive examination of user preferences for digital content and the impact of authorship disclosure in a digital marketing context.

3.2 Philosophical assumptions

In research, a paradigm is the overarching philosophical framework or worldview that shapes how a researcher understands reality, what counts as knowledge, and how inquiry should be carried out (Kivunja & Kuyini, 2017). This study is grounded in a pragmatist paradigm, which carries distinct philosophical positions across three interconnected dimensions: ontology, epistemology, and methodology.

Ontologically, pragmatism holds that reality is neither fixed and singular nor entirely subjective. In fact, it is plural, context-dependent, and shaped by practical human experience (Morgan, 2014). This assumption is particularly fitting for a study examining user perceptions of GenAI-authored and human-written content, since those perceptions are not absolute truths. In reality, and as this study will reveal, they are shaped by individual experience, context, and the conditions under which content is encountered. Epistemologically, pragmatism rejects the idea that knowledge must be

accessed exclusively through either objective measurement or subjective interpretation (Johnson & Onwuegbuzie, 2004). Instead, it accepts that while objective knowledge can be sought through measurement, perceptions are also influenced by subjective interpretations (Creswell & Clark, 2017). This position supports the integration of numerical data with participants' narrative explanations through the mixed methods approach, allowing the study to capture both the “what” and the “why” of user evaluations (Feilzer, 2010). Methodologically, this implies that the research question drives the choice of methods rather than adherence to a single philosophical tradition, resulting in a study design that measured context-dependent perceptions, both with and without authorship disclosure. The primary source of knowledge is an electronic survey that yielded structured numerical data, providing a systematic view of how participants evaluated landing page content across constructs such as authenticity and credibility. Open-ended qualitative responses added interpretive depth by illuminating the reasoning behind participants' ratings. Together, these ontological, epistemological, and methodological commitments ensure that the study's findings remain grounded, contextually relevant, and oriented towards practical understanding. Grounded in this pragmatist orientation, the study further draws on an established theoretical framework to explain the impact of source awareness on user perceptions of GenAI-authored content.

3.2.1 Source Credibility Theory

Section 2.3.1 established Source Credibility Theory (SCT) (Hovland et al., 1953) as the theoretical basis for two constructs: credibility and author knowledge, both grounded in SCT's dimensions of perceived expertise and trustworthiness. This section considers what SCT predicts for the present study. Recent research extends SCT into digital contexts, showing that perceived authorship, human or AI, shapes how users judge website and social media content (Parshakov et al., 2025; Zhu et al., 2025). SCT concerns perceptions of a source, not its actual identity. If this holds, a user's belief about who authored a message should activate similar evaluative processes, whether the source is a person or a system. This assumption is not taken for granted here. SCT was developed around human communicators, whose expertise and trustworthiness are anchored in verifiable credentials, accountability, and reputational stakes. Hovland et al.'s (1953) original work concerned speakers whose credibility could be gained or lost through consequence. A GenAI system has no comparable history or stake in its output. Whether SCT's two dimensions transfer equally to such a source is unclear. Author expertise may be judged on demonstrated competence regardless of who or what produced it, whereas trustworthiness has traditionally depended on accountability that a GenAI source cannot possess. This is an open empirical question, not a settled

theoretical premise, and it is what the present study's disclosure manipulation is designed to test.

This framing also clarifies why a mixed-methods design is required. Johnson and Grayson's (2005) distinction between cognitive trust, based on rational assessment of competence, and affective trust, based on emotional and relational response to the communicator, maps onto this study's dual evaluative framework. Cognitive trust corresponds to functional competence, while affective trust corresponds to relational authenticity. A quantitative measure alone could capture whether functional evaluations shift once authorship is disclosed. It cannot explain why relational trust responds differently. Reflexive thematic analysis is therefore used to capture that reasoning, which SCT alone does not resolve when the source is non-human.

3.3 Mixed methods approach

This study employed a mixed methods design, collecting quantitative and qualitative data simultaneously through a single online survey (Creswell & Clark, 2017). Mixed methods research is one of the most common procedures for understanding complex research problems (Shorten & Smith, 2017), providing both the breadth of quantitative evidence and the depth of qualitative insight. Yardley et al. (2016) assert that users' connection with digital technologies and their response to them should be measured over time using constructs that include subjective experience (how connected they feel, what it means to them) and objective behaviour (how they use, act, or react). In this study, quantitative measures captured patterns in content evaluations of perceived authenticity, credibility, persuasiveness, author knowledge, writing connection, human-likeness, product interest, brand genuineness, and overall quality. In addition, qualitative data were used to gain insights into participants' interpretations of GenAI-authored and human-written marketing content in the context of a landing page. Integrating both data types enabled the study to move beyond identifying differences to explaining them, particularly regarding authorship disclosure effects, which was crucial to answering this study's research questions.

Another advantage of using a mixed methods design is that it is well-suited to complex research questions in the social sciences (Dawadi et al., 2021). It allows confirmatory and exploratory questions to be addressed simultaneously and helps explain contradictory findings. As noted in Sections 2.4 and 2.5, prior research on GenAI content has produced mixed results with users favouring GenAI content when it is effective and useful (Parshakov et al., 2025), but preferring human output when authorship is disclosed (Y. Zhang & Gosline, 2023). Furthermore, Section 2.6 highlights how contextual factors like task type and digital environment influence user

evaluations, and how users' attitudes towards AIGC are evolving as it becomes more mainstream (Sew et al., 2025). The mixed methods approach adopted in this study, therefore, enables both the validation of prior findings and the exploration of user perceptions in structured digital environments, such as website landing pages, that are increasingly adopting GenAI tools for content creation.

3.4 Research design

The study used a 2×2 mixed-design experimental survey. The design included two independent variables: content authorship and authorship disclosure. Content authorship had two levels: GenAI-authored content, produced using ChatGPT, and human-authored content, written by a human digital content specialist. The GenAI landing page was created using ChatGPT-4 with the following prompt:

“Create a website landing page for a real estate agent targeting first home buyers.”

GPT-4 is an LLM developed by OpenAI to generate text in response to user prompts and can assist with tasks such as drafting, editing, and refining written content (*GPT-4 | OpenAI*, n.d.). In this study, ChatGPT was used to produce a marketing-style landing page based on the same product context and core information used for the human-authored version. The pages were designed for a real estate service targeting first-home buyers, a domain known to heighten sensitivity to credibility and authenticity cues (Longoni et al., 2019). The generated text was then used as the AI-authored stimulus for comparison with the human-written landing page.

The human-authored landing page was written by the researcher, who has over 7 years of professional experience as a Digital Content Specialist. This experience involved producing and refining digital content for online platforms, with a focus on audience engagement, structure, and brand-aligned communication. This background made the researcher suitable for creating a realistic example of professionally written marketing content for comparison with GenAI content. To support consistency and reduce bias, both the GenAI-authored and human-written landing pages were based on the same product context and core information. Authorship disclosure also had two levels: disclosed and undisclosed. The study used two types of comparison to address the research questions. First, it used a within-subjects design (Ledolter et al., 2020), in which each participant evaluated both versions of the same website landing page: one GenAI-authored and one human-authored. This allowed the study to directly compare responses to the two content types and reveal which they rate higher and prefer. This element of the design addressed the first research question:

Research question 1: Do users prefer GenAI-authored or human-authored website landing page content when authorship is not disclosed?

Second, the study used a between-subjects design (Hazra & Gogtay, 2016), in which responses were compared between two separate participant groups. Participants were randomly assigned to either a control group or a treatment group. The control group viewed both landing pages without being told whether the content was produced by ChatGPT or by a human writer. The treatment group viewed the same landing pages with clear authorship labels identifying the source of each page. This allowed the study to examine whether, and how, making authorship visible changed user preferences, addressing the second research question:

Research question 2: How does authorship disclosure influence users' preferences for GenAI-authored and human-authored website landing page content?

At the page level, the design therefore produced four conditions: GenAI content with disclosure, GenAI content without disclosure, human-authored content with disclosure, and human-authored content without disclosure. This structure made it possible to examine both baseline content preferences and the effect of authorship disclosure on those preferences. Aside from the presence or absence of authorship labels, all survey materials, instructions, landing page layouts, and evaluation measures were identical across both groups. This ensured that any observed differences between the two groups could be attributed to authorship disclosure rather than to differences in design, wording, or survey procedure.

Participants evaluated the landing pages using the same set of measures, including Likert-scale items assessing credibility, trust, authenticity, perceived quality, persuasiveness, author knowledge, human-likeness, brand genuineness, and writing connection. They also completed preference-based comparison items and open-ended questions explaining the reasons for their choices. These qualitative responses provided additional insight into how participants interpreted and justified their evaluations. In addition, randomisation was implemented via the online survey platform, which automatically assigned participants to either the control group (without authorship disclosure) or the treatment group (with authorship disclosure) at a 1:1 allocation ratio. The group assignment was concealed from both participants and the researcher during data collection. This reduced allocation bias and helped ensure a balanced distribution of participants across the two experimental groups (Odgaard-Jensen et al., 2011).

Having established the experimental conditions and evaluation measures, the following section describes how participants were recruited and the ethical procedures governing their involvement.

3.5 Participants

Participants took part through CloudResearch Connect, an online crowdsourcing platform used in social and behavioural science research to identify and recruit participants (CloudResearch, n.d.). Participants were eligible to take part if they were adults aged 18 years or older and had access to the internet and a computer or mobile device. Individuals aged 18 and above are generally recognised as legal adults and able to provide independent informed consent (Mathews, 2022). Data were collected from participants worldwide, with no restrictions based on location or digital literacy. No additional exclusion criteria were applied. Participation was voluntary, and individuals were free to choose whether to take part and could withdraw at any stage without penalty.

Before publishing the survey and carrying out the study, ethical approval was obtained from AUTECH (Appendix A: AUTECH Approval Letter). All project details were displayed on their CloudResearch Connect dashboards, where they could access the Participant Information Sheet (Appendix B: Participant Information Sheet) and review the study information at their own pace. No deception was used as participants were informed in the study information sheet that the research involved comparison groups and that the author's information might or might not be shown. This ensured transparency without undermining the integrity of the experimental design. Moreover, consent to complete the survey was obtained electronically via Qualtrics (Appendix D: Survey Consent Preview), in which participants confirmed they had read the information sheet and voluntarily agreed to participate. Participants who consented were then automatically directed to the survey.

A target sample size of 100 completed surveys was selected to align with the study's exploratory aims and analytical goals. Previous research suggests that surveys focused on central tendencies and common response patterns can achieve stable results with samples of approximately 80 participants (Manyara et al., 2024). On this basis, a sample of 100 participants was selected, with the additional margin included to account for attrition and potential exclusion of low-quality responses during data cleaning. It also meant providing reliable descriptive insight while remaining feasible within the study's time and resource constraints. This approach also reflects pragmatic survey research guidance, which cautions against collecting unnecessarily large

samples when additional responses are unlikely to meaningfully improve insight (Lakens, 2021).

The study setup on CloudResearch Connect included the study description, eligibility requirements, and payment information. Fair payment was ensured through platform-standard compensation. Participants were paid USD \$3.50 for a survey expected to take approximately 20 minutes, consistent with Connect's recommended hourly rate of around \$10.50. Payment details were visible to participants before enrolment, supporting voluntary and informed participation without undue influence. Moreover, all participant interaction occurred through the CloudResearch platform rather than through direct personal contact. Individuals registered on Connect were shown the study on their dashboard if they met the eligibility criteria set by the researcher. The use of CloudResearch Connect meant the sample consisted of self-selected online users familiar with completing digital tasks. As a result, the findings are most directly generalisable to digitally literate adult internet users rather than to the wider general population. This sampling approach aligns with the study's aims, as the research focuses on how online users evaluate web-based marketing content.

CloudResearch Connect managed participant accounts and contact information throughout the recruitment and data collection process. This personal information was not shared with the researcher to maintain the participants' anonymity. The researcher monitored participation through the Connect dashboard and tracked completions in real time. Completed responses were submitted through Connect and recorded through Qualtrics. The survey remained open for two weeks, giving eligible participants sufficient time to consider the invitation. No follow-up reminders were issued.

Privacy and confidentiality were strictly protected. CloudResearch Connect assigned anonymised participant IDs and did not share any identifying information with the researcher. No names, email addresses, or contact details were collected at any point, and all findings were reported in aggregate form. All researcher-participant communication occurred within the Connect platform, ensuring that contact details were never disclosed. Consent records existed only as timestamped confirmations within Qualtrics, accessible solely to the primary researcher. Furthermore, the study posed minimal risk. The primary burden was the short time commitment required to complete the survey. The possibility that participants might accidentally include identifying details in free-text responses was mitigated by anonymity through Connect and clear instructions advising against providing personal information. No sensitive topics were included, no vulnerable groups were targeted, and no conflicts of interest

were present. All materials used in the study were created specifically for the project, and any participant quotations reported in the findings were fully anonymised.

The following section describes the online survey instrument used by these participants to evaluate the landing page content.

3.6 Survey

Data were collected using an online survey on Qualtrics (Appendix E: Survey). Online surveys are widely used in digital communication research due to their ability to efficiently reach large and geographically diverse participant groups (Menon & Muraleedharan, 2020). They are cost-effective, quick to administer, and well-suited to examining digital behaviours and attitudes in contemporary online environments. In this study, the fully digital format of the Qualtrics survey reduced logistical demands and allowed respondents to complete it in familiar settings, making it particularly appropriate for studying online behaviour (Evans & Mathur, 2018). Moreover, the survey elements, such as the landing page that participants evaluated, were designed in accordance with web usability standards and recommended reading levels (Fogg et al., 2001), ensuring accessibility without the need for additional instructions or training.

Online survey platforms also support multimedia content, adaptive question formats, and compatibility across devices (Singh & Sagar, 2021). This made it easier for participants to complete the survey on Qualtrics, view the landing page images within it, and answer different question formats (Likert scale, open-ended). Automated data collection further reduced the risk of entry errors and enabled the use of a structured experimental design. Collectively, these features made online surveys a robust and flexible method for the present study.

3.6.1 Survey design

The survey questionnaire was designed to ascertain the participants' preferences for GenAI-authored or human-written landing page content. The evaluation questions utilised a 5-point Likert scale asking participants to rate how strongly they agreed with each statement (Joshi et al., 2015). Likert-type questions are widely used to measure perceptions in marketing research, largely because they convert subjective views into analysable numeric data and are easy for respondents to understand (Sullivan & Anthony R. Artino, 2013).

The survey consisted of 3 sections. Section 1 was titled '*Webpage content review*' and collected information about how participants responded to the website pages on their own. Participants were asked to rate website landing page content written by ChatGPT-

4 and a human digital content specialist (See Section 3.4 for authorship details). After viewing the assigned landing page, participants evaluated the content across nine User Perception Measures: authenticity, credibility, persuasiveness, author knowledge, writing connection, human-likeness, perceived brand genuineness, interest in product, and overall content quality. These were based on the literature on users' evaluations of digital content across varying contexts (see Sections 2.2, 2.3.1, and 2.4).

Section 2 of the survey was titled '*Comparison between the webpages*'. It included open-ended survey questions that asked participants to describe whether and why they preferred one page over the other, and to consider it more credible, human-sounding, or higher quality than the other. It also asked them which page they would act on and why. This qualitative data helped augment their numeric responses from Section 1. Both sections 1 and 2 were completed under the same authorship condition for each participant. The control group evaluated and compared the pages without knowing which were GenAI-authored and which were human-authored, providing data that directly addresses Research Question 1. The treatment group completed the same sections with authorship clearly labelled throughout, providing data that addresses Research Question 2. Together, the two sections and two groups allow the study to examine both baseline content preferences and how those preferences shift when authorship is known.

In addition, Section 3 of the survey was titled '*Your experience with AI-generated content*', and focused on participants' experience with tools like ChatGPT, and their attitude towards GenAI-authored content. It focused on their openness to brands using such content and the importance of transparency about its origins (See Section 2.7.2.7), to further explain the reason for their preferences. To minimise potential order effects, the presentation order of the two landing pages (GenAI-authored and human-authored) was counterbalanced. Qualtrics randomised which version each participant saw first, ensuring that evaluations were not unintentionally shaped by sequence.

3.7 Data cleaning, storage and processing

Before analysing the data, an initial check was conducted to verify the dataset's quality and identify any errors or structural issues (Jebb et al., 2017). Checks included identifying and excluding three incomplete responses, extremely short completion times (under 5 minutes), and straight-lined answers to ensure reliability of the dataset (Krosnick, 1991). All analyses were then conducted using a cleaned dataset. The survey data were then exported from Qualtrics into Microsoft Excel and screened for quality and completeness. Preview responses generated during testing were identified using the Qualtrics "Response Type" indicator and removed, along with the first-row

question text labels and system-generated metadata (e.g., timestamps, IP addresses, and location fields) irrelevant to the research questions. Presentation order (human-written webpage first versus GenAI-authored webpage first) and condition assignment were verified using Qualtrics block randomisation variables. Survey items were organised into conceptual groupings corresponding to evaluations of human-written and GenAI-authored webpages. Likert-scale responses were converted to numeric values (1 = strongly disagree to 5 = strongly agree) to facilitate descriptive analysis. The dataset was stored securely on the primary researcher's AUT OneDrive account, with access restricted to the primary researcher and supervisor. As the survey was fully anonymous and collected no personal or identifying information, no sensitive data were held at any point. The Qualtrics dataset was deleted shortly after the survey closed, and the OneDrive copy will be permanently deleted after six years in accordance with AUT's data-retention requirements.

3.7.1 Data validity

Several strategies were employed to enhance the validity and reliability of this research. Construct validity was supported by grounding the nine perception measures in established literature on credibility, persuasion, and human-AI interaction (Flanagin & Metzger, 2007). Internal validity was strengthened through random assignment and consistent presentation of the landing pages (Shadish, 2002). Reliability was reinforced by standardising instructions and ensuring consistent wording across items (Pallant, 2020). External validity, or generalisability, is acknowledged as limited given the online convenience sample (Creswell & Clark, 2017). Nonetheless, the study offers valuable insights into broader patterns of user preferences for GenAI content in a structured digital marketing context. Potential limitations, such as sample bias or self-reporting effects, are discussed further in the concluding chapter (See Section 6.5). With these validity measures in place, the following section describes the quantitative analytical procedures applied to the survey data.

3.8 Quantitative data analysis

The main data source for the survey was quantitative. Therefore, descriptive statistics were generated for all variables to provide a clear summary of the survey dataset and interpret results (Downie & Starry, 2019). Central tendency, including the mean, median, and mode, was calculated, while the variability of the data was measured using the range, standard deviation, and interquartile range (Downie & Starry, 2019). These descriptive summaries provided a clear, accessible overview of participants' responses before any further comparisons were made. To support the interpretation of numerical data, results were visualised using bar charts, boxplots, and comparative

graphs (Vetter, 2017). These helped illustrate differences between groups (disclosure versus no-disclosure conditions) and within subjects (GenAI-authored versus human-written pages), highlighting patterns in user perceptions and preferences. Simple tables and visualisations were used to summarise central tendencies, variability, and overall data patterns.

3.9 Qualitative data analysis

Qualitative data from participants' open-ended responses were analysed using the six-step reflexive thematic analysis (TA) framework developed by Braun & Clarke (2019). Rather than applying a pre-structured coding scheme, reflexive TA allowed themes to emerge inductively from participants' own language and reasoning about GenAI-authored and human-written content. This approach is better suited to capturing the nuanced and often contradictory ways in which users articulate their perceptions of authenticity, trust, and credibility (Byrne, 2022). This was particularly important given that user attitudes towards GenAI-authored content are still evolving and are unlikely to conform neatly to predetermined categories (Jia et al., 2025) (See Section 2.7). Since participants were responding to content under different disclosure conditions, their explanations required interpretive judgement rather than mechanical coding. Reflexive TA accommodates this by treating researcher subjectivity as an analytic resource rather than a source of bias (Braun & Clarke, 2019)

The following Reflexive TA steps, as recommended by Braun and Clarke (n.d.), were followed:

1. Familiarisation

The researcher read all survey responses, noting initial ideas and reflections on AIGC content.

2. Generating codes

Interpretive codes were then produced to capture meaningful features relevant to the research questions, such as content structure, writing style, human presence, and tone.

3. Constructing initial themes

The codes were grouped to form candidate themes.

4. Reviewing themes

The researcher then checked the themes against coded data and the dataset to refine boundaries.

5. Defining and naming themes

Each theme was named to highlight its central meaning and its relationship to the research questions. These were human anchoring, functional clarity of content, and context-dependent interpretations of human-likeness.

6. Producing the report

A narrative was built around the collected data and themes to generate meaningful findings, and it is discussed in more detail in Chapter 5.

3.10 Chapter conclusion

This chapter outlined the methodological approach used to examine user preferences for GenAI-authored and human-authored digital marketing content. The study was grounded in a pragmatist paradigm and informed by Source Credibility Theory, allowing user perceptions to be examined through both measurable evaluation patterns and participants' own explanations. A mixed methods experimental survey design was adopted to address the two research questions. The 2×2 design enabled comparison between GenAI-authored and human-authored landing page content, as well as comparison between disclosed and undisclosed authorship conditions. Quantitative survey items captured patterns in participants' evaluations, while open-ended responses provided further insight into the reasoning behind those preferences. Analysis of the captured data was conducted using descriptive statistics and thematic analysis to produce findings that illuminate how and why users evaluate GenAI-authored content. The chapter also described how participants took part via CloudResearch Connect, how the survey was structured and administered through Qualtrics, and the ethical safeguards and validity strategies that governed the study. Together, these methodological choices were designed to produce findings that are both empirically grounded and interpretively rich, supporting the examination of context-sensitive and evolving user perceptions of GenAI-authored content in digital marketing. The following chapter presents the quantitative and qualitative results of this study.

Chapter 4 Results

4.1 Chapter overview

The purpose of this chapter is to present the empirical findings of the study examining users' perceptions of GenAI-authored and human-written marketing content. Two central research questions guide the analysis:

Research question 1 (RQ1): Do users prefer GenAI-authored or human-authored website landing page content when authorship is not disclosed?

Research question 2 (RQ2): How does authorship disclosure influence users' preferences for GenAI-authored and human-authored website landing page content?

This chapter presents the study's findings in four parts. First, descriptive statistics are used to compare participant evaluations of the GenAI-authored and human-written landing pages. Second, the effect of authorship disclosure on those evaluations is examined, including any order effects observed across experimental conditions. Third, thematic analysis of participants' qualitative responses reveals recurring patterns in how and why they evaluate content. Finally, findings related to participants' prior experience with generative AI and their attitudes toward its use in marketing contexts are reported.

4.2 Sample characteristics

As detailed in Section 3.5, 100 participants were recruited through the CloudResearch Connect platform and completed the survey. Five entries were excluded for incomplete or invalid responses, resulting in a final dataset of 95 participants. All analyses reported in this chapter are based on these 95 cases. The participant sample broadly reflected an adult, primarily North American population. It was gender balanced, with a slight male majority (52%). Participants ranged from 18 to over 65 years of age, though most were between 25 and 44 years, with the largest group falling within the 35–44 age bracket (38%). The sample was predominantly based in the United States, with a small proportion of respondents from Canada and the United Kingdom. Moreover, participants were randomly assigned to either a treatment condition (authorship was disclosed) or a control condition (authorship not disclosed), with approximately equal numbers in each group. There were 43 participants in the treatment group and 52 in the control group. In addition, to account for potential order effects, participants were randomly assigned to one of two presentation sequences: 46 participants viewed the human-authored landing page first, and 49 viewed the GenAI-authored page first.

4.3 Content evaluations and user preferences without authorship disclosure (RQ1)

The study used a survey to examine how users evaluated GenAI-authored and human-authored landing page content, and which they preferred, across both disclosed and undisclosed authorship conditions (See details of the Research design in Section 3.4). This section presents findings from the group that evaluated both landing pages without knowing which was GenAI-authored and which was human-authored, using descriptive statistics as outlined in Section 3.8. Participants rated each page across nine perception constructs (See Section 3.6.1), answering the following question for each of the two webpages:

Please read the webpage carefully and indicate how much you agree or disagree with the statements that follow.

1. *Based on this page, the content is authentic.*
2. *Based on this page, the content is credible.*
3. *Based on this page, the content is persuasive.*
4. *The author appears knowledgeable about the topic.*
5. *The writing on this page sounds human.*
6. *I connected with the writing style.*
7. *I am interested in learning more about this product.*
8. *Based on this page, the brand seems genuine.*
9. *Overall, the content on this page is high quality.*

When authorship was not disclosed, both webpages received broadly positive ratings, with means ranging from 3.23 to 4.23 across all nine constructs (See Table 4.1). The GenAI-authored page received higher mean ratings than the human-authored page across most constructs, with the largest mean differences observed in overall quality (4.02 vs. 3.67), credibility (4.16 vs. 3.90), product interest (3.75 vs. 3.48), and author knowledge (4.23 vs. 4.13). Standard deviations were generally lower for the GenAI-authored page, indicating greater consistency in participant ratings. Writing connection and human-likeness were exceptions, showing wider score dispersion across both pages, particularly for the human-authored page.

Table 4.1*Descriptive Statistics for Webpage Ratings Without Authorship Disclosure*

	Human-authored page		GenAI-authored page	
	Mean ratings	SD*	Mean ratings	SD*
Authenticity	4.04	0.82	4.10	0.83
Credibility	3.90	0.91	4.16	0.76
Persuasiveness	3.77	1.00	3.90	0.88
Author knowledge	4.13	0.79	4.23	0.78
Human-likeness	3.79	1.05	3.88	0.92
Writing connection	3.23	1.10	3.42	0.98
Product interest	3.48	1.26	3.75	1.05
Brand genuineness	3.85	0.96	4.02	0.85
Overall quality	3.67	1.10	4.02	0.90

Table 4.1 Mean ratings and standard deviations across nine perceptual constructs for human-authored and GenAI-authored webpages, evaluated by participants prior to authorship disclosure (N = 52; scale: 1–5). Without authorship disclosure, the GenAI-authored page received higher ratings across most constructs, with the most pronounced advantages observed in credibility, author knowledge, and overall quality.

Note. *SD = Standard deviation

Response distributions across both pages, presented in Figure 4.1, illustrate the generally positive evaluations while exposing notable variation across individual constructs. Credibility, author knowledge, and overall quality recorded the strongest agreement, with approximately 70% to 85% of participants selecting "Agree" or "Strongly Agree," and this pattern was slightly more pronounced for the GenAI-authored page. Writing connection and human-likeness returned weaker agreement across both pages. Writing connection recorded approximately 55% to 60% agreement, alongside over 25% neutral responses and around 20% disagreement. Human-likeness followed a similar pattern, though the human-authored page returned marginally higher agreement at approximately 65% to 70%, compared with 60% to 65% for the GenAI-authored page.

Figure 4.1

Proportional Distribution of Ratings for the Human-authored Page and GenAI-authored Page Recorded in Survey Section 1 (Without Authorship Disclosure)




Figure 4.1 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for the nine perceptual constructs comparing human-authored and GenAI-authored webpages evaluated by participants prior to authorship disclosure (N = 52). GenAI-authored pages received a higher proportion of agree and strongly agree responses across most constructs.

While Section 1 of the survey asked participants to rate each webpage independently, Section 2 asked participants to make direct comparisons between the two. In their responses, participants indicated which page they would act on and which they

preferred across each of the nine constructs. These comparative judgements directly contribute to Research Question 1 by revealing explicit preference patterns in the absence of authorship disclosure. The survey asked the following question:

Q. The following statements are aspects to compare between the human-authored page (Home Start) and the ChatGPT-generated page (FirstNest Realty). For each one, indicate which page was better in your opinion:

1. *Authenticity of the content*
2. *Credibility of the information*
3. *Persuasiveness of the page*
4. *How knowledgeable the content appears*
5. *Writing style resonated with me*
6. *Human-sounding writing*
7. *Made me want to learn more about the product*
8. *How genuine the brand seems (based on the page)*
9. *Overall quality of the page's content.*

As shown in the grouped bar chart in Figure 4.2, when authorship was not disclosed, 56% of participants preferred the GenAI-authored page, compared with 35% for the human-authored page. This pattern was consistent across all nine constructs, with the GenAI-authored page receiving higher preference proportions in every case. Preference for the GenAI-authored page was strongest for persuasiveness (54%) and overall quality (48%). At the same time, the narrowest gaps were observed for human-likeness and product interest, where human-authored page preference reached approximately 33% and 29%, respectively. The human-authored page's highest preference share was recorded for authenticity (38%) and its lowest for credibility (27%). Neutral responses were most prominent for product interest (29%), credibility (27%), and human-likeness (23%). Unclear responses were negligible across all constructs, ranging from 0% to 2%.

Figure 4.2

Proportional Distribution of Preferences for the Human-Authored Page and GenAI-authored Page Recorded in Survey Section 2 (Without Authorship Disclosure)




Figure 4.2 Grouped bar chart showing the proportional distribution of participant preferences across nine perceptual constructs and an overall webpage preference measure, comparing human-authored and GenAI-authored webpages prior to authorship disclosure (N = 52). The GenAI-authored page was preferred across all constructs, with the largest margins observed in persuasiveness, overall quality, and overall webpage preference.

4.4 Impact of authorship disclosure on content evaluations and user preferences (RQ2)

This section presents findings from the group that evaluated both landing pages, with knowledge of which was GenAI-authored and which was human-authored. Participants were shown a label identifying the source of each page, presented as “Webpage A was written by a human” or “Webpage B was written by ChatGPT” (See Appendix E: Survey). To examine whether authorship disclosure influenced content evaluations and to help answer Research Question 2, their ratings across the same nine constructs are analysed in this section using descriptive statistics.

When authorship was disclosed, participants gave the human-authored page numerically higher mean ratings than the GenAI-authored page across most constructs (See Table 4.2). The largest mean difference was observed for human-likeness (4.12 vs. 3.56), followed by brand genuineness (3.88 vs. 3.69) and authenticity (4.09 vs.

3.93). Mean differences were minimal for credibility (4.00 vs. 3.98), persuasiveness (3.86 vs. 3.81), and author knowledge (4.19 vs. 4.14). The GenAI-authored page returned higher ratings for writing connection (3.60 vs. 3.44) and product interest (3.51 vs. 3.49), while overall quality was identical across both pages at 3.77.

Table 4.2

Descriptive Statistics for Webpage Ratings With Authorship Disclosure

	Human-authored page		GenAI-authored page	
	Mean ratings	SD*	Mean ratings	SD*
Authenticity	4.09	0.75	3.93	0.77
Credibility	4.00	0.69	3.98	0.86
Persuasiveness	3.86	0.91	3.81	0.82
Author knowledge	4.19	0.66	4.14	0.68
Human-likeness	4.12	0.93	3.56	0.91
Writing connection	3.44	1.08	3.60	1.09
Product interest	3.49	1.05	3.51	1.03
Brand genuineness	3.88	0.82	3.69	0.84
Overall quality	3.77	0.97	3.77	0.90

Table 4.2 Mean ratings and standard deviations across nine perceptual constructs for human-authored and GenAI-authored webpages, evaluated by participants following authorship disclosure (N = 43; scale: 1–5). Following disclosure, the human-authored page received higher ratings across most constructs, with the most pronounced advantage observed in human-likeness, authenticity, and brand genuineness.

Note. *SD = Standard deviation

A comparison of ratings for the GenAI-authored page across the disclosure and no-disclosure conditions is presented in Table 4.3. Across most constructs, mean ratings for the GenAI-authored page were lower among participants in the disclosed condition than the undisclosed condition. The largest mean differences were observed for brand genuineness (4.02 vs. 3.69), human-likeness (3.88 vs. 3.56), overall quality (4.02 vs. 3.77), and product interest (3.75 vs. 3.51). Credibility and authenticity showed smaller mean differences. Writing connection was the only construct in which mean ratings for the GenAI-authored page were higher in the disclosed condition than the undisclosed condition (3.60 vs. 3.42).

Table 4.3*Descriptive Statistics for Webpage Ratings Across Conditions*

	GenAI page (without authorship disclosure)		GenAI page (with authorship disclosure)	
	Mean ratings	SD*	Mean ratings	SD*
Authenticity	4.10	0.83	3.93	0.77
Credibility	4.16	0.76	3.98	0.86
Persuasiveness	3.90	0.88	3.81	0.82
Author knowledge	4.23	0.78	4.14	0.68
Human-likeness	3.88	0.92	3.56	0.91
Writing connection	3.42	0.98	3.60	1.09
Product interest	3.75	1.05	3.51	1.03
Brand genuineness	4.02	0.85	3.69	0.84
Overall quality	4.02	0.90	3.77	0.90

Table 4.3 Mean ratings and standard deviations across nine perceptual constructs for the GenAI-authored webpage, comparing participant evaluations before and after authorship disclosure (N = 43 (with authorship disclosure); N=52 (without authorship disclosure); scale: 1–5). Ratings declined across most constructs following disclosure, with the most pronounced drops observed in human-likeness, brand genuineness, and authenticity.

Note. *SD = Standard deviation

Extending these evaluations to direct preference judgements, Figure 4.3 presents comparative preference distributions that address RQ2 by illustrating how preference patterns varied between the disclosed and undisclosed conditions for user preferences for GenAI-authored and human-authored website landing page content. When participants knew the source of the GenAI-authored page, their results were closely divided on overall webpage preference: 42% preferred the human-authored page, 40% preferred the GenAI-authored page, and 18% selected neutral. However, they selected the GenAI page as their preference across six of the nine constructs, including credibility (47%), persuasiveness (51%), author knowledge (42%), writing connection

(44%), product interest (40%), and overall quality (42%). The human-authored page received higher preference proportions only for authenticity (49%) and human-likeness (40%), while brand genuineness was evenly split between the pages at 33% each. Neutral responses were most prominent for human-likeness (28%), author knowledge (23%), and credibility (21%). Unclear responses were negligible across all constructs (0%–2%).

Figure 4.3

Proportional Distribution of Preferences for the Human-authored Page and GenAI-authored Page Recorded in Survey Section 2 (With Authorship Disclosure)




Figure 4.3 Grouped bar chart showing the proportional distribution of participant preferences across nine perceptual constructs and an overall webpage preference measure, evaluated following authorship disclosure (N = 43). With full knowledge of authorship, participants favoured the GenAI-authored page across most constructs, with the human-authored page receiving higher preference only on authenticity and human-likeness.

4.5 Changes in ratings across disclosure conditions

To directly address Research Question 2 and examine how authorship disclosure influenced user preferences, this section compares rating distributions for human- and GenAI-authored pages across the two participant groups. By placing these distributions

side by side, it is possible to examine whether knowing the source of each page meaningfully influenced participants' ratings across the nine constructs.

Figure 4.4

Proportional Distribution of Ratings for Human-Authoring Page Recorded in Survey Section 1 (Disclosed Versus Undisclosed Authorship)




Figure 4.4 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for nine perceptual constructs, comparing ratings of the human-authored page under undisclosed and disclosed authorship conditions (N = 43 (with authorship disclosure); N=52 (without authorship disclosure)). Ratings for the human-authored page remained largely stable or improved following disclosure, with the most notable gains observed in authenticity and human-likeness.

Broadly similar rating patterns were observed for the human-authored page across both conditions, as shown in Figure 4.4. In most constructs, response distributions overlapped closely, with agreement concentrated in the "Agree" and "Strongly Agree"

categories regardless of whether authorship was disclosed. The most visible differences were in the intensity of agreement rather than its direction. Participants in the disclosure condition showed slightly higher proportions of "Strongly Agree" responses for authenticity, brand genuineness, and overall quality. In comparison, those in the no-disclosure condition showed marginally higher proportions of "Agree" responses across the same constructs. Modest variation was also visible for writing connection and persuasiveness, though neither showed a meaningful shift in overall agreement levels across conditions.

Similar to the human-authored page, ratings for the GenAI-authored page showed broadly similar distributions across both conditions, with responses concentrated in the "Agree" and "Strongly Agree" categories across most constructs (See Figure 4.5). Again, differences were most visible in the intensity of agreement rather than its direction. When authorship was not disclosed, authenticity, credibility, and overall quality showed higher proportions of "Strongly Agree" responses, estimated at approximately 45%, 40%, and 45%, respectively. When authorship was disclosed, these constructs recorded lower proportions of "Strongly Agree" responses (approximately 35%, 30%, and 30%, respectively), with a correspondingly higher proportion of "Agree" responses. Notably, neutral responses were also higher in the disclosed condition, most notably for brand genuineness (approximately 12% versus 28%), overall quality (approximately 10% versus 20%), and human-likeness (approximately 20% versus 30%).

Figure 4.5

Proportional Distribution of Ratings for GenAI-Authored Page Recorded in Survey Section 1 (Disclosed Versus Undisclosed Authorship)




Figure 4.5 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for nine perceptual constructs comparing ratings of the GenAI-authored page under undisclosed and disclosed authorship conditions (N = 43 (with authorship disclosure); N=52 (without authorship disclosure)). In contrast to the human-authored page, ratings for the GenAI-authored page declined across most constructs following disclosure.

4.6 Qualitative findings

In addition to the structured survey measures, participants were asked a series of open-ended questions to capture the reasoning underlying their evaluations. The inclusion of qualitative responses enables a deeper examination of how participants interpreted the webpages, what specific cues shaped their judgements, and how they

justified their preferences. This is particularly important for addressing the research questions, as both RQ1 (preferences under undisclosed authorship) and RQ2 (the effect of authorship disclosure on user preferences) involve perceptual processes that cannot be captured through numerical ratings alone.

The following open-ended questions were asked in Section 2 of the survey:

Q. In one or two sentences, what was the most notable difference between the two pages?

Q. In one sentence, what, if anything, increased or reduced your trust in the page you liked more?

Q. What specific cues made the writing feel more human or more AI-like (e.g., phrasing, repetition, formality, or something else)? One example is enough.

Q. In one or two sentences, what specific thing about the page you would act on most influenced your decision?

Participants' responses to these questions were analysed using reflexive thematic analysis (Braun & Clarke, 2019), resulting in the following three themes.

Theme 1: Functional clarity as the primary basis for evaluation and action

Across responses, assessments were consistently grounded in how effectively information was presented and how easily the page supported understanding and decision-making. This theme, therefore, captures the extent to which participants prioritised clarity, structure, and usability when evaluating the landing pages. Participants frequently attributed trust and preference to aspects such as “flowed better” and “easier to digest all the info” rather than to a “wall of text”. For the GenAI-authored page in particular, participants who knew the source noted that it provided “clear, actionable information”, “focused on providing information in a clear and logical way”, and “page layout less cluttered”, perceiving it as more credible and useful. One participant noted that “the clear, step-by-step roadmap... made the overwhelming process feel manageable,” highlighting the role of structured guidance in reducing cognitive effort.

It was observed that, even for participants who were aware of the author, evaluations were primarily driven by functional performance rather than source-based considerations. In the context of a conversion-oriented landing page, clarity and organisation appear to operate as key indicators of competence, directly influencing both perceived credibility and behavioural intention. The prominence of these factors

suggests that participants engaged with the content in a task-oriented manner, prioritising usability and informational value over questions of authorship.

Theme 2: Human presence as a central mechanism of trust and reassurance

Identifiable human anchoring consistently functioned as a credibility cue across all responses. Participants' responses consistently reflected the importance of visible human elements in shaping perceptions of trust, credibility, and willingness to act. They cited cues such as named individuals, contact details, and identifiable agents as increasing the page's perceived reliability. For the human-written page in particular, responses highlighted the reassurance associated with "a real name and a direct phone number," "an actual real estate agent," "done by humans", and "someone I could call anytime." The reference to a specific individual on that landing page was described as "conversational and real", "more believable," "more personal," and easier to "connect" with. Conversely, the absence of such cues on the GenAI-authored page contributed to perceptions of the page as "generic" or less trustworthy.

In the context of landing pages, it appears that trust was constructed not solely through informational content but also through signals of accountability and relational presence. Therefore, the ability to identify and potentially contact a real person appeared to reduce perceived risk in a decision-focused digital environment. Importantly, this mechanism was observed both with and without authorship disclosure, indicating that human anchoring functions as a fundamental credibility cue independent of explicit source information.

Theme 3: Context-dependent and variable interpretations of human-likeness

Participants' judgements of the human or AI-likeness of the writing itself were inconsistent and context-dependent. Rather than relying on consistent linguistic markers, participants evaluated tone, phrasing, and structure in relation to contextual expectations and perceived appropriateness. Some participants associated human-like writing with a "more personal and friendly feel," "less formal" tone, or conversational flow, while others preferred language that was "professional" and "business-like." Notably, similar features were interpreted differently across responses. For example, the use of informal expressions such as "FOMO" on the human-written page was described by some as natural and relatable, and by others as "trying too hard", "forced", "robotic", or inauthentic. Likewise, "perfect grammar" and highly structured formatting on the GenAI-authored page were, at times, viewed as indicators of artificiality, particularly when they appeared overly regular or "mechanical."

In essence, perceptions of human-likeness were not determined by fixed stylistic features but by the extent to which the communication style aligned with contextual expectations. In a digital decision-making environment such as a landing page, participants appeared to evaluate tone based on its appropriateness and credibility rather than its degree of informality. As a result, judgements of humanness or AI-likeness were fluid and context-sensitive, reflecting individual interpretations rather than a shared standard.

4.7 Order effects

This section examines whether the sequence in which the webpages were presented influenced participants' evaluations. As participants viewed both webpages consecutively, there is potential for earlier exposure to shape subsequent judgements. Assessing order effects helps determine whether observed differences in ratings are attributable to webpage characteristics or are partially influenced by presentation sequence (Meier & Song, 2026). Figure 4.6 and Figure 4.7 compare participants' responses based on the first webpage viewed. When the GenAI-authored page was shown first, the subsequently viewed human-authored page displayed higher proportions of "Agree" and "Strongly Agree" responses across overall quality, brand genuineness, connection, and human-likeness. Conversely, when the human-authored page was shown first, the GenAI-authored page showed stronger agreement levels across many of the same constructs. In both presentation orders, disagreement remained low, and the primary difference lay in the increased intensity of agreement for the second page viewed.

Figure 4.6

Proportional Distribution of Ratings for the Webpages Among Participants who Viewed the Human-Authored Page First Recorded in Survey Section 1




Figure 4.6 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for nine perceptual constructs, comparing human-authored and GenAI-authored pages among participants who viewed the human-authored page first (N = 43 (with authorship disclosure); N=52 (without authorship disclosure)). The GenAI-authored page received a higher proportion of agree and strongly agree responses across most constructs.

Figure 4.7

Proportional Distribution of Ratings for the Webpages Among Participants who Viewed the GenAI-authored Page First Recorded in Survey Section 1




Figure 4.7 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for nine perceptual constructs, comparing human-authored and GenAI-authored pages among participants who viewed the GenAI-authored page first (N = 43 (with authorship disclosure); N=52 (without authorship disclosure)). The human-authored page received a higher proportion of agree and strongly agree responses across most constructs.

4.8 Participants' GenAI experience and attitudes

This section reports participants' self-reported experience with GenAI and their broader attitudes towards its use in marketing contexts. While not directly addressing the research questions, these measures provide important explanatory context for the

evaluative patterns observed in earlier sections. In particular, they offer insight into how pre-existing familiarity with, and attitudes towards, AI may shape initial content preferences (RQ1), as noted in more detail in Section 2.7. They also contextualise responses to authorship disclosure (RQ2), particularly where awareness of AI authorship may have activated pre-existing beliefs or biases that shaped participants' judgements throughout the survey.

Participants were asked how frequently they use generative AI tools.

Q. Which best describes how often you personally use generative AI tools (e.g., ChatGPT, Gemini, Claude)?

- *Never*
- *Less than once a month*
- *1 to 3 times a month*
- *1 to 2 times a week*
- *to 6 times a week*
- *Daily or more*

For reporting purposes, responses were consolidated into three categories. Infrequent users included those who indicated never using AI or using it less than once a month to one to three times per month. Moderate users were defined as those reporting use one to two times per week. Frequent users included participants who reported using generative AI three to six times per week or daily. Based on this classification, Figure 4.8 illustrates that 47% of participants were frequent users, 36% were infrequent users, and 17% reported moderate weekly use. Overall, these figures indicate that nearly half of the sample regularly engage with generative AI tools, while just over one-third report only occasional or minimal use. The distribution was largely similar within the treatment condition (49% frequent, 30% infrequent, 21% moderate use) and the control condition (46% frequent, 40% infrequent, 14% moderate use).

Figure 4.8

Distribution of Self-Reported GenAI Usage Frequency Among Participants Recorded in Survey Section 3




Figure 4.8 Pie chart showing the distribution of self-reported GenAI usage frequency among participants, categorised as infrequent, moderate, or frequent (N = 95). The majority of participants reported frequent use of GenAI, indicating a relatively experienced sample.

In addition, participants' attitudes towards the use of AI in marketing were examined through descriptive analyses of Likert-scale responses and open-ended responses.

The following statements are about brands using AI to create marketing content. Please indicate how much you agree or disagree with each statement.

Q. I am comfortable with brands using AI to create their marketing content. (Comfort with AI use in Marketing)

Q. Knowing who created marketing content helps me judge it fairly. (Perceived value of source information)

Q. I trust marketing content less when it is labelled as AI-generated. (Reduced trust in AI-labelled content)

Q. I want brands to clearly state when marketing content is AI-generated. (Preference for authorship disclosure)

Figure 4.9

Proportional Distribution of Participants' Attitudes Towards Generative AI Use in Marketing Recorded in Survey Section 3




Figure 4.9 Stacked bar chart showing the proportional distribution of participant responses across a five-point agreement scale (1 = Strongly Disagree to 5 = Strongly Agree) for four attitudinal constructs (comfort with AI use in marketing, perceived value of source information, reduced trust in AI-labelled content, and preference for authorship disclosure) reflecting participant attitudes towards GenAI use in marketing (N = 95). Participants expressed strong support for authorship disclosure and a high perceived value of source information. At the same time, responses to comfort with AI use in marketing and reduced trust in AI-labelled content were more divided.

Figure 4.9 shows that responses across all four AI attitude items are concentrated primarily in the “Agree” and “Strongly Agree” categories. The two variables “Comfort with AI use in Marketing” and “Reduced trust in AI-labelled content” show a mix of agreement and neutrality, with relatively little disagreement. “Perceived value of source information” and “Preference for authorship disclosure” show stronger clustering in the upper agreement categories, with minimal disagreement.

Towards the end of the survey, participants' views regarding brands using AI for marketing communications were recorded by asking:

Q. In one or two sentences, how do you feel about brands using AI to write narrative marketing content (e.g., social posts, ads, email sequences)?

Responses indicated that a central tension between efficiency and human authenticity shapes attitudes towards GenAI-authored marketing content. Many participants expressed concern that AI-produced content lacks a “human touch,” describing it as “generic,” “robotic,” or “disingenuous.” Marketing was frequently framed as a fundamentally human activity requiring originality and creativity, and the use of AI was sometimes interpreted as a signal that brands prioritise cost-cutting or output over genuine connection. These concerns extended beyond writing style to broader issues of integrity, labour displacement, and brand values. Participants argued that brands “should be hiring and paying real people with the skills to create this content,” viewing the replacement of human writers as dehumanising.

At the same time, scepticism was often conditional rather than absolute. Concerns about authenticity were intertwined with worries about accuracy and accountability. Some participants noted that “the accountability is not guaranteed” and that they would “trust it less if AI was used to create everything.” However, many articulated pragmatic acceptance provided that AI use is transparent and subject to human oversight. Statements such as “I have no problem with it as long as it is disclosed” and “a human always needs to review AI generated content” reflect a preference for AI as a supportive tool rather than a full substitute for human creativity. A minority also expressed clearly positive views, describing AI as efficient or “the way of the future”. However, overall responses suggest that legitimacy depends on disclosure, accountability, and the retention of a recognisable human voice.

4.9 Chapter conclusion

The results show that, without authorship disclosure (RQ1), participants rated both webpages positively, with slightly higher mean ratings and higher proportions preferring the GenAI-authored page across most constructs. Comparative evaluations further indicate that a greater proportion of participants selected the GenAI-authored page as superior overall. However, when authorship was disclosed (RQ2), evaluation patterns shifted. The human-authored page received higher mean ratings on human-likeness, authenticity, and brand genuineness. At the same time, other constructs, such as credibility, persuasiveness, and author knowledge, remained closely aligned between the two pages. Comparative preference data under disclosure showed a more balanced distribution. However, the GenAI-authored page was still preferred for six of nine constructs. In contrast, the human-authored page was preferred for authenticity and human-likeness. Neutral responses increased across several measures following disclosure.

Qualitative findings indicate that evaluations were primarily grounded in clarity, structure, and usability. In addition, trust was strongly influenced by cues of human presence, such as identifiable individuals and contact information. Perceptions of human-likeness varied across participants and were shaped by contextual expectations rather than consistent stylistic features. Furthermore, order effects analysis shows that the second webpage viewed tended to receive higher agreement intensity, regardless of whether it was GenAI-authored or human-written. Finally, participants reported generally high familiarity with generative AI. They expressed mixed attitudes towards its use in marketing, with responses indicating both acceptance of its utility and concerns regarding authenticity and trust. These findings collectively set the stage for the discussion in Chapter 5, which interprets these patterns in relation to existing literature on GenAI content and considers their implications for digital marketing practice.

Chapter 5 Discussions

5.1 Chapter overview

This chapter discusses the findings presented in Chapter 4 in relation to the existing literature on user perceptions of GenAI-authored content in digital marketing contexts. The discussion is organised around four key themes emerging from the data. These include two findings related to the undisclosed condition: the evaluative advantage of GenAI-authored content when authorship is undisclosed, and the neutralisation of relational evaluation when source cues are absent. A further theme discusses the redistributive rather than reductive effect of authorship disclosure. Lastly, broader contextual factors that shape user responses are examined, including participants' prior experience with generative AI and their attitudes toward its use in marketing. Taken together, the findings from these themes point to a dual evaluative framework, grounded in functional competence and relational authenticity, through which users appear to assess functional and persuasive landing page content authored by both GenAI and humans. The implications of this framework are discussed throughout the chapter and synthesised in the concluding section.

5.2 GenAI-authored content performs favourably when authorship is unknown

In the absence of authorship information, GenAI-authored website landing page content cannot only meet but also exceed the perceived effectiveness of human-authored content. Evaluated on its own merits, the GenAI page performs at least as well and, on several dimensions, more favourably (See Table 4.1), despite participants being unable to distinguish between the two reliably. Comparative evaluations reinforced this inclination, with the GenAI-authored page preferred more frequently across most measures and associated with a greater likelihood of acting on its content (See Figure 4.2). This pattern is consistent with Ou and Ho's (2024) distinction between message credibility and source credibility. When authorship is undisclosed, and no source cue is available, evaluation defaults to message credibility, defined as how believable and trustworthy the information itself is judged to be, based on its perceived quality and fluency. In this condition, content is assessed by what it communicates rather than by who produced it, and, on that basis, the GenAI-authored page performed strongly.

The favourable evaluation of GenAI-authored landing page content can also be understood through the distinction between cognitive and affective trust (Johnson &

Grayson, 2005), outlined in Section 2.3.1. Without source attribution, participants primarily relied on cognitive trust, evaluating both pages on observable features such as clarity, organisation, and persuasiveness, which produced consistent judgements of competence. Through this lens, it appeared that competence-related dimensions, including credibility, persuasiveness, author knowledge, and overall quality, were consistently rated more favourably for the GenAI-authored page, with notably high agreement among participants (See Figure 4.1). Evidence from persuasion-oriented digital marketing contexts strongly supports this finding. Across product advertisements, campaign messages, social media posts, and email marketing, GenAI-authored content has been found to match or outperform human-authored equivalents on dimensions directly relevant to marketing effectiveness, including quality, persuasiveness, clarity, and professionalism (Acatrinei & Miu, 2025; Meguellati et al., 2025; Park et al., 2024; Y. Zhang & Gosline, 2023). When classical persuasion principles were applied, this advantage became more pronounced, and personalised GenAI-authored messages similarly demonstrated greater influence on consumer decision-making (Matz et al., 2024; Meguellati et al., 2025). While these studies vary in format and platform, they converge on a consistent pattern in which GenAI-authored content is not merely comparable to human-authored content on functional dimensions. However, they may, in fact, hold a measurable advantage.

The concentration of favourable evaluations on competence-related dimensions also points to the possible role of heuristic processing in shaping user judgements. In addition to informational scrutiny, users rely on cognitive heuristics, defined as mental shortcuts triggered by observable, surface-level signals, to form rapid evaluative judgements (Metzger et al., 2010). This is consistent with the survey's qualitative responses, which clustered around a theme of functional clarity as a dominant evaluative dimension. Participants consistently attributed trust and preference to presentational features rather than authorship, describing the GenAI-authored content as "easier to digest all the info" and as "flowed better" without becoming a "wall of text." These responses align with the presentational cues that Sundar (2008) identifies as triggering automatic, heuristic-based judgements. When applied to machine-generated outputs, this suggests users may associate syntactic cleanliness and structural consistency with efficiency and accuracy, and the present findings suggest that GenAI-authored landing page content may satisfy these triggers.

The functional nature of landing pages as a content format further helps explain why heuristic processing may have played a role in user engagement. Landing pages are goal-oriented digital environments designed to communicate information quickly and guide users towards action, prioritising clarity and usefulness over emotional

expression (Dvir & Gafni, 2018). In their research on web usability, Nielsen and Pernice (2009) assert that users scan rather than read website content, relying on headings, layout, and concise text to extract relevant information efficiently. In this environment, well-organised, easy-to-process content is more likely to be perceived as credible and useful, not necessarily because users have evaluated it deeply, but because it reduces cognitive effort and aligns with their task. This aligns with the content quality constructs established in Section 2.3, in which quality is defined as material that is valuable, relevant, and engaging for a specific target audience. Research also shows that high-quality website content is strongly correlated with improved usability and conversion rates (Almaghlouth, 2024; Bubphapant & Brandão, 2024). In particular, GenAI-authored content appears well-positioned to satisfy these expectations. Hartmann et al. (2025) demonstrate that in structured digital marketing contexts, GenAI can reproduce the persuasive logic, stylistic conventions, and communicative structure associated with effective marketing copy. Similarly, Ouyang et al. (2022) explain that generative models trained on large volumes of professional text can produce outputs that meet the formal expectations of a given genre without requiring deliberate creative effort. On a landing page, users primarily ask whether the content helps them understand the offer, assess its relevance, and decide what to do next (Thomas & Mathew, 2023). In a format where conventions are predictable and expectations are clear, GenAI's capacity to meet these expectations becomes a measurable advantage. The qualitative results in Section 4.6 reflect this pattern. The GenAI-authored page was described as offering "clear, actionable information" and as being "focused on providing information in a clear and logical way," with a "page layout less cluttered" than the alternative. One participant noted that "the clear, step-by-step roadmap... made the overwhelming process feel manageable." These responses speak directly to cognitive effort reduction, and are consistent with the possibility that users were engaging in heuristic rather than deliberative processing. In this sense, favourable evaluations in the undisclosed condition appear to reflect the functional competence of the content rather than any deliberate inference about its origin. It should be noted that processing mode was not directly measured in this study. The heuristic account offered here is an interpretation of outcome and qualitative data, not a tested mechanism.

5.3 Relational constructs remain suspended in the absence of source cues

While functional and competence-based criteria drove favourable evaluations of the GenAI-authored page, a markedly different pattern emerged across relational dimensions. Unlike cognitive trust, affective trust reflects emotional reassurance and identity-based responses to the communicator rather than rational assessments of

competence (Johnson & Grayson, 2005). Applied to the findings, relational dimensions such as authenticity, human-likeness, and brand genuineness revealed a markedly different pattern than cognitive dimensions. Relational constructs were evaluated far more closely across both pages, with a substantial proportion of neutral responses indicating difficulty in detecting meaningful differences (See Figure 4.1). This pattern carries theoretical significance. Within digital marketing literature, authenticity is not an inherent property of content but a judgement formed by users based on their experiences, expectations, and perceptions of the source (Zniva et al., 2023). When no authorship cue is provided, participants lacked the evaluative anchor through which such judgements are typically formed. This comparability did not reflect GenAI's ability to signal humanness, but the absence of any benchmark against which to assess it. A similar logic applies to brand genuineness, understood as the extent to which a brand is perceived as honest, legitimate, and consistent in its communication (Wymer & Akbar, 2017). Without source disclosure, participants had no basis on which to form judgements of the brand's integrity. The result was a high percentage of neutral responses, rather than positive or negative attribution on relational dimensions.

The inability to reliably distinguish between GenAI and human-written content is consistent with evidence from research on short stories, social media posts, and news articles, where users struggle to differentiate GenAI-authored from human-authored content on affective or stylistic grounds when source information is withheld (Doshi & Hauser, 2024; Hussain, 2024). Digital marketing research confirms that this attributional pattern operates independently of content origin (Parshakov et al., 2025; Zhu et al., 2025), suggesting that users reward functional competence rather than authorship under undisclosed conditions. Applied to the landing page context, this was directly reflected in participants' qualitative responses (See Section 4.6), revealing a lack of stable criteria for evaluating authenticity or humanness. Language described as "conversational and real" by one participant was labelled "generic" or "trying too hard" by another, with subjective interpretations projected onto largely similar material. When pressed to justify preferences, participants consistently defaulted to concrete, functional criteria. This is consistent with Ouyang et al. (2022), who established that generative models that replicate predictable persuasive conventions leave users without clear reference points for distinguishing authorship on stylistic grounds alone. This finding diverges from research suggesting that users can detect and respond negatively to the perceived absence of human intent and emotional depth in GenAI content (Kirk & Givi, 2025; Stein et al., 2024). The constructs of human-likeness and writing connection, with high neutral responses (See Table 4.1), appear to require a source cue to activate (See Section 2.6). Without one, relational qualities that might otherwise distinguish human-authored content are effectively neutralised, leaving

functional attributes as the primary basis for evaluation. Without authorship cues, GenAI-authored content is judged on what it does rather than who produced it. As will be discussed in the following section, these relational constructs do not disappear from user evaluation entirely. However, they remain suspended until a source cue becomes available to activate them.

5.4 Authorship disclosure shifts rather than reduces GenAI content evaluations

When authorship is disclosed, results show a decline in the perceived quality of GenAI-authored website landing page content. However, this decline is not a direct rejection of GenAI content, but a redistribution of evaluative focus. This distinction is theoretically significant, as the findings refute the pattern of AI aversion described by Dietvorst et al. (2015) as users' mistrust of algorithms due to their error-prone nature. Neither do they align with research suggesting that users systematically devalue content once it is known to be GenAI-authored (Camilleri, 2024), even when indistinguishable from human writing (Casal & Kessler, 2023; Darda et al., 2023). Most directly, the findings diverge from Longoni et al.'s (2019) disclosure penalty, which holds that identical content is rated less favourably once AI authorship is known. If this penalty operated as documented, a consistent and meaningful decline in evaluations of the GenAI-authored page would be expected upon disclosure. The modest magnitude of the observed differences suggests otherwise (See Table 4.2 and Table 4.3). The findings equally challenge prior research suggesting that disclosure uniformly reduces trust, engagement, and perceived brand authenticity in general marketing contexts (Arango et al., 2023; Brüns & Meißner, 2024; Zhang et al., 2024). In fact, overall webpage preference remained nearly evenly split under disclosure, with no decisive shift away from the GenAI-authored page (See Figure 4.3). This pattern is more consistent with research positing that authorship disclosure prompts effortful evaluation rather than outright dismissal (Nakano et al., 2025). Rather than rejecting the content outright, participants seemed to engage more deliberately and critically, introducing greater caution without consistent negative judgement.

In conversion-oriented digital environments, trust is constructed not solely through informational content but also through signals of accountability and relational presence (Davidavičius & Limba, 2022). A key finding of this study reflects this directly. Authorship disclosure did not reveal a decline in GenAI content evaluations but rather an increase in human-authored content on relational dimensions. Despite balanced overall preferences, the GenAI-authored page continued to perform strongly on competence-related attributes such as clarity, structure, and informativeness (See

Figure 4.3). The more pronounced shift occurred on relational attributes such as authenticity, human-likeness, and brand genuineness, where the human-authored page gained a measurable advantage under disclosure (See Figure 4.4). This pattern aligns with Y. Zhang and Gosline (2023), who found that disclosure activated source-based bias, or human favouritism, rather than penalising GenAI content. The qualitative responses gathered in this study capture the mechanism directly. Participants consistently cited visible human elements, such as named individuals, contact details, and identifiable agents, as credibility cues that increased perceived reliability. For the human-authored page specifically, participants described the reassurance of "a real name and a direct phone number," "an actual real estate agent," and "someone I could call anytime," with such cues described as "more believable," "more personal," and easier to "connect" with. Conversely, the absence of these cues contributed to perceptions of the GenAI-authored page as "generic" or "robotic." Crucially, human anchoring functioned as a credibility cue independent of explicit source information. Disclosure, in this sense, did not undermine the recognised competence of GenAI-authored content. It introduced a relational benchmark and the expectation of human presence.

5.4.1 Dual evaluative framework

As the preceding discussion establishes, the findings diverge from existing literature in two significant ways. First, comparability between GenAI and human-authored content in the undisclosed condition was uneven, with competence-based constructs favouring the GenAI page, whereas relational constructs yielded neutrality. Second, the expected disclosure penalty (Longoni et al., 2019) did not materialise uniformly, with decline concentrated on relational rather than functional dimensions. These divergences suggest that a single evaluative lens, whether message credibility or source credibility alone, is insufficient to account for the full pattern of user responses. As established in Section 2.3.1, Source Credibility Theory (Hovland et al., 1953) distinguishes between perceived expertise and perceived trustworthiness. Camilleri (2024) demonstrates that source disclosure triggers an affective, trustworthiness-based response that can operate independently of competence-based assessment. Together, these divergences point to the need for an evaluative framework capable of attending to both dimensions simultaneously.

These findings are proposed here as evidence of a dual evaluative framework, in which GenAI-authored and human-authored landing page content appear to have been judged against two overlapping but distinct criteria: functional competence and relational authenticity. When authorship was unknown, evaluations were largely

performance-based, with participants prioritising clarity, structure, informativeness, and ease of use. In this mode, the GenAI page performed strongly because it aligned with the functional expectations of landing page content. However, once authorship was disclosed, competence-based judgements were not abandoned. A second layer of evaluation, based on source, human presence, and perceived authenticity, appears to have been introduced alongside them. This proposed framework offers one account of why GenAI-authored content continued to be rated positively on functional measures while human-authored content gained an advantage on relational dimensions. The findings, therefore, suggest that users do not evaluate GenAI landing page content through a simple acceptance-versus-rejection lens. Rather, evaluation is better characterised by a dual framework in which GenAI can be recognised as competent and useful, while human-authored content is granted additional relational value when authenticity, trustworthiness, and human connection become salient. As a single-study finding, this framework is offered as a proposition rather than an established model, and its application beyond this landing page context has not yet been tested.

5.5 Context and familiarity shape responses to GenAI-authored content

User responses to GenAI-authored content are not uniform. In fact, they vary by context, task type, and individual familiarity. Prior research establishes that functional and task-oriented environments produce more favourable evaluations, as users prioritise clarity and informational utility over authorship, while creative or emotionally expressive contexts tend to trigger greater resistance (Castelo et al., 2019). The present findings sit at the intersection of these two tendencies. As an action-oriented format that asks users to commit, enquire, or engage, the landing page introduced a level of perceived risk that heightened source-based scepticism under disclosure. Liu-Thompkins et al. (2022) posit that high-conversion contexts make signals of human involvement more important for establishing trust and accountability, and this was reflected directly in participants' responses. Some expressed a preference for content "done by humans" not because the GenAI-authored content failed functionally, but because the landing page environment raised the threshold for relational trust. Under disclosure, features that previously signalled competence, such as 'perfect grammar' and structured formatting, were reframed as 'mechanical,' a pattern consistent with broader research showing that users perceive AI as lacking the intentionality and lived experience required for authentic expression (Arango et al., 2023; Kirkby et al., 2023).

Individual familiarity with GenAI and attitudes towards its use in marketing further moderated these responses. Prior research indicates that more experienced or

positively inclined users are less susceptible to source-based bias (Kubovics, 2025a; Li et al., 2025). Nearly half of the participants in this study were frequent GenAI users (See Figure 4.8), which may help explain why disclosure did not produce stronger or more consistent rejection of the GenAI-authored page. Attitudinal data revealed a broader tension. Participants expressed strong support for disclosure and placed high value on source information, yet comfort with AI use in marketing and trust in AI-labelled content were more divided (See Figure 4.9). Scepticism was rarely absolute. Many described GenAI content as "generic," "robotic," or lacking a "human touch," with some arguing that brands "should be hiring and paying real people." Yet conditional acceptance was common, with participants stating "I have no problem with it as long as it is disclosed" and "a human always needs to review AI generated content." These findings are also temporally situated. The rapid normalisation of GenAI since ChatGPT's release in 2022 (Cheng et al., 2025) means user familiarity and evaluative frameworks are shifting at a pace that outstrips much of the existing literature, suggesting the disclosure penalty documented in earlier research may already be weakening. Acceptance of GenAI-authored content in digital marketing settings is therefore neither fixed nor unconditional. Rather, the context of the digital environment mediates it, along with individual attitudes towards AI, and a rapidly shifting baseline of familiarity that is shaping those attitudes.

A layer of caution is warranted in interpreting these findings. Order effects were observed across the study, with the second webpage viewed generally rated more favourably than the first, regardless of authorship (See Figure 4.6 and Figure 4.7). This aligns with research on additive and subtractive order effects, where the sequence in which similar items are presented can systematically raise or lower evaluations (Moore, 2002). In this study, the second page may have benefited from participants' prior familiarity with the task and comparison frame. While the findings support a modest preference for the GenAI-authored page under unknown authorship, this preference should therefore be interpreted as a small but meaningful pattern rather than a definitive authorship effect.

5.6 Chapter conclusion

Taken together, the findings discussed in this chapter reveal a consistent and theoretically significant pattern in how users evaluate GenAI-authored website landing page content. Without authorship information, GenAI-authored content performed at least as well as, and in several respects more favourably than, human-authored content. A preference for GenAI itself did not drive this advantage. In fact, it was message credibility and heuristic processing that rewarded surface-level clarity,

structure, and functional coherence of the website landing page. In parallel, relational dimensions including authenticity, human-likeness, and brand genuineness produced evaluative neutrality rather than meaningful differentiation, as the absence of source cues left participants without the evaluative anchor through which such judgements are typically formed.

Authorship disclosure did not reverse this pattern but redistributed evaluative focus. The expected disclosure penalty did not materialise uniformly. Instead, source credibility became an active input alongside message credibility, introducing a relational benchmark against which the GenAI-authored page was measured. The disclosure that was activated was not algorithm aversion but human favouritism, reflected in higher ratings of the human-authored page on relational dimensions. In contrast, competence-based evaluations of the GenAI-authored page remained largely intact. These responses were further shaped by the landing page's conversion-focused format, which heightened perceived risk and raised the threshold for relational trust when authorship was known.

Collectively, these findings support a proposed dual evaluative framework for understanding how both pages were assessed. One grounded in functional competence, or the extent to which content is clear, credible, and persuasive, and the other in relational authenticity, or the extent to which content signals human presence and connection. GenAI-authored content performed strongly on the former. It was on the latter that human authorship retained a meaningful advantage under disclosure. Within the scope of this study, this framework offers a more precise account of user engagement with GenAI-authored landing page content than a simple acceptance-versus-rejection model allows. However, whether it extends to other digital marketing formats, content types, or platforms remains to be tested. Responses were further moderated by individual familiarity with GenAI, attitudes towards its use in marketing, and a rapidly shifting baseline of exposure that is gradually normalising GenAI content as part of everyday digital experience.

Chapter 6 Conclusions

6.1 Chapter overview

This thesis examined how users evaluate GenAI-authored and human-written digital marketing content, with particular attention to the role of authorship disclosure. Using a mixed methods survey design, the study compared participants' responses to two website landing pages in a controlled experimental setting: one written by GenAI and one by a human digital content specialist. Participants evaluated the pages under either undisclosed or disclosed authorship conditions, allowing the study to investigate both preference formation and the effect of source awareness on those preferences.

This chapter brings the thesis to a close by drawing together the key findings, contributions, and implications of the study. It begins by summarising the key patterns observed across both experimental conditions in response to the two research questions guiding this study.

Research question 1: Do users prefer GenAI-authored or human-authored website landing page content when authorship is not disclosed?

Research question 2: How does authorship disclosure influence users' preferences for GenAI-authored and human-authored website landing page content?

It then considers what these findings contribute to theoretical and practical debates around GenAI-authored content, authorship disclosure, and user trust. The chapter proceeds to examine the practical implications for marketers, before acknowledging the limitations of the study and the directions these open for future research.

6.2 Summary of key findings

The findings indicate that user evaluations of GenAI-authored and human-authored marketing content were broadly positive across both conditions. However, the basis of those evaluations shifted meaningfully depending on whether authorship was disclosed. Without authorship context, the GenAI-authored landing page performed comparably to, and in several respects marginally better than, the human-authored page. Addressing RQ1, these results suggest that users favour GenAI-authored content when source information is absent. Rather than presuming authorship, participants appeared to evaluate content on functional grounds: how clearly it communicated, how well it was structured, and how effectively it supported decision-making in a goal-oriented context of a website landing page.

When authorship was disclosed, the human-authored page gained a clearer advantage on relational dimensions such as authenticity, human-likeness, and brand genuineness. In contrast, the GenAI-authored page attracted more guarded and less consistently positive responses. Addressing RQ2, authorship disclosure influenced user preferences not by reversing judgements entirely, but by redistributing evaluative emphasis. It left competence-based assessments largely intact while heightening sensitivity to the perceived presence or absence of human authorship. Qualitative responses reinforced this pattern across both conditions, where participants consistently valued clarity, structure, and usability, but associated trust and reassurance with visible human involvement. Taken together, these findings show that participants did not evaluate GenAI-authored content uniformly positively or negatively. Instead, their responses depended on whether they were judging the content's usefulness or its perceived human quality. This dynamic has important implications for how generative AI is deployed in digital marketing practice and for how well marketers balance the need for source transparency against content effectiveness.

6.3 Contributions to knowledge and theory

This study contributes to the growing body of literature on GenAI-authored content in several interconnected ways. The most substantive theoretical contribution is the articulation of a dual evaluative framework through which users appear to assess GenAI-authored content. Rather than supporting a straightforward model of algorithm aversion, the findings reveal that AI authorship does not produce negative judgements. Instead, users distinguish between competence-based dimensions and relational dimensions. This framework suggests that user responses to GenAI-authored content tend to be shaped not only by source-based biases but also by the functional expectations associated with the communicative environment. This is a distinction that existing models of algorithm aversion do not fully accommodate.

Second, the study advances current understanding of disclosure effects by demonstrating that authorship disclosure does not uniformly reduce evaluations of GenAI-authored content. Instead, awareness of the source of digital content heightens sensitivity to relational concerns such as authenticity and human connection while leaving many competence-based judgements relatively intact. This contributes a more nuanced account of disclosure effects than models that treat transparency as straightforwardly harmful to trust or acceptance. In doing so, it reinforces the context-dependent nature of user responses. It illustrates that the interaction between content characteristics, disclosure cues, and the broader communicative setting shapes perceptions of GenAI-authored content. By situating the analysis within website landing

pages as a functional, conversion-oriented environment, the study also contributes a context-specific empirical insight. It moves the discussion beyond journalism, social media, and general communication settings, where much of the existing literature has been concentrated, towards a commercially applied digital marketing context.

Finally, the study makes a methodological contribution by integrating multiple evaluative approaches, including independent ratings, comparative judgements, and qualitative explanations of participants' reasoning. The divergence observed across these measures reinforces the argument that user perceptions of GenAI-authored content are multifaceted and cannot be fully captured through a single methodological lens. The controlled comparison between disclosed and undisclosed conditions further allowed the study to identify not only patterns of preference but also the reasoning underlying them, offering a fuller account of how users interpret GenAI-authored marketing content under different informational conditions.

6.4 Implications for policy and practice

The findings carry meaningful implications for both marketing practice and the broader governance of GenAI-authored communication. The most immediate practical implication is that generative AI can be used effectively in structured, goal-oriented communication environments. Marketers operating in similar performance-driven formats, including email campaigns, product category pages, onboarding sequences, help centre content, and landing pages, are well-positioned to benefit from its deployment at scale. Moreover, human oversight is needed in sectors where perceived authenticity and human accountability are core expectations, such as financial services providers, healthcare organisations, legal and professional services firms, and mission-driven entities, including charities and advocacy groups. For such organisations, deploying GenAI-authored content without adequate disclosure or visible human oversight can pose a major reputational risk.

Regarding disclosure, the findings suggest it should be treated as an ethical imperative rather than a discretionary choice. Disclosure strategies should be calibrated to the type of communication, the audience's expectations, and the extent to which relational trust is central to the interaction. The absence of transparency, by contrast, risks undermining perceptions of fairness and accountability in ways that may prove more damaging in the long term. Taken together, the findings point towards the use of hybrid content production models, in which organisations can benefit from GenAI's capacity to produce content at scale while still maintaining signs of human review.

6.5 Limitations of the study

Several limitations should be considered when interpreting the findings of this study. First, the study was conducted in a controlled survey environment rather than a natural browsing context. This design was appropriate for isolating the effect of authorship disclosure, but it also simplified the conditions under which users normally evaluate digital marketing content. Real-world website judgements are shaped by a wider set of cues, including brand familiarity, visual design, reviews, endorsements, prior expectations, search intent, and competing information sources (Metzger et al., 2010). The findings, therefore, speak most directly to content-level evaluations, rather than to the full range of contextual factors that influence digital decision-making in practice.

Moreover, only one landing page was created for each authorship condition. This constrains the range of content evaluated and limits generalisability across product categories and writing styles. The human-authored page was written by the researcher rather than an independent copywriter. This introduces a risk of researcher bias in content design, despite efforts to align tone and structure across conditions. The researcher's dual role as content creator and study designer cannot be fully separated from the interpretation of results. Future research should draw stimuli from multiple GenAI systems and independent human authors across several items, to strengthen robustness and generalisability.

Second, the study relied on a relatively small sample of 95 participants recruited through an online platform. While this sample was sufficient for identifying descriptive patterns and comparing responses across the two experimental conditions, it limits the extent to which the findings can be generalised to broader populations of internet users (Cao et al., 2024). The participant group also appeared relatively digitally literate, with nearly half reporting frequent use of generative AI. Prior familiarity with these tools may have shaped participants' responses to GenAI-authored content (Li et al., 2025). Although AI familiarity appeared broadly distributed across the sample, the findings should still be interpreted with caution when considering less digitally experienced audiences.

Moreover, participants were not required to identify as first-home buyers, despite the stimulus content targeting this audience. Landing pages addressing property decisions typically presume personal relevance or financial stake. Participants without this relevance may have evaluated the content on generic persuasive qualities rather than personal applicability. This limitation may have understated disclosure effects among genuine first-home buyers, for whom the financial stakes are higher. Future research

should recruit participants matching the stimulus's target audience, to test whether personal relevance moderates the effect of disclosure.

Third, the study measured self-reported evaluations rather than observed behavioural outcomes. Participants indicated which landing page they perceived as more credible, persuasive, authentic, or preferable. However, the study did not measure whether these perceptions would translate into actual behaviours such as clicking, signing up, requesting more information, or making a purchase. As a result, the findings provide insight into perceived content effectiveness, but not into conversion behaviour in real marketing environments.

Fourth, the disclosure manipulation was intentionally simple. Authorship labels were presented to participants in the survey as “*Webpage A was written by a human*” or “*Webpage B was written by ChatGPT*” (See Appendix E: Survey), and they received no further details about the content production process or editorial oversight. This distinction created a clear experimental comparison, but it may not reflect the more complex ways GenAI is used in practice. In real marketing workflows, content may be drafted by AI, edited by humans, approved by specialists, or produced through several stages of human-AI collaboration. Participants’ responses to a simple authorship label may therefore differ from their responses to more detailed or contextualised disclosure formats.

A further limitation concerns reliance on GPT-4 as the sole GenAI system used to produce experimental stimuli. Findings may not generalise to other GenAI models, which vary in tone, structure, and perceived output quality. The prompt used to generate the GenAI-authored page was also brief, providing minimal guidance on tone, structure, or persuasive strategy. Prompt specificity is known to shape output quality and style in large language models, and a more detailed or iteratively refined prompt may have produced content closer to professional copywriting standards (White et al., 2023). The brevity of the prompt therefore represents a lower bound on GenAI performance rather than its full potential. This may have disadvantaged the GenAI condition relative to the human-authored page, which benefited from the researcher's contextual judgement and revision. Future research should test multiple prompting strategies, including iterative refinement, to assess how prompt engineering affects perceptions of GenAI-authored content.

Finally, the study focused on a single content format and scenario: textual website landing pages in a first-home buyer context. This limits the extent to which the findings can be applied to other forms of digital marketing communication, particularly formats that are more emotional, creative, visual, interactive, or high-stakes. The findings

should therefore be understood as context-specific rather than as a general account of user responses to all GenAI-authored marketing content. They also reflect user perceptions during a period when GenAI is still becoming normalised. As these tools become more familiar and more deeply embedded in everyday digital environments, users' expectations of AI-authored communication may continue to change.

6.6 Directions for future research

The findings of this study point to several directions that future research could productively pursue. A first and perhaps most immediate direction is to replicate this work in more valid digital environments. Future studies could test GenAI-authored and human-authored content within more comprehensive website environments that include visual design, brand identity, testimonials, reviews, navigation features, and other contextual cues that users typically rely on when making judgements online. Moving beyond self-reported evaluations towards behavioural measures such as click-through rates, time on page, enquiry submissions, sign-up behaviour, or purchase decisions would also help determine whether stated preferences align with actual user behaviour.

A second direction is to investigate authorship disclosure as a design variable rather than as a simple binary label. Future research could compare different disclosure formats, including brief authorship labels, detailed process explanations, human-in-the-loop statements, editorial review notices, and graduated disclosures that specify the degree of AI involvement. This would help identify which forms of transparency support informed user judgement without unnecessarily reducing trust. It would also allow researchers to examine whether users respond differently to fully GenAI-authored content, human-edited AI content, and content produced through collaborative human-AI workflows.

Third, future studies should examine how responses to GenAI-authored content vary across industries and communication goals. The present study focused on a functional landing page in the first-home-buyer context. However, user expectations may differ across sectors such as healthcare, finance, education, legal services, charities, tourism, and entertainment. Similarly, reactions may vary across content types, including social media posts, email campaigns, product descriptions, chatbot interactions, video scripts, brand storytelling, and visual advertising. Comparing these settings would help clarify when users prioritise efficiency and clarity, and when they place greater weight on authenticity, empathy, creativity, or human accountability.

Finally, longitudinal research is needed to understand how user perceptions evolve as generative AI becomes more familiar. Current responses may still be influenced by novelty, uncertainty, or emerging social norms around AI use. Studies that track users over time or repeatedly expose participants to GenAI-authored marketing content could examine whether preferences for human authorship persist, weaken, or become more context-dependent. This would be particularly valuable for understanding whether concerns about authenticity and disclosure remain stable or shift as GenAI-authored communication becomes a routine part of digital marketing practice.

6.7 Concluding reflection

This thesis set out to examine whether users prefer GenAI-authored or human-authored marketing content, and whether disclosure changes those preferences. The findings show that users are neither simply sceptical of GenAI nor uncritically accepting of it. Rather, judgements operate across two distinct dimensions: functional competence and relational authenticity. While users value GenAI content on largely functional grounds, disclosure activates relational evaluations, shifting user judgements in ways that purely competence-based models cannot explain.

The central implication is that disclosure is not just a transparency obligation, but a credibility signal that reshapes how content is viewed. As GenAI becomes embedded in digital marketing, organisations will need to show not only when AI has been used, but also how human judgement remains involved. Trust in GenAI-authored communication will depend less on whether content is artificial or human, and more on whether its production feels transparent, accountable, and appropriate to the context.

References

- Acatrinei, C., & Miu, C. M. (2025). AI vs. human-written email marketing: Consumer insights for sustainable communication. *Amfiteatru Economic*, 27(S19), 1256–1271. <https://doi.org/10.24818/ea/2025/s19/1256>
- Al-kfairy, M., Mustafa, D., Kshetri, N., Insiew, M., & Alfandi, O. (2024). Ethical challenges and solutions of generative AI: An interdisciplinary perspective. *Informatics*, 11(3), Article 58. <https://doi.org/10.3390/informatics11030058>
- Almaghlouth, A. (2024). Impact of the qualified digital content on digital experience optimization and conversion. *European Journal of Computer Science and Information Technology*, 12(5), 1–16. <https://doi.org/10.37745/ejcsit.2013/vol12n5116>
- Altay, S., & Gilardi, F. (2024). People are skeptical of headlines labeled as AI-generated, even if true or human-made, because they assume full AI automation. *PNAS Nexus*, 3(10). <https://doi.org/10.1093/pnasnexus/pgae403>
- Appelman, A., & Sundar, S. S. (2016). Measuring message credibility: Construction and validation of an exclusive scale. *Journalism & Mass Communication Quarterly*, 93(1), 59–79. <https://doi.org/10.1177/1077699015606057>
- Arango, L., Singaraju, S. P., & Niininen, O. (2023). Consumer responses to AI-generated charitable giving ads. *Journal of Advertising*, 52(4), 486–503. <https://doi.org/10.1080/00913367.2023.2183285>
- Bandi, A., Adapa, P. V. S. R., & Kuchi, Y. E. V. P. K. (2023). The power of generative AI: A review of requirements, models, input–output formats,

evaluation metrics, and challenges. *Future Internet*, 15(8), Article 260.

<https://doi.org/10.3390/fi15080260>

Banh, L., & Strobel, G. (2023). Generative artificial intelligence. *Electronic Markets*, 33(1). <https://doi.org/10.1007/s12525-023-00680-1>

Benges, S., El-Sayed, H., Sarker, M. K., Houkpati, Y., Irungu, J., & Oladunni, T. (2023). Advancements in generative AI: A comprehensive review of GANs, GPT, autoencoders, diffusion model, and transformers. *IEEE Access*, 12, 69812–69837. <https://doi.org/10.1109/access.2024.3397775>

Bogert, E., Lauharatanahirun, N., & Schechter, A. (2022). Human preferences toward algorithmic advice in a word association task. *Scientific Reports*, 12(1), Article 14501. <https://doi.org/10.1038/s41598-022-18638-2>

Bordas, A., Le Masson, P., Thomas, M., & Weil, B. (2024). What is generative in generative artificial intelligence? A design-based perspective. *Research in Engineering Design*, 35(4), 427–443. <https://doi.org/10.1007/s00163-024-00441-x>

Braun, V., & Clarke, V. (2019). Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4), 589–597. <https://doi.org/10.1080/2159676x.2019.1628806>

Braun, V., & Clarke, V. (n.d.). *Doing reflexive TA*. Thematic Analysis. Retrieved May 25, 2026, from <https://www.thematicanalysis.net/doing-reflexive-ta/>

Brüns, J. D., & Meißner, M. (2024). Do you create your content yourself? Using generative artificial intelligence for social media content creation diminishes perceived brand authenticity. *Journal of Retailing and Consumer Services*, 79, Article 103790. <https://doi.org/10.1016/j.jretconser.2024.103790>

- Bubphapant, J., & Brandão, A. (2024). Content marketing research: A review and research agenda. *International Journal of Consumer Studies*, 48(1), e12984. <https://doi.org/10.1111/ijcs.12984>
- Buchanan, J., & Hickman, W. (2024). *Do people trust humans more than ChatGPT?* *Journal of Behavioral and Experimental Economics*, 112, 102239. <https://doi.org/10.1016/j.socec.2024.102239>
- Burton, J. W., Stein, M. K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239. <https://doi.org/10.1002/bdm.2155>
- Byrne, D. (2022). A worked example of Braun and Clarke's approach to reflexive thematic analysis. *Quality & Quantity*, 56(3), 1391–1412. <https://doi.org/10.1007/s11135-021-01182-y>
- Camilleri, M. A. (2024). Factors affecting performance expectancy and intentions to use ChatGPT: Using SmartPLS to advance an information technology acceptance framework. *Technological Forecasting and Social Change*, 201, Article 123247. <https://doi.org/10.1016/j.techfore.2024.123247>
- Cao, Y., Chen, R. C., & Katz, A. J. (2024). *Why is a small sample size not enough?* *The Oncologist*, 29(9), 761–763. <https://doi.org/10.1093/oncolo/oyae162>
- Casal, J. E., & Kessler, M. (2023). Can linguists distinguish between ChatGPT/AI and human writing? A study of research ethics and academic publishing. *Research Methods in Applied Linguistics*, 2(3), Article 100068. <https://doi.org/10.1016/j.rmal.2023.100068>

- Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809–825.
<https://doi.org/10.1177/0022243719851788>
- Çetinkaya, Y. M., Külah, E., Toroslu, İ. H., & Davulcu, H. (2024). Targeted marketing on social media: Utilizing text analysis to create personalized landing pages. *Social Network Analysis and Mining*, 14(1), Article 77.
<https://doi.org/10.1007/s13278-024-01213-0>
- Cheah, C. W., Koay, K. Y., & Lim, W. M. (2024). Social media influencer over-endorsement: Implications from a moderated-mediation analysis. *Journal of Retailing and Consumer Services*, 79, Article 103831.
<https://doi.org/10.1016/j.jretconser.2024.103831>
- Cheng, A., Calhoun, A., & Reedy, G. (2025). Artificial intelligence-assisted academic writing: Recommendations for ethical use. *Advances in Simulation*, 10(1), Article 22. <https://doi.org/10.1186/s41077-025-00350-6>
- Choudhury, A., & Shamszare, H. (2023). Investigating the impact of user trust on the adoption and use of ChatGPT: Survey analysis. *Journal of Medical Internet Research*, 25, Article e47184. <https://doi.org/10.2196/47184>
- CloudResearch. (n.d.). *Connect for researchers: Vetted participants on-demand*. Retrieved May 25, 2026, from <https://www.cloudresearch.com/products/connect-for-researchers/>
- Corrêa, N. K., Galvão, C., Santos, J. W., Del Pino, C., Pinto, E. P., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., & de Oliveira, N. (2024). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns*, 4(10), Article 100857.
<https://doi.org/10.1016/j.patter.2023.100857>

- Creswell, J. W., & Clark, V. L. P. (2017). Choosing a mixed method design. *Designing and Conducting Mixed Methods Research*, 53–106.
- Curina, I., Francioni, B., Hegner, S. M., & Cioppi, M. (2020). Brand hate and non-repurchase intention: A service context perspective in a cross-channel setting. *Journal of Retailing and Consumer Services*, 54, Article 102031. <https://doi.org/10.1016/j.jretconser.2019.102031>
- Dabbous, A., & Barakat, K. A. (2020). Bridging the online offline gap: Assessing the impact of brands' social network content quality on brand awareness and purchase intention. *Journal of Retailing and Consumer Services*, 53, Article 101966. <https://doi.org/10.1016/j.jretconser.2019.101966>
- Darda, K., Carre, M., & Cross, E. (2023). Value attributed to text-based archives generated by artificial intelligence. *Royal Society Open Science*, 10(2), Article 220915. <https://doi.org/10.1098/rsos.220915>
- Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- Davidavičius, S., & Limba, T. (2022). Recognition of digital content needs for inbound marketing solutions. *Social Sciences*, 11(8), Article 351. <https://doi.org/10.3390/socsci11080351>
- Dawadi, S., Shrestha, S., & Giri, R. A. (2021). Mixed-methods research: A discussion on its types, challenges, and criticisms. *Journal of Practical Studies in Education*, 2(2), 25–36. <https://doi.org/10.46809/jpse.v2i2.20>
- Desai, V. (2019). Digital marketing: A review. *International Journal of Trend in Scientific Research and Development*, Special Issue(Special Issue-FIIIPM2019), 196–200. <https://doi.org/10.31142/ijtsrd23100>

- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126.
<https://doi.org/10.1037/xge0000033>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), Article eadn5290.
<https://doi.org/10.1126/sciadv.adn5290>
- Downie, N., & Starry, A. R. (2019). Descriptive and inferential statistics. In A. M. Colman (Ed.), *Companion encyclopedia of psychology* (Vol. 1, pp. 55–95). Routledge. <https://doi.org/10.4324/9781315542096>
- Dvir, N., & Gafni, R. (2018). How content volume on landing pages influences consumer behavior. *Proceedings of the Informing Science and IT Education Conference*, pp. 35–53. <https://doi.org/10.28945/4016>
- Dwiphalala, F., & Astuti, R. D. (2024). The role of source credibility and fairness on the formation of parasocial relationships and product interest and their influence on purchase intention and e-WOM of beauty brands in Indonesia. *Eduvest - Journal of Universal Studies*, 4(8), 7056–7073.
<https://doi.org/10.59188/eduvest.v4i8.1770>
- Ermüt, S. (2026, April 20). LLM market share: Compare usage & adoption. AIMultiple. Retrieved 20 May, 2026, from <https://aimultiple.com/llm-market-share>
- Evans, J. R., & Mathur, A. (2018). The value of online surveys: A look back and a look ahead. *Internet Research*, 28(4), 854–887.
<https://doi.org/10.1108/intr-03-2018-0089>

- Feilzer, M. Y. (2010). Doing mixed methods research pragmatically: Implications for the rediscovery of pragmatism as a research paradigm. *Journal of Mixed Methods Research*, 4(1), 6–16.
<https://doi.org/10.1177/1558689809349691>
- Flanagin, A. J., & Metzger, M. J. (2007). The role of site features, user attributes, and information verification behaviors on the perceived credibility of web-based information. *New Media & Society*, 9(2), 319–342. <https://doi.org/10.1177/1461444807075015>
- Fogg, B. J., Marshall, J., Laraki, O., Osipovich, A., Varma, C., Fang, N., Paul, J., Rangnekar, A., Shon, J., Swani, P., & Treinen, M. (2001). What makes web sites credible? A report on a large quantitative study. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 61–68. <https://doi.org/10.1145/365024.365037>
- Grewal, D., Saturnino, C. B., Davenport, T., & Guha, A. (2025). How generative AI is shaping the future of marketing. *Journal of the Academy of Marketing Science*, 53(3), 702–722. <https://doi.org/10.1007/s11747-024-01064-3>
- Gu, C., Jia, S., Lai, J., Chen, R., & Chang, X. (2024). Exploring consumer acceptance of AI-generated advertisements: From the perspectives of perceived eeriness and perceived intelligence. *Journal of Theoretical and Applied Electronic Commerce Research*, 19(3), 2218–2238.
<https://doi.org/10.3390/jtaer19030108>
- Guzman, A. L., & Lewis, S. C. (2024). What generative AI means for the media industries, and why it matters to study the collective consequences for advertising, journalism, and public relations. *Emerging Media*, 2(3), 347–355. <https://doi.org/10.1177/27523543241289239>

- Hartmann, J., Exner, Y., & Domdey, S. (2025). *The power of generative marketing: Can generative AI create superhuman visual marketing content? International Journal of Research in Marketing*, 42(1), 13–31. <https://doi.org/10.1016/j.ijresmar.2024.09.002>
- Haupt, M., Freidank, J., & Haas, A. (2025). Consumer responses to human-AI collaboration at organizational frontlines: Strategies to escape algorithm aversion in content creation. *Review of Managerial Science*, 19(2), 377–413. <https://doi.org/10.1007/s11846-024-00748-y>
- Hazra, A., & Gogtay, N. (2016). Biostatistics series module 3: Comparing groups: Numerical variables. *Indian Journal of Dermatology*, 61(3), 251–260. <https://doi.org/10.4103/0019-5154.182416>
- Hollebeek, L. D., & Macky, K. (2019). Digital content marketing's role in fostering consumer engagement, trust, and value: Framework, fundamental propositions, and implications. *Journal of Interactive Marketing*, 45, 27–41. <https://doi.org/10.1016/j.intmar.2018.07.003>
- Hovland, C. I., Janis, I. L., & Kelley, H. H. (1953). *Communication and persuasion: Psychological studies of opinion change*. Yale University Press.
- Huang, D., Zhu, Y., & Mustafaraj, E. (2019). *How dependable are “first impressions” to distinguish between real and fake news websites? HT 2019 - Proceedings of the 30th ACM Conference on Hypertext and Social Media*, 201–210. <https://doi.org/10.1145/3342220.3343670>
- Hussain, A. (2024). The impact of artificial intelligence on digital media content creation. *International Journal of Innovative Science and Research Technology (IJISRT)*, 998–1003. <https://doi.org/10.38124/ijisrt/ijisrt24jul927>

- Ismagilova, E., Slade, E., Rana, N. P., & Dwivedi, Y. K. (2020). The effect of characteristics of source credibility on consumer behaviour: A meta-analysis. *Journal of Retailing and Consumer Services*, 53, Article 101736. <https://doi.org/10.1016/j.jretconser.2019.01.005>
- Jacob, M. E., & Johnson, J. (2021). Conceptualising digital content marketing for greater consumer brand engagement. *Colombo Business Journal*, 12(2), 80–102. <https://doi.org/10.4038/cbj.v12i2.83>
- Jain, G., Pareek, S., & Carlbring, P. (2024). Revealing the source: How awareness alters perceptions of AI and human-generated mental health responses. *Internet Interventions*, 36, Article 100745. <https://doi.org/10.1016/j.invent.2024.100745>
- Jebb, A. T., Parrigon, S., & Woo, S. E. (2017). Exploratory data analysis as a foundation of inductive research. *Human Resource Management Review*, 27(2), 265–276. <https://doi.org/10.1016/j.hrmr.2016.08.003>
- Jenkins, E. L., Ilicic, J., Barklamb, A. M., & McCaffrey, T. A. (2020). Assessing the credibility and authenticity of social media content for applications in health communication: Scoping review. *Journal of Medical Internet Research*, 22(7). <https://doi.org/10.2196/17296>
- Jia, M., Zhao, Y., & Zhang, X. (2025). Will potential threats persist? Attitudes and behavioral responses of creative professionals toward GenAI. *Information Research*, 30(iConf), 1073–1083. <https://doi.org/10.47989/ir30iconf47098>
- Joan Isibor, N., Attipoe, V., Oyeyipo, I., Christiana Ayodeji, D., Apiyo Mayienga, B., Alonge, E., & Clement Onwuzulike, O. (2025). Analyzing successful content marketing strategies that enhance online engagement and sales for digital brands. *International Journal of Advanced Multidisciplinary*

Research and Studies, 5(2), 842–851.

<https://doi.org/10.62225/2583049x.2025.5.2.3906>

Johnson, D., & Grayson, K. (2005). Cognitive and affective trust in service relationships. *Journal of Business Research*, 58(4), 500–507.

[https://doi.org/10.1016/s0148-2963\(03\)00140-1](https://doi.org/10.1016/s0148-2963(03)00140-1)

Johnson, R. B., & Onwuegbuzie, A. J. (2004). Mixed methods research: A research paradigm whose time has come. *Educational Researcher*, 33(7), 14–26. <https://doi.org/10.3102/0013189X033007014>

Joshi, A., Kale, S., Chandel, S., & Pal, D. (2015). Likert scale: Explored and explained. *British Journal of Applied Science & Technology*, 7(4), 396–403. <https://doi.org/10.9734/bjast/2015/14975>

Jung, J., Song, H., Kim, Y., Im, H., & Oh, S. (2017). Intrusion of software robots into journalism: The public's and journalists' perceptions of news written by algorithms and human journalists. *Computers in Human Behavior*, 71, 291–298. <https://doi.org/10.1016/j.chb.2017.02.022>

Kalota, F. (2024). A primer on generative artificial intelligence. *Education Sciences*, 14(2), Article 172. <https://doi.org/10.3390/educsci14020172>

Kashive, N., Powale, L., & Kashive, K. (2020). Understanding user perception toward artificial intelligence (AI) enabled e-learning. *International Journal of Information and Learning Technology*, 38(1), 1–19.

<https://doi.org/10.1108/ijilt-05-2020-0090>

Kaufmann, E., Chacon, A., Kausel, E. E., Herrera, N., & Reyes, T. (2023). Task-specific algorithm advice acceptance: A review and directions for future research. *Data and Information Management*, 7(3), 100040.

<https://doi.org/10.1016/j.dim.2023.100040>

- Kirk, C. P., & Givi, J. (2025). The AI-authorship effect: Understanding authenticity, moral disgust, and consumer responses to AI-generated marketing communications. *Journal of Business Research*, 186, 114984. <https://doi.org/10.1016/j.jbusres.2024.114984>
- Kirkby, A., Baumgarth, C., & Henseler, J. (2023). To disclose or not disclose, is no longer the question – effect of AI-disclosed brand voice on brand authenticity and attitude. *Journal of Product and Brand Management*, 32(7), 1108–1122. <https://doi.org/10.1108/jpbm-02-2022-3864>
- Kivunja, C., & Kuyini, A. B. (2017). Understanding and applying research paradigms in educational contexts. *International Journal of Higher Education*, 6(5), 26–41. <https://doi.org/10.5430/ijhe.v6n5p26>
- Koob, C. (2021). Determinants of content marketing effectiveness: Conceptual framework and empirical findings from a managerial perspective. *PLOS ONE*, 16(4), Article e0249457. <https://doi.org/10.1371/journal.pone.0249457>
- Kothari, H., Choudhary, A., Jain, A., Singh, S., Prasad, K. D. V., & Vani, U. K. (2025). Impact of social media advertising on consumer behavior: Role of credibility, perceived authenticity, and sustainability. *Frontiers in Communication*, 10, Article 1595796. <https://doi.org/10.3389/fcomm.2025.1595796>
- Kotler, P. (1972). A generic concept of marketing. *Journal of Marketing*, 7(3), 48. <https://doi.org/10.1177/002224297203600209>
- Kreps, S., McCain, R. M., & Brundage, M. (2022). All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. *Journal of Experimental Political Science*, 9(1), 104–117. <https://doi.org/10.1017/xps.2020.37>

- Krosnick, J. A. (1991). Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology*, 5(3), 213–236. <https://doi.org/10.1002/acp.2350050305>
- Kshetri, N., Dwivedi, Y. K., Davenport, T. H., & Panteli, N. (2024). Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda. *International Journal of Information Management*, 75, Article 102716. <https://doi.org/10.1016/j.ijinfomgt.2023.102716>
- Kubovics, M. (2025a). Influencing the perception and acceptance of AI-generated marketing content as part of SGD. *Journal of Lifestyle and SDGs Review*, 5(6), e06929. <https://doi.org/10.47172/2965-730x.sdgsreview.v5.n06.pe06929>
- Kubovics, M. (2025b). Willingness of users to accept AI content creation. *European Business & Management*, 11(2), 40–47. <https://doi.org/10.11648/j.ebm.20251102.12>
- Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, 8(1), Article 33267. <https://doi.org/10.1525/collabra.33267>
- Ledolter, J., Gramlich, O. W., & Kardon, R. H. (2020). Parametric statistical inference for comparing means and variances. *Investigative Ophthalmology & Visual Science*, 61(8), Article 25. <https://doi.org/10.1167/iovs.61.8.25>
- Lermann Henestrosa, A., & Kimmerle, J. (2024). The effects of assumed AI vs. human authorship on the perception of a GPT-generated text. *Journalism and Media*, 5(3), 1085–1097. <https://doi.org/10.3390/journalmedia5030069>

- Li, F., Yang, Y., & Yu, G. (2025). Nudging perceived credibility: The impact of AIGC labeling on user distinction of AI-generated content. *Emerging Media*, 3(2), 275–304. <https://doi.org/10.1177/27523543251317572>
- Liu-Thompkins, Y., Okazaki, S., & Li, H. (2022). Artificial empathy in marketing interactions: Bridging the human-AI gap in affective and social customer experience. *Journal of the Academy of Marketing Science*, 50(6), 1198–1218. <https://doi.org/10.1007/s11747-022-00892-5>
- Logg, J., Minson, J., & Moore, D. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629–650. <https://doi.org/10.1093/jcr/ucz013>
- Lou, C., & Yuan, S. (2019). Influencer marketing: How message value and credibility affect consumer trust of branded content on social media. *Journal of Interactive Advertising*, 19(1), 58–73. <https://doi.org/10.1080/15252019.2018.1533501>
- Manyara, A. M., Purvis, A., Ciani, O., Collins, G. S., & Taylor, R. S. (2024). *Sample size in multistakeholder Delphi surveys: at what minimum sample size do replicability of results stabilize?* *Journal of Clinical Epidemiology*, 174. <https://doi.org/10.1016/j.jclinepi.2024.111485>
- Mathews, B. (2023). Adolescent capacity to consent to participate in research: A review and analysis informed by law, human rights, ethics, and developmental science. *Laws*, 12(1), Article 2. <https://doi.org/10.3390/laws12010002>

Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., & Cerf, M.

(2024). The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14, Article 4692.

<https://doi.org/10.1038/s41598-024-53755-0>

Meguelliati, E., Civelli, S., Han, L., Bernstein, A., Sadiq, S., & Demartini, G.

(2025). *LLM-generated ads: From personalization parity to persuasion superiority*. arXiv. <https://arxiv.org/abs/2512.03373>

Meier, S. J., & Song, H. (2026). Question order effects in multidimensional risk perception measurement. *Risk Analysis*, 46(1), e70164.

<https://doi.org/10.1111/risa.70164>

Menon, V., & Muraleedharan, A. (2020). Internet-based surveys: Relevance, methodological considerations and troubleshooting strategies. *General Psychiatry*, 33(5), e100264. <https://doi.org/10.1136/gpsych-2020-100264>

Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and heuristic approaches to credibility evaluation online. *Journal of Communication*, 60(3), 413–439. <https://doi.org/10.1111/j.1460-2466.2010.01488.x>

Metzler, H., & Garcia, D. (2024). Social drivers and algorithmic mechanisms on digital media. *Perspectives on Psychological Science*, 19(5), 735–748. <https://doi.org/10.1177/17456916231185057>

Moore, D. W. (2002). Measuring new types of question-order effects: Additive and subtractive. *Public Opinion Quarterly*, 66(1), 80–91. <https://doi.org/10.1086/338631>

Morales, N. M. (2025). Adoption and usability of artificial intelligence in marketing content creation. *Journal of Posthumanism*, 5(3), 1107–1118. <https://doi.org/10.63332/joph.v5i3.852>

- Moran, G., & Muzellec, L. (2017). eWOM credibility on social networking sites: A framework. *Journal of Marketing Communications*, 23(2), 149–161.
<https://doi.org/10.1080/13527266.2014.969756>
- Munteanu, T., Kannis-Dymand, L., Millear, P., & Jones, M. (2024). Analysis of cognitive behavioural therapy apps for generalised anxiety disorder: Evidence-based content and user experience. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 18(4).
<https://doi.org/10.5817/cp2024-4-6>
- Nakano, H., Takezawa, J., Matulic, F., Yang, C.-L., & Yatani, K. (2025). Understanding reader perception shifts upon disclosure of AI authorship. *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3742413.3789076>
- Ngo, T. T. A., Bui, C. T., Chau, H. K. L., & Tran, N. P. N. (2024). Electronic word-of-mouth (eWOM) on social networking sites (SNS): Roles of information credibility in shaping online purchase intention. *Heliyon*, 10(11).
<https://doi.org/10.1016/j.heliyon.2024.e32168>
- Nielsen, J., & Pernice, K. (2009). *Eye tracking techniques for web usability*. Nielsen Norman Group.
- Odgaard-Jensen, J., Vist, G. E., Timmer, A., Kunz, R., Akl, E. A., Schünemann, H., Briel, M., Nordmann, A. J., Pregno, S., & Oxman, A. D. (2011). Randomisation to protect against selection bias in healthcare trials. *Cochrane Database of Systematic Reviews*, 2011(4).
<https://doi.org/10.1002/14651858.mr000012.pub3>
- Ogah, A. I., & Abutu, D. O. (2022). Theoretical analysis on persuasive communication in advertising and its application in marketing

- communication. *EJOTMAS: Ekpoma Journal of Theatre and Media Arts*, 8(1–2), 313–331. <https://doi.org/10.4314/ejotmas.v8i1-2.17>
- Ohanian, R. (1990). Construction and validation of a scale to measure celebrity endorsers' perceived expertise, trustworthiness, and attractiveness. *Journal of Advertising*, 19(3), 39–52.
- Okoso, A., Yang, M., & Baba, Y. (2025). Do expressions change decisions? Exploring the impact of AI's explanation tone on decision-making. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3706598.3713744>
- OpenAI. (n.d.). *GPT-4*. Retrieved May 25, 2026, from <https://openai.com/index/gpt-4-research/>
- Osborne, M. R., & Bailey, E. R. (2025). Me vs. the machine? Subjective evaluations of human- and AI-generated advice. *Scientific Reports*, 15, Article 3980. <https://doi.org/10.1038/s41598-025-86623-6>
- Ou, M., & Ho, S. S. (2024). Factors associated with information credibility perceptions: A meta-analysis. *Journalism and Mass Communication Quarterly*, 101(2), 346–372. <https://doi.org/10.1177/10776990231222556>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Aspell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35. <https://arxiv.org/abs/2203.02155>
- Pallant, J. (2020). *SPSS survival manual: A step by step guide to data analysis using IBM SPSS (7th ed.)*. Routledge. <https://doi.org/10.4324/9781003117452>

- Park, J., Oh, C., & Kim, H. Y. (2024). AI vs. human-generated content and accounts on instagram: User preferences, evaluations, and ethical considerations. *Technology in Society*, 79, Article 102705. <https://doi.org/10.1016/j.techsoc.2024.102705>
- Parshakov, P., Naidenova, I., Paklina, S., Matkin, N., & Nesseler, C. (2025). *Users favor LLM-generated content—until they know it's AI*. arXiv. <https://arxiv.org/abs/2503.16458>
- Pellas, N. (2023). The effects of generative AI platforms on undergraduates' narrative intelligence and writing self-efficacy. *Education Sciences*, 13(11), Article 1155. <https://doi.org/10.3390/educsci13111155>
- Radivojevic, K., Chou, M., Badillo-Urquiola, K., & Brenner, P. (2024). *Human perception of LLM-generated text content in social media environments*. arXiv. <https://arxiv.org/abs/2409.06653>
- Rae, I. (2024, May 11). The effects of perceived AI use on content perceptions. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3613904.3642076>
- Rashmi, & Adhana, D. K. (2023). Content marketing in B2C: Unraveling the key drivers of effectiveness. *Asian Journal of Management and Commerce*, 4(2), 243–250. <https://doi.org/10.22271/27084515.2023.v4.i2c.216>
- Raso, J. (2024). *Responsible AI: Binaries that bind*. *McGill Law Journal*, 69(4), 417–439. <https://doi.org/10.26443/law.v69i4.1627>
- Renier, L. A., Schmid Mast, M., & Bekbergenova, A. (2021). To err is human, not algorithmic – robust reactions to erring algorithms. *Computers in Human Behavior*, 124, Article 106879. <https://doi.org/10.1016/j.chb.2021.106879>
- Saima, & Khan, M. A. (2020). Effect of social media influencer marketing on consumers' purchase intention and the mediating role of credibility.

Journal of Promotion Management, 27(4), 503–523.

<https://doi.org/10.1080/10496491.2020.1851847>

Sands, S., Campbell, C., Ferraro, C., Demsar, V., Rosengren, S., & Farrell, J.

(2024). Principles for advertising responsibly using generative AI.

Organizational Dynamics, 53(2), Article 101042.

<https://doi.org/10.1016/j.orgdyn.2024.101042>

Saura, J. R., Palos-Sánchez, P., & Suárez, L. M. C. (2017). Understanding the

digital marketing environment with KPIs and web analytics. *Future*

Internet, 9(4), Article 76. <https://doi.org/10.3390/fi9040076>

Schickmair, V., Rudisch-Sommer, R., & Stöckl, A. (2025). From words to

conversions: Leveraging large language models for dynamic landing

page generation. *International Conference on Artificial Intelligence,*

Computer, Data Sciences, and Applications, ACDSA 2025, 1–8.

<https://doi.org/10.1109/acdsa65407.2025.11166477>

Sengar, S. S., Hasan, A. B., Kumar, S., & Carroll, F. (2024). *Generative artificial intelligence: A systematic review and applications*. arXiv.

<https://arxiv.org/abs/2405.11029>

Sew, C., Schömb, S., Kelly, R. M., & Goncalves, J. (2025). User willingness to

engage with AI-generated influencer content. *Proceedings of the CHI*

Conference on Human Factors in Computing Systems.

<https://doi.org/10.1145/3706599.3719916>

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-*

experimental designs for generalized causal inference. Houghton Mifflin.

Shorten, A., & Smith, J. (2017). Mixed methods research: Expanding the

evidence base. *Evidence-Based Nursing*, 20(3), 74–75.

<https://doi.org/10.1136/eb-2017-102699>

- Singh, S., & Sagar, R. (2021). A critical look at online survey or questionnaire-based research studies during COVID-19. *Asian Journal of Psychiatry*, 65, Article 102850. <https://doi.org/10.1016/j.ajp.2021.102850>
- Stamkou, C., Saprikis, V., Fragulis, G. F., & Antoniadis, I. (2025). User experience and perceptions of AI-generated e-commerce content: A survey-based evaluation of functionality, aesthetics, and security. *Data*, 10(6), Article 89. <https://doi.org/10.3390/data10060089>
- Stein, J. P., Linda Breves, P., & Anders, N. (2024). Parasocial interactions with real and virtual influencers: The role of perceived similarity and human-likeness. *New Media and Society*, 26(6), 3433–3453. <https://doi.org/10.1177/14614448221102900>
- Sullivan, G. M., & Anthony R Artino, J. (2013). Analyzing and interpreting data from Likert-type scales. *Journal of Graduate Medical Education*, 5(4), 541. <https://doi.org/10.4300/jgme-5-4-18>
- Sundar, S. S. (2008). The MAIN model: A heuristic approach to understanding technology effects on credibility. In M. J. Metzger & A. J. Flanagin (Eds.), *Digital media, youth, and credibility* (pp. 73–100). MIT Press. <https://doi.org/10.1162/dmal.9780262562324.073>
- Swart, J. (2021). Experiencing algorithms: How young people understand, feel about, and engage with algorithmic news selection on social media. *Social Media + Society*, 7(2). <https://doi.org/10.1177/20563051211008828>
- Thomas, M. R., & Mathew, J. (2023). Landing page via perceived interactivity for four generational cohorts. *Ushus Journal of Business Management*, 22(3), 27–46. <https://doi.org/10.12725/ujbm.64.3>

- Thon, F. M., & Jucks, R. (2017). Believing in expertise: How authors' credentials and language use influence the credibility of online health information. *Health Communication, 32*(7), 828–836.
<https://doi.org/10.1080/10410236.2016.1172296>
- Toff, B., & Simon, F. M. (2025). "Or they could just not use it?": The dilemma of AI disclosure for audience trust in news. *International Journal of Press/Politics, 30*(4), 881–903.
<https://doi.org/10.1177/19401612241308697>
- Vetter, T. R. (2017). *Descriptive statistics: Reporting the answers to the 5 basic questions of who, what, why, when, where, and a sixth, so what?* *Anesthesia and Analgesia, 125*(5), 1797–1802.
<https://doi.org/10.1213/ane.00000000000002471>
- Wahid, R., Mero, J., & Ritala, P. (2023). Editorial: Written by ChatGPT, illustrated by Midjourney: Generative AI for content marketing. *Asia Pacific Journal of Marketing and Logistics, 35*(8), 1813–1822.
<https://doi.org/10.1108/apjml-10-2023-994>
- Wang, P., & Huang, Q. (2023). Digital influencers, social power and consumer engagement in social commerce. *Internet Research, 33*(1), 178–207.
<https://doi.org/10.1108/intr-08-2020-0467>
- White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with ChatGPT. arXiv.
<https://arxiv.org/abs/2302.11382>
- Wymer, W., & Akbar, M. M. (2017). Brand authenticity, its conceptualization, and its relevance to nonprofit marketing. *International Review on Public and*

Nonprofit Marketing, 14(3), 359–374. <https://doi.org/10.1007/s12208-017-0177-z>

Yao, Y., Duan, J., Xu, K., Cai, Y., Sun, Z., & Zhang, Y. (2024). A survey on large language model (LLM) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing*, 4(2), Article 100211. <https://doi.org/10.1016/j.hcc.2024.100211>

Yardley, L., Spring, B. J., Riper, H., Morrison, L. G., Crane, D. H., Curtis, K., Merchant, G. C., Naughton, F., & Blandford, A. (2016). Understanding and promoting effective engagement with digital behavior change interventions. *American Journal of Preventive Medicine*, 51(5), 833–842. <https://doi.org/10.1016/j.amepre.2016.06.015>

Yeadon, W., Agra, E., Inyang, O.-O., Mackay, P., & Mizouri, A. (2024). *Evaluating AI and human authorship quality in academic writing through physics essays*. arXiv. <https://arxiv.org/abs/2403.05458>

Ytre-Arne, B., & Moe, H. (2021). Folk theories of algorithms: Understanding digital irritation. *Media, Culture & Society*, 43(5), 807–824. <https://doi.org/10.1177/0163443720972314>

Zhang, J. J., Wang, Y. W., Ruan, Q., & Yang, Y. (2024). Digital tourism interpretation content quality: A comparison between AI-generated content and professional-generated content. *Tourism Management Perspectives*, 53, Article 101279. <https://doi.org/10.1016/j.tmp.2024.101279>

Zhang, X., & Choi, J. (2022). The importance of social influencer-generated contents for user cognition and emotional attachment: An information relevance perspective. *Sustainability*, 14(11), Article 6676. <https://doi.org/10.3390/su14116676>

- Zhang, Y., & Gosline, R. (2023). Human favoritism, not AI aversion: People's perceptions (and bias) toward generative AI, human experts, and human–GAI collaboration in persuasive content generation. *Judgment and Decision Making*, 18. <https://doi.org/10.1017/jdm.2023.37>
- Zhe, C., & Srijinda, P. (2024). The impact of AI-generated content on content consumption habits of Chinese social media users through Xiaohongshu application. *Edelweiss Applied Science and Technology*, 8(6), 1504–1516. <https://doi.org/10.55214/25768484.v8i6.2268>
- Zhou, T., & Lu, H. (2025). The effect of trust on user adoption of AI-generated content. *The Electronic Library*, 43(1), 61–76. <https://doi.org/10.1108/el-08-2024-0244>
- Zhu, T., Weissburg, I., Zhang, K., & Wang, W. Y. (2025). Human bias in the face of AI: Examining human judgment against text labeled as AI generated. *Findings of the Association for Computational Linguistics: ACL 2025*, 25907–25914. <https://doi.org/10.18653/v1/2025.findings-acl.1329>
- Zniva, R., Weitzl, W. J., & Lindmoser, C. (2023). Be constantly different! How to manage influencer authenticity. *Electronic Commerce Research*, 23(3), 1485–1514. <https://doi.org/10.1007/s10660-022-09653-6>

Appendix

Appendix A: AUTECH Approval Letter



Auckland University of Technology Ethics Committee (AUTECH)

23 September 2025

Matthew Guinibert
Faculty of Design and Creative Technologies

Dear Matthew

Re Ethics Application: **25/317 Users' perception of AI-generated digital marketing content.**

Thank you for your responses to AUTECH's conditions.

Your ethics application has been approved for three years until 23 September 2028.

Non-Standard Conditions of Approval

1. The survey can be open longer than 2 weeks, the committee was just highlighting that 'indefinitely' might be unrealistic.
2. Instead of using google docs as the mechanism for the summary of research findings either, the 'thank you' page on Qualtrics is suitable for this purpose or a url in the Information Sheet to the supervisors AUT professional web page (where the findings can be posted) might be more appropriate in this instance.

Non-standard conditions do not need to be submitted to or reviewed by AUTECH unless requested but must be completed before commencing your study.

Standard Conditions of Approval

1. The research is to be undertaken in accordance with the [Auckland University of Technology Code of Conduct for Research](#) and as approved by AUTECH.
2. All public facing documents must have the AUTECH approval number and be of a high standard of spelling and grammar. Dates on the Information Sheet(s) and Consent Form(s) must be consistent.
3. Any amendments to the project must be approved by AUTECH prior to being implemented.
4. A progress report is due annually on the anniversary of the approval date.
5. A final report is due at the expiration of the approval period, or, upon completion of project.
6. Any serious or adverse events must be reported to AUTECH, this includes unforeseen issues that might affect continued ethical acceptability of the project.
7. AUTECH grants ethical approval only. You are responsible for obtaining management permission for access from any institution or organisation at which your research is being conducted and you need to meet all ethical, legal, public health, and locality obligations or requirements for the jurisdictions in which the research is being undertaken.

The application number and title need to be referenced on all correspondence related to this project.

All forms are available online <http://www.aut.ac.nz/research/researchethics>

For any enquiries, please contact the Secretariat at ethics@aut.ac.nz
(This is a computer-generated letter for which no signature is required)

The AUTECH Secretariat
Auckland University of Technology Ethics Committee

Cc: rnw8831@aut.ac.nz

Appendix B: Participant Information Sheet



Participant Information Sheet

Date that data collection will start:

1 October 2025 (subject to Ethics approval)

Project Title: Users' perception of AI-generated digital marketing content

Hello,

You are invited to participate in a research study on how users perceive AI-generated digital marketing content. This study is being conducted by Hira Mirza from the Auckland University of Technology, New Zealand. Other research team members include Dr Matthew Guinibert. The study is being carried out as a requirement for the thesis component of a Master of Communication Studies.

What is the purpose of this research?

This study investigates how users perceive AI-generated and human-written digital marketing content, and whether knowing the author influences those perceptions.

Participants will be asked to evaluate short website landing pages created either by ChatGPT (an AI tool) or by a professional digital content specialist. The findings will help researchers and marketers understand how generative AI can be used responsibly in customer-facing communication, without undermining trust.

Approximately 100 participants will be recruited internationally via CloudResearch Connect.

How was I identified and why am I being invited to participate in this research?

You have been invited because you are a registered member of CloudResearch Connect. The only inclusion criteria are that you are 18 years or older and have internet access. There are no exclusion criteria.

How do I agree to participate in this research?

Your participation in this research is voluntary (it is your choice). Whether you choose to participate will neither advantage nor disadvantage you. Consent will be obtained electronically through the Connect platform. Before accessing the survey, you will be asked to click "I agree" to confirm that you have read and understood this Information Sheet and that you consent to take part. You are free to withdraw at any time. To do this, simply close your browser window the survey is presented on. Any information you have entered up to that point will be deleted from the data set. As this is an anonymous survey, it will not be possible to withdraw your information after you have completed the survey.

What will my participation involve?

You will be asked to complete an online survey that takes approximately 20 minutes.

- You will review two short website pages.
- You will be randomly assigned to one of two groups:
 - Group 1: Reviews the pages without knowing who wrote them.
 - Group 2: Reviews the same pages but with authorship disclosed (AI and human).
- After reviewing each page, you will answer a short series of questions about your perceptions.
- You will then complete comparative questions and provide your views on AI use in marketing.

At the end of the survey, you will be redirected back to the Connect platform to confirm completion.

What are the benefits?

By taking part in this study, you will contribute to research on the ethical and effective use of AI in digital communication. Your participation may also help you reflect on how you personally engage with online content. More broadly, the findings will provide practical guidance for marketers and businesses on how to

use AI in ways that are transparent and trustworthy, while also contributing to academic knowledge through a Master's thesis, publications, and conference presentations.

What are the costs?

There are no costs to you apart from the 20 minutes of your time. You will receive \$3.50 USD for completing the survey, paid through the Connect platform. This equates to approximately \$10.50 per hour, in line with Connect's fair compensation guidelines.

Will the results of the study be published?

The results of this research will be published in a Master's thesis. This thesis will be available to the general public through the AUT library. Results may be published in peer-reviewed, academic journals. Results will also be presented during conferences or seminars to wider professional and academic communities. You will not be identifiable in any publication.

What are the discomforts and risks?

There are no known risks beyond the inconvenience of time. You are advised not to enter personal or identifying details in any free-text responses.

What will happen to information about me?

Your responses will remain completely anonymous. CloudResearch Connect assigns each participant an ID number and does not share personal details with the researchers. Survey data will be collected in Qualtrics and then stored securely on the researcher's AUT OneDrive account, which is password-protected and access-controlled. Only the researcher and supervisor will have access to the data. All information will be kept for six years before being permanently deleted, and data in Qualtrics will be deleted shortly after analysis is complete. The results will only ever be reported in aggregate form, so no individual can be identified.

What opportunity do I have to consider this invitation?

This survey will remain open on the Connect platform for two weeks.

Will I receive feedback on the results of this research?

The results of this research will be published in a Master's thesis, which will be available to the public through [AUT's open access repository, Tuwhera](#). Further, a summary of the findings will be available on the [Google drive link](#).

What do I do if I have concerns about this research?

Any concerns regarding the nature of this project should be notified in the first instance to the Project Supervisor, Matthew Guinibert, matt.guinibert@aut.ac.nz, (+649) 921 9999 ext 8437.

Concerns regarding the conduct of the research should be notified to the Executive Secretary of AUTC, ethics@aut.ac.nz, (+649) 921 9999 ext 6038.

Researcher Contact Details:

Hira Mirza
rnw8831@autuni.ac.nz

Project Supervisor Contact Details:

Dr Matthew Guinibert
matt.guinibert@aut.ac.nz

Approved by the Auckland University of Technology Ethics Committee on *type the date final ethics approval was granted*, AUTC Reference number 25/317.

Appendix C: Participant Recruitment on CloudResearch Connect



Appendix D: Survey Consent Preview

Restart Survey

Place Bookmark

Tools

Share Preview



Please read the [information sheet](#) about this survey before you proceed.

By clicking the "I agree", I confirm that I have read the information sheet and indicate my willingness to voluntarily take part in this survey.

☐ I agree

☐ I disagree

12:29



Please read the [information sheet](#) about this survey before you proceed.

By clicking the "I agree", I confirm that I have read the information sheet and indicate my willingness to voluntarily take part in this survey.

☐ I agree

☐ I disagree

Appendix E: Survey

Survey: Users perception of AI-generated marketing content.

Start of Block: Consent

Survey on users' perception of AI-generated digital content.

Consent Please read the information sheet about this survey before you proceed. By clicking the "I agree", I confirm that I have read the information sheet and indicate my willingness to voluntarily take part in this survey.

☐ I agree (1)

☐ I disagree (2)

End of Block: Consent

Start of Block: Section 1A: Human - Treatment group

Page Break

Section 1A: Webpage content review

Webpage A is for a real estate agency called 'Home Start' that caters to first-home buyers. Please read the webpage carefully and indicate how much you agree or disagree with the statements that follow.

Webpage A was written by a human.

FirstNest Realty

First-Home Buyers

Your first home, minus the guesswork.

We help first-home buyers get finance-ready, shortlist the right homes, and make confident offers. Plain-English guidance, buyer-only viewings, and support at every step.

Book a free buyer consult

Browse starter homes

Finance-ready checklist

Understand your budget, deposits, and typical costs. We'll review your plan so you move fast when the right home appears.

Smart shortlist

We curate homes that match your must-haves: commute, schools, size, and price. No endless scrolling.

Offer support

From LIMs and building reports to auction prep, we explain the paperwork and guide you through the offer.

Getting started is simple

- 1 15-minute consult.** Tell us your goals and timeline.
- 2 Finance plan.** We share a buyer checklist and connect you with lenders or brokers if needed.
- 3 View shortlist.** Private, buyer-only viewings of homes that actually fit.
- 4 Make an offer.** We help with reports, tactics, and next steps through settlement.

FAQs

Can I buy with a small deposit?

Possibly. Options vary by lender and policy. We'll outline realistic pathways and typical trade-offs.

Do I need pre-approval before viewing?

It helps you move quickly and set a clear budget. If you don't have one yet, we can point you to next steps.

What does your service cost?

Your initial consult is free. For most purchases the seller pays the listing agent's commission. We discuss any buyer-side fees in advance.

What buyers say

"I felt totally lost at the start. FirstNest explained each step and helped us win our first home." — Aaliyah & Omar

Starter costs guide

- Deposit and legal fees
- Inspections and reports
- Moving and setup

Open-home checklist

- Moisture, light, ventilation
- Noise, neighbours, parking
- Maintenance red flags

Neighbourhood snapshot

- Walkability and transport
- Schools and childcare
- Local amenities

Indicative page prepared 2025-08-25 for research ethics review. Content is fictional; no tracking, no cookies, and no forms are active.

First, think about the content on the page.

Q1 Q. Based on this page, the content is authentic.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q2 Q. Based on this page, the content is credible.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q3 Q. Based on this page, the content is persuasive.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Instruction 2 Now, think about the author and their writing style.

Q4 Q. The author appears knowledgeable about the topic.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q5 Q. The writing on this page sounds human.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q6 Q. I connected with the writing style.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Instruction 3 Finally, based on the content of the page, indicate how you feel about the product and brand.

Q7 Q. I am interested in learning more about this product.

- ☐ Strongly Agree
- ☐ Agree
- ☐ Neither agree nor disagree
- ☐ Disagree
- ☐ Strongly disagree

Q8 Q. Based on this page, the brand seems genuine.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q9 Q. Overall, the content on this page is high quality.

- ☐ Strongly Agree
- ☐ Agree
- ☐ Neither agree nor disagree
- ☐ Disagree
- ☐ Strongly disagree

End of Block: Section 1A: Human - Treatment group

Start of Block: Section 1B: ChatGPT - Treatment group

Section 1B: Content review

Webpage B is for a real estate agency called 'FirstNest' that caters to first-home buyers. Please read the webpage carefully and indicate how much you agree or disagree with the statements that follow.

Page Break

Webpage B was written by ChatGPT.

FirstNest Realty

First-home buyers — you're closer than you think.

Let's make your first home happen.

If you've got a deposit plan, a few saved listings, and a hundred tabs open, you don't need more noise—you need a clear path and a calm guide. That's us.

Start with a friendly 15-minute chat. Bring your questions; leave with a plan.

[Book free consult](#)[See buyer roadmap](#)

Picture this

You walk out of an open home. It's sunny, the kitchen's perfect, and the price guide is...mystifying. Do you offer? Wait? Walk away? You want someone who speaks plain-English and has your back. That's the whole point of FirstNest.

"I was sure we'd mess it up. FirstNest turned the unknowns into a checklist. We got pre-approval, viewed fewer homes, and won one." — **Amelia & Yusuf**

What you get with us

Money, made simple

We translate deposits, fees, and reports into a one-page plan. No jargon. No surprises.

A shortlist that fits your life

Commute, schools, sunlight, storage—we filter the noise so viewings feel purposeful.

Offer with confidence

Reports, clauses, auction basics—we walk through each step, then help you play it smart.

We're buyer-first

Honest timelines, clear trade-offs, and heads-up on red flags. No hard sell. Ever.

Your 7-day kickstart

- 1 **Day 1:** 15-minute chat. Goals, budget, questions.
- 2 **Day 2:** Money map. Deposit, true costs, next steps.
- 3 **Day 3:** Shortlist agreed. 3–5 homes to view.
- 4 **Day 4–5:** Private, buyer-only viewings.
- 5 **Day 6:** Reports reviewed. Offer strategy set.
- 6 **Day 7:** Make a confident offer (or refine).

Real talk

Good homes move fast. Fear and FOMO are real. We'll keep you grounded with facts and a plan you can explain to your future self.

Who we're great for

- First-time buyers who value straight talk
- Busy people who want a trusted guide
- Anyone who wants to understand each step—not guess it

Ready to feel sure about your next step?

[Book free consult](#)[Get the 7-day plan](#)

Instruction 4 First, think about the content on the page.

Q10 Q. Based on this page, the content is authentic.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q11 Q. Based on this page, the content is credible.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q12 Q. Based on this page, the content is persuasive.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Instruction 5 Now, think about the author and their writing style.

Q13 Q. The author appears knowledgeable about the topic.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q14 Q. The writing on this page sounds human.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q15 Q. I connected with the writing style.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Instruction 6 Finally, based on the content of the page, indicate how you feel about the product and brand.

Q16 Q. I am interested in learning more about this product.

- ☐ Strongly Agree
- ☐ Agree
- ☐ Neither agree nor disagree
- ☐ Disagree
- ☐ Strongly disagree

Q17 Q. Based on this page, the brand seems genuine.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither agree nor disagree
 - ☐ Disagree
 - ☐ Strongly disagree
-

Q18 Q. Overall, the content on this page is high quality.

- ☐ Strongly Agree
- ☐ Agree
- ☐ Neither agree nor disagree
- ☐ Disagree
- ☐ Strongly disagree

End of Block: Section 1B: ChatGPT - Treatment group

Start of Block: Section 2: Comparison

Section 2: Comparison between the webpages.

Q37 Q. In one or two sentences, what was the most notable difference between the two pages?

Instruction 13 The following statements are aspects to compare between Webpage A (Home Start) and Webpage B (FirstNest Realty). For each one, indicate which page was better in your opinion.

Q38 Q. Authenticity of the content

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Q39 Q. Credibility of the information

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Display this question:

*If Q. Credibility of the information = Webpage A is much better
Or Q. Credibility of the information = Webpage A is slightly better
Or Q. Credibility of the information = Webpage B is slightly better
Or Q. Credibility of the information = Webpage B is much better*

Q39A Q. In one sentence, what, if anything, increased or reduced your trust in the page you liked more?

Q40 Q. Persuasiveness of the page

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Q41 Q. How knowledgeable the content appears

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Q42 Q. Writing style resonated with me

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Q43 Q. Human-sounding writing

- ☐ Webpage A is much better
- ☐ Webpage A is slightly better
- ☐ No difference between A and B
- ☐ Webpage B is slightly better
- ☐ Webpage B is much better
- ☐ Can't tell

Display this question:

If Q. Human-sounding writing = Webpage A is much better

Or Q. Human-sounding writing = Webpage A is slightly better

Or Q. Human-sounding writing = Webpage B is slightly better

Or Q. Human-sounding writing = Webpage B is much better

Q43A Q. What specific cues made the writing feel more human or more AI-like (e.g., phrasing, repetition, formality, or something else)? One example is enough.

Q44 Q. Made me want to learn more about the product

- ☐ Webpage A is much better
 - ☐ Webpage A is slightly better
 - ☐ No difference between A and B
 - ☐ Webpage B is slightly better
 - ☐ Webpage B is much better
 - ☐ Can't tell
-

Q45 Q. How genuine the brand seems (based on the page)

- ☐ Webpage A is much better
 - ☐ Webpage A is slightly better
 - ☐ No difference between A and B
 - ☐ Webpage B is slightly better
 - ☐ Webpage B is much better
 - ☐ Can't tell
-

Q46 Q. Overall quality of the page's content

- ☐ Webpage A is much better
 - ☐ Webpage A is slightly better
 - ☐ No difference between A and B
 - ☐ Webpage B is slightly better
 - ☐ Webpage B is much better
 - ☐ Can't tell
-

Page Break

Q47 Q. Which page would you be more likely to act on (e.g., click through, sign up, or purchase) today?

- ☐ Webpage A much more likely
- ☐ Webpage A slightly more likely
- ☐ No difference
- ☐ Webpage B slightly more likely
- ☐ Webpage B much more likely

Display this question:

If Q. Which page would you be more likely to act on (e.g., click through, sign up, or purchase) today? = Webpage A much more likely

Or Q. Which page would you be more likely to act on (e.g., click through, sign up, or purchase) today? = Webpage A slightly more likely

Or Q. Which page would you be more likely to act on (e.g., click through, sign up, or purchase) today? = Webpage B slightly more likely

Or Q. Which page would you be more likely to act on (e.g., click through, sign up, or purchase) today? = Webpage B much more likely

Q47A Q. In one or two sentences, what specific thing about the page you would act on most influenced your decision?

End of Block: Section 2: Comparison

Start of Block: Section 3: Use of AI

Section 3: Your experience with AI-generated content.

Q48 Q. Which best describes how often you personally use generative AI tools (e.g., ChatGPT, Gemini, Claude)?

- ☐ Never
 - ☐ Less than once a month
 - ☐ 1 to 3 times a month
 - ☐ 1 to 2 times a week
 - ☐ 3 to 6 times a week
 - ☐ Daily or more
-

Instruction 14 The following statements are about brands using AI to create marketing content. Please indicate how much you agree or disagree with each statement.

Q49 Q. I want brands to clearly state when marketing content is AI-generated.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither Agree nor Disagree
 - ☐ Disagree
 - ☐ Strongly Disagree
-

Q50 Q. I trust marketing content less when it is labeled as AI-generated.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither Agree nor Disagree
 - ☐ Disagree
 - ☐ Strongly Disagree
-

Q51 Q. Knowing who created marketing content helps me judge it fairly.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither Agree nor Disagree
 - ☐ Disagree
 - ☐ Strongly Disagree
-

Q52 Q. I am comfortable with brands using AI to create their marketing content.

- ☐ Strongly Agree
 - ☐ Agree
 - ☐ Neither Agree nor Disagree
 - ☐ Disagree
 - ☐ Strongly Disagree
-

Page Break

Q53 Q. In one or two sentences, how do you feel about brands using AI to write narrative marketing content (e.g., social posts, ads, email sequences)?

End of Block: Section 3: Use of AI

Start of Block: Section 1AA: Human - Control group

Page Break