

**Exploring How AI Disclosure in Blog Posts Affects the Perceptions of Brand Warmth and
Competence**

Mohamad Shayto

A Thesis

In the John Molson School of Business

Department of Marketing

Presented in Partial Fulfillment of the Requirements

For the Degree of

Master of Science (Marketing)

at Concordia University

Montreal, Quebec, Canada

August 2024

© Mohamad Shayto, 2024

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis

Prepared By: Mohamad Shayto

Entitled: Exploring How AI Disclosure in Blog Posts Affects the Perceptions of Brand Warmth and Competence

and submitted in partial fulfillment of the requirements for the degree of

Master of Science (Marketing)

Complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final Examining Committee:

_____ Examiner

Dr. Darlene Walsh

_____ Examiner

Dr. Pierre-Yann Dolbec

_____ Thesis Supervisor

Dr. Sharlene He

Approved by:

_____ Graduate Program Director

Dr. Mrugank Thakor

_____ Dean, John Molson School of Business

Dr. Anne-Marie Croteau

ABSTRACT

Exploring How AI Disclosure in Blog Posts Affects the Perceptions of Brand Warmth and Competence

Mohamad Shayto

Marketers increasingly employ generative AI technologies in content creation. However, whether and how disclosing AI authorship may affect consumers' brand perceptions remains underexplored. This thesis investigates the impact of disclosing whether marketing content was AI-generated or human-written, specifically on consumer perceptions of the brand's warmth and competence. In general, I hypothesize that brand perceptions will be more negative when content (e.g., a blog post) is disclosed to be authored by AI compared to by a human. However, I theorize the effect may depend on whether the content is informational or narrative in nature, such that the negative impact of AI disclosure is expected to be greater for narrative content than informational content. These hypotheses were examined across two experiments. The results of Study 1 revealed that human-written content was perceived as warmer and more competent, enhancing brand credibility, brand attitude, and purchase intentions. Content type did not moderate these effects. Study 2 replicated the effect of disclosure on warmth, and additionally found that content type moderated this effect (although the moderation did not operate as expected). No significant effects on competence emerged in this study. Theoretical and managerial contributions are discussed, especially regarding the strategic use of human authorship in content marketing.

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Sharlene He, who has provided patient guidance for me throughout this process. Thank you for your time in answering my questions and helping walk me through the steps of writing a thesis for the very first time.

I would also like to thank my family overseas, who have always had my back in a time which involved many personal-life and health challenges to overcome.

Last but not least, thank you to my partner, who's love and presence mean the world.

Table of Contents

List of Figures	vi
Introduction	1
Theoretical Background.....	2
The Role of AI in Content Creation.....	3
Consumer Attitudes Towards AI.....	4
Perceptions of Brand Warmth and Competence.....	6
The Moderating Role of Content Type: Narrative vs. Informational.....	7
Overview of Experiments	12
Experiment 1	12
Methods	13
Results and Discussion	17
Experiment 2	25
Methods	25
Results and Discussion	30
General Discussion	37
Limitations and Future Directions for Research.....	38
Practical Contributions.....	42
Theoretical Contributions.....	43
Conclusion.....	44
References.....	45
Appendix A: Experimental Stimuli.....	49

List of Figures

Figure 1. Average Brand Warmth Rating by Content Type and Disclosure (Experiment 1)

Figure 2. Average Brand Competence Rating by Content Type and Disclosure (Experiment 1)

Figure 3. Average Brand Attitude Rating by Content Type and Disclosure (Experiment 1)

Figure 4. Average Purchase Intentions Rating by Content Type and Disclosure (Experiment 1)

Figure 5. Average Brand Credibility Rating by Content Type and Disclosure (Experiment 1)

Figure 6. Average Brand Connection Rating by Content Type and Disclosure (Experiment 1)

Figure 7. Average Brand Warmth Rating by Content Type and Disclosure (Experiment 2)

Figure 8. Average Brand Competence Rating by Content Type and Disclosure (Experiment 2)

Figure 9. Average Brand Attitude Rating by Content Type and Disclosure (Experiment 2)

Figure 10. Average Brand Credibility Rating by Content Type and Disclosure (Experiment 2)

Figure 11. Average Brand Perceptions of Technological Abilities Rating by Content Type and Disclosure (Experiment 2)

INTRODUCTION

Artificial intelligence (AI) is rapidly transforming the way that modern businesses operate, allowing for efficiency and innovation across a plethora of different sectors in a way never witnessed before. AI has been shown to have the potential to optimize supply chains, excel in task automation, enhance customer experiences, and improve decision making by providing valuable data insights (Allioui et al., 2023). For example, some of the sectors that ChatGPT (a natural language processing model developed by OpenAI) has had the most transformative impact on are e-commerce (reducing costs and improving user experience), healthcare (streamlining patient interactions and improving communication between healthcare providers), and digital marketing (increased efficiency through automation and better decision-making) (George et al., 2023).

In the world of marketing, especially digital marketing, AI has presented itself as a powerful tool that enables brands to reduce the time and cost of creating their marketing campaigns, as well as increasing their effectiveness (Roggeveen & Sethuraman, 2020; Van Esch & Stewart Black, 2023). Thanks to novel AI tools available to the public, such as OpenAI's ChatGPT, Microsoft's Copilot, and Google's Gemini, brands across the world now enjoy the ability to generate blog posts, social media updates, as well as product description listings at scale, all of which enable these brands to maintain a consistent and engaging online presence at a fraction of the resources previously required to do so. In this day and age, being able to generate timely and personalized content on demand is vital for businesses, and generative AI has shown potential to be an ideal solution for brands looking to meet these ever increasing demands.

As consumers become increasingly knowledgeable about the use of AI, they may become more and more aware of the different ways that AI is used by companies and brands they interact with on a daily basis, such as Starbucks (Palumbo & Edelman, 2023). One question that arises is how consumers' brand perceptions may be affected by the knowledge that marketing content is AI-generated. In particular, I explore how disclosure of AI-generated content (e.g., blog posts by brands) may affect consumers' perceptions of the brand. I seek to identify conditions under which disclosure may be beneficial vs. detrimental to brand perceptions.

The impact of AI disclosure on consumer perceptions remains under-researched in academic literature. While there has been considerable attention given to the broader implications of AI use in marketing (Paschen et al., 2019)—such as its role in personalizing customer experiences, optimizing advertising campaigns, and improving customer service—very few studies have specifically examined how consumers' brand perceptions change once they are made aware of AI use in the content marketing practices of brands they interact with. Existing literature has predominantly focused on consumer trust in AI-driven decision-making, such as studies which aimed to explore the conditions under which consumers are more likely to trust and accept AI-generated recommendations (Kim et al., 2021), as well as studies on the ethical implications of AI use in marketing (Gonçalves et al., 2023).

In this thesis, I aim to fill this gap in existing literature by examining how the disclosure of AI authorship affects consumer perceptions of brand warmth and competence, two critical dimensions of brand perception that affect trust and loyalty (Kervyn et al., 2022). In this thesis, I theorize that different types of content (i.e., narrative vs. informational blog posts) may affect the

impact of disclosure on brand perceptions. This research aims to provide valuable insights for marketers (both practitioners and academics alike) regarding how to best navigate the complex landscape of AI-driven content strategies, in the midst of all the tremendous opportunities this new technology provides.

THEORETICAL BACKGROUND

The Role of AI in Content Creation

Artificial intelligence (AI) has fundamentally transformed the industry of content creation, emerging as a powerful tool for brands striving to engage with their audiences in an increasingly competitive digital environment. In order to stand out in a sea of intensified competition among brands in the digital age, each trying to capture as much of the consumer attention pie as possible in a crowded online space, businesses are required to produce personalized, high-quality content at an unprecedented scale.

Existing academic literature sheds light on how AI tools are revolutionizing content generation by automating processes that once required significantly more resources, namely human effort and creativity. As Esch & Black (2021) note, AI can significantly reduce the time and costs associated with content creation, and by automating repetitive tasks that would typically require hands-on human effort, AI enables marketing professionals to focus on higher-value activities. In their work, Kose & Sert (2016) highlight how AI is increasingly becoming crucial to content marketing efforts, and introduce several intelligent content marketing models that leverage the power of AI in order to improve the adaptability, flexibility, and interactivity of

marketing strategies.

However, there are also downsides of using AI for content creation. In their research, Brüns & Meißner (2024) explored the implications of Generative AI adoption for brands as part of their social media marketing, and their research revealed that while Generative AI can significantly enhance the efficiency and scale of social media content production, followers on social media platforms tend to perceive brands using Generative AI for content creation as less authentic, which can lead to negative attitudes and perceptions towards the brand, as well as negative behavioral reactions.

Consumer Attitudes Towards AI

As AI technologies become increasingly prevalent across various sectors in our day to day lives, understanding consumer attitudes towards AI is important for businesses looking to integrate these technologies into their operations. The literature has generally found evidence of algorithm aversion, which describes a reluctance to rely on algorithms in making judgments and decisions (Burton et al., 2020). This term was coined by Dietvorst et al. (2015), who found that participants refused to put their trust in a forecasting model after observing it work, and preferred to put their trust in a human forecast instead, even when human errors exceeded those made by the algorithm. According to Dietvorst et al. (2015), this is due to people more quickly losing confidence in algorithms than they lose confidence in human forecasters after seeing them make the same mistake. Research by Madhavan & Wiegmann (2007) experimentally varied the algorithm's expertise (novice vs expert) and the human agent's expertise (novice vs expert), and

found that decision makers prefer the algorithm labeled as novice over the human novice, but prefer the human expert over the algorithm labeled as an expert, meaning that users are more likely to develop algorithm aversion if they could interact with human agents that had high expertise. However, it is worth noting that other research by Logg et al. (2019) found that people will prefer to rely on algorithmic advice under the right conditions. Algorithmic advice may not have identical effects as generative AI—depending on the situation, some effects may be similar, while others may differ. Thus far, more work has been done in the domain of algorithms than in the domain of AI.

Consumer trust in AI is a crucial factor when it comes to increasing its acceptance and widespread adoption. Kim et al. (2019) examine consumers' trust of AI-generated information and how this trust impacts their behavioral and evaluative responses. This research specifically focuses on the aspect of preciseness when delivering AI-generated information to the consumer, showing that consumers tend to have higher evaluations and behavioral intentions when AI-generated information is presented in a precise format (such as stating there is a 59.85% chance of rain) as compared to an imprecise format (such as saying there is a 60% chance of rain). All in all, this research sheds light on how consumers might perceive precise information as more credible and reliable.

Despite the growing body of research on consumer attitudes towards AI, a notable gap exists in the literature in regards to how these attitudes help shape consumer perceptions of AI-generated content, especially within marketing contexts. Namely, the effects of AI disclosure on perceived brand warmth and competence remain underexplored. In this thesis, I aim to bridge

this gap by looking into how consumer attitudes toward AI influence their perceptions of content disclosed as AI-generated compared to that disclosed as human-written.

Perceptions of Brand Warmth and Competence

Brand warmth and competence are two fundamental dimensions of brand perception that play a crucial role in shaping consumer behavior. In their research, Fiske et al. (2006) state that warmth refers to traits related to perceived intent, such as friendliness, trustworthiness, and morality, while competence refers to traits related to perceived ability, such as intelligence, skill, and efficiency. Aaker et al. (2010) built on the work of Fiske et al. (2006) and extended this understanding to the domain of consumer behavior, where they showed that the dimensions of warmth and competence also govern how consumers perceive firms. These attributes are not only pivotal in forming overall brand perception, but they also significantly influence consumer trust, loyalty, and purchasing decisions.

I suggest that warmth and competence are especially relevant brand dimensions to consider in exploring the impact of AI-use disclosure in content marketing. I theorize that the dimension of warmth is important because AI is generally perceived as differing in authenticity, relatability, and emotional resonance compared to humans. In their research, Castelo et al. (2019) found that even though consumers think of algorithms as less useful for subjective tasks, these effects are diminished when individuals see that algorithms have high affective human-likeness.

I also theorize that the dimension of competence is important because while AI

technologies have made notable strides in generating contextually appropriate content about most topics, it is still viewed that AI does not have the deep understanding and expertise that human writers bring to the table. Salvagno et al. (2023), note that AI should not replace human judgment and expertise, since the guidance and supervision of human experts in the field is crucial in order to ensure the accuracy, coherence, and credibility of the content before publishing it.

I hypothesize that:

H1: Disclosure of a blog post as being AI-generated (vs. written by a human) will lead to lower perceived brand warmth and competence.

Warmth and competence were selected as the primary dependent variables in this thesis because they are exceptionally fitting to capture the nuances of consumer reactions to AI-generated content. Moreover, they are fundamental dimensions of social perception that have also been shown to be relevant to brand perception, but despite their importance, these dimensions have not yet been examined in existing literature in the context of AI use.

The Moderating Role of Content Type: Narrative vs. Informational

Brands' content marketing strategies often involve posting different types of content, each serving a unique purpose in the overall marketing mix. Two of the most common types of content are narrative and informational blog posts.

Informational content is designed to provide factual information, educate the consumer,

and help establish the brand's expertise in a particular domain. By providing detailed, accurate and reliable information on topics of interest to target customers, brands can improve their perceived competence and establish themselves as authorities in their respective field. For example, a cybersecurity company might publish blog posts that delve into how to best choose the right antivirus for the consumer's unique needs. Through this type of content, brands seek to position themselves as reliable experts in their fields. Wang et al. (2009) differentiate between high-involvement consumers, who they define as those interested in the product and respond better to informational appeal, and low-involvement consumers, who they define as those more influenced by emotional appeals that require less cognitive effort.

On the other hand, through narrative content, which Pulizzi (2012) considers to be essential to modern day content marketing, brands seek to tell stories which engage consumers on an emotional level and resonate with the reader's own experiences, aspirations, or values. Through this this type of content, brands often seek to build a connection between themselves and their audience by "humanizing" the brand and making it relatable to the reader. For example, a mattress brand might share a blog post detailing a customer's own experience with one of their products, emphasizing everything they experienced along the way of the buying journey. Narrative content revolves around storytelling, through which the brand seeks to engage the reader emotionally by presenting a sequence of events or experiences. Through narratives, brands aim to create an emotional connection with the audience, making the brand appear more relatable, human, and warm. Narratives often involve characters, plots, and themes which resonate with the personal experiences and interests of the audience, thereby creating a base for

empathy and emotional involvement. Narratives can be highly effective because they can transport the consumer. Transportation refers to psychological absorption into a story, where the reader becomes deeply engaged with the narrative Green & Brock (2000). This often results in changes in beliefs and attitudes of the reader in a manner that is consistent with the story's content.

In Chu & Liu's (2023) research on the use of AI-generated content in narrative storytelling (particularly OpenAI's ChatGPT), it was found that AI-generated narratives might actually perform better than human narratives when it comes to reducing counterarguing (the cognitive process in which people generate opposing thoughts in response to persuasive messages they are exposed to) in certain contexts. However, the results suggest that when disclosed as AI-generated narratives, they may be met with skepticism, which is referred to as a "labeling effect". In other words, simply informing readers that a narrative was AI-generated and not written by a human being led to lower scores on narrative transportation and higher counterarguing.

I hypothesize that disclosure of AI use will negatively impact perceptions of brand warmth for narrative blog posts but not for informational blog posts. Previous research has shown that consumers still exhibit a reluctance to trust AI systems, especially when it comes to performing tasks that require emotional intelligence (Glikson & Woolley, 2020). Narratives or stories are often associated with emotionality, and are strictly perceived as a human activity, due to the deep meaning, empathy, and personal connection they contain. Fisher (1984) notes that contrary to the rational world paradigm, which views humans as inherently rational beings whose

communicate and make decisions based on logical arguments, the narrative paradigm where humans are seen as "homo narrans", or storytelling beings who make sense of the world through stories rather than abstract logic, is a more accurate representation. Escalas (2004) note that advertisements which tell a story and engage consumers in narrative processing result in stronger SBCs (self brand connections) as compared to advertisements which do not use a narrative structure. Disclosure of AI use in generating a narrative blog post may reduce perceived brand warmth because consumers might perceive the brand as less personable, less empathetic, and ultimately less warm, as the narrative's impact is minimized by knowing that it was created by a machine instead of a human being. On the other hand, disclosure of AI use in generating an informational blog post may not affect perceived brand warmth as much, since informational content is typically valued for its accuracy and usefulness rather than its emotional resonance. Informational blog posts are designed to relay facts, suggestions and guidelines, all of which do not necessarily require the emotional and empathetic characteristics that narratives must have. Therefore, the disclosure of AI use in informational content might not lead to a significant reduction in perceived brand warmth.

I also hypothesize that disclosure of AI use will negatively impact perceptions of competence for informational blog posts but not for narrative blog posts. Consumers often seek informational content when looking for expert and reliable advice on a topic, and competence is a critical dimension in the evaluation of such content. When consumers are made aware that an AI system has generated this informational content, this may lead to skepticism regarding the depth and accuracy of the information presented, which could result in a diminished perception

of the brand's competence. However, because narrative blog posts are designed to tell stories which resonate on a personal and emotional level with the reader, competence may not be the main factor that consumers consider when evaluating narrative content. Instead, consumers may focus more on the emotional aspect of the story. In other words, AI may not be as strongly penalized by the consumer when the content's purpose is to entertain or emotionally engage the audience. Therefore, the disclosure of AI use in narrative blog posts may not significantly impact perceptions of brand competence.

These hypotheses are formally stated as follows:

H2a: A blog post disclosed as AI-generated will negatively affect perceptions of brand warmth more for narrative posts than for informational posts.

H2b: A blog post disclosed as AI-generated will negatively affect perceptions of brand competence more for informational posts than for narrative posts.

OVERVIEW OF EXPERIMENTS

This thesis comprises two main experiments designed to investigate how AI-generated versus human-written content impacts consumer perceptions, particularly in the context of brand warmth, competence, and overall brand attitude. The second experiment builds upon the findings of the first, allowing for a comprehensive exploration of the research questions and a more nuanced understanding of the different moving parts that are involved. Experiment 1 aimed to determine whether consumers perceive AI-generated content differently from human-written content across various dimensions of brand perception, and also investigated how content type (narrative vs. informational) influences these perceptions. Experiment 2 builds on the findings of Experiment 1, but the prompts were based on a completely different product category. In Experiment 2, In this experiment, participants were exposed to blog posts from another fictional brand, CyberGuardian, which offers cybersecurity solutions. This experiment delves deeper into the perceived technological capabilities of the brand.

EXPERIMENT 1

The main objective of Experiment 1 was to investigate how the disclosure of content authorship (AI-generated vs. human-written) affects consumer perceptions of warmth, competence (H1), and how blog post type (narrative vs. informational) may moderate these effects (H2). Consequences for overall brand evaluation and purchase intentions were also examined. Additional outcomes—brand credibility and brand connection—were explored.

Methods

Two-hundred and ten U.S. participants were recruited through the Connect platform and were compensated US\$0.63 for a 4-minute study. Participants who failed the attention check were excluded. After exclusions, 202 participants were included in the analyses (MAge = 36.89, SD = 10.60; 110 Male).

Participants first provided informed consent and were then presented with instructions for completing the study. Participants were then randomly assigned to one of four conditions in a 2 (content type: narrative vs. informational) by 2 (authorship disclosure: AI-generated vs. human-written) between-participants design. They were then shown a blog post about the same topic but different content type and disclosure conditions (i.e., informational post & AI-generated, informational post & human written, narrative post & AI-generated, and narrative post & human written).

The blog post was posted by a fictional brand, with the fictional brand being called SerenitySnooze, which offers premium mattresses and sleep solutions designed to enhance sleep quality and comfort, including memory foam mattresses, adjustable bed frames, and sleep accessories. The content of the posts and the brands were carefully crafted to be representative of typical blog posts in their respective categories. In developing the blog posts used in this study, ChatGPT (an AI language model developed by OpenAI) was used in order to assist in generating the content. The AI was used in order to create initial drafts of both the narrative and informational blog post types, which I then manually reviewed and edited in order to ensure they

aligned with the intended purpose of the experiment and maintained consistency in tone and style.

The decision to use a mattress brand for Study 1 stemmed from the desire to explore AI disclosure in a product category where consumer decisions are heavily influenced by both factual education and narratives/testimonials, as well as a category that is relevant to most participants. Almost everybody needs a mattress to sleep on, and not only do people rely on informational guides in order to make the right decision for their needs, they also rely on other people's personal experiences, narratives, and testimonials.

The title of the informational blog post was the following: "Choosing the Right Mattress: Your Guide to a Better Night's Sleep", while the title for the storytelling blog post was the following: "Personal Story: 'The Mattress That Changed Everything'". For both blog posts, the title and content were identical across the AI and human conditions. Only the authorship disclosure changed.

In the narrative condition, the blog post was a story, told from a first-person perspective, of a user who had been suffering from poor sleep quality due to an old and uncomfortable mattress. The narrative detailed the user's decision to invest in a new mattress from SerenitySnooze, highlighting the significant improvements in sleep quality and overall well-being experienced after the switch. The story emphasized the comfort and support provided by the new mattress and the positive impact it has had on the user's daily life. In the AI condition, participants were provided with a disclosure at the top of the blog post that read: "This SerenitySnooze blog post, written by our specialized AI system, highlights how finding the

perfect mattress for your needs can significantly improve daily life and well-being.” In the human condition, the disclosure instead read: “This SerenitySnooze blog post, written by our satisfied customer Miles Lindsay, highlights how finding the perfect mattress for your needs can significantly improve daily life and well-being.”.

In the informational condition, the blog post was crafted to provide readers with comprehensive guidance on how to choose the right mattress for optimal sleep. It offered detailed insights into different types of mattresses, as well as factors to consider when selecting a mattress for the best night’s sleep. The content aimed to educate consumers on the benefits of various mattress features and how to make an informed decision when purchasing a mattress. In the AI condition, participants were provided with a disclosure at the top of the blog post that read: “This SerenitySnooze blog post was generated by our specialized AI system to ensure you receive the most up-to-date and well-rounded advice. Our AI blends the latest sleep research with practical insights to guide you in choosing the perfect mattress.”. In the human condition, the disclosure instead read: “This SerenitySnooze blog post was written by our mattress specialist Miles Lindsay to ensure you receive the most up-to-date and well-rounded advice. Miles blends the latest sleep research with practical insights to guide you in choosing the perfect mattress.”. Appendix A provides the stimuli used in this study.

Participants were first asked manipulation and attention check questions. They were asked to rate on a 7-point scale (1 = *Not at all* to 7 = *Very Much*): "To what extent does the blog post offer useful facts and information?" (i.e., usefulness manipulation check) and "To what

extent does the blog post tell a story?" (i.e., storytelling manipulation check). Additionally, they were asked, "What was the source of the blog post?" with options "The blog post was written by a person named Miles Lindsay" and "The blog post was generated by artificial intelligence (AI)."

The goal of the storytelling manipulation check was to validate that the narrative blog posts were successfully perceived as narrative content by participants, as intended in the design of the study. If participants rated the narrative content high on the storytelling scale, it would indicate that the manipulation was successful, meaning the content was correctly interpreted as a story rather than just informational text. Similarly, the usefulness manipulation check was designed to confirm that the informational content was perceived by the participants as providing useful and practical information. A successful manipulation would mean that participants rated the informational content high on usefulness, thereby validating that the content was correctly interpreted as educational and practical.

Participants were asked to rate their perceptions of the brand's warmth and competence on a 7-point scale (1 = *Not at all* to 7 = *Very Much*). Following prior research, warmth was measured using two items (friendly, warm) while competence was measured using two items (competent, capable; Aaker et al., 2012). Participants were also asked to rate their perceptions of the blog post on a 7-point scale (1 = *Not at all* to 7 = *Very Much*) using the following terms: credible, believable, and compelling.

Participants were then asked about their evaluation of the SerenitySnooze brand on three 7-point scale items with the question "What is your evaluation of the SerenitySnooze brand?"

with the items being “*Do not like at all*” vs “*Like a lot*”, “*Very negative*” vs “*very positive*”, and “*very unfavorable*” vs “*very favorable*”.

Additionally, participants were asked to report their connection to the blog post and the brand on a 7-point scale (1 = *Not at all* to 7 = *Very Much*) with the questions: "How much did you connect with the blog post?" and "How connected do you feel to the SerenitySnooze brand?".

To assess purchase intent, participants were asked on a 7-point scale (1 = *Not at all* to 7 = *Very Much*): "If you had a need for a mattress in the next 30 days, how likely would you be to buy a SerenitySnooze mattress?" and "How likely would you be to consider the SerenitySnooze brand for a future mattress purchase?".

Finally, participants were offered an open-ended section where they could share any comments they may have had about this blog post study, and were then asked standard demographics questions (age and gender).

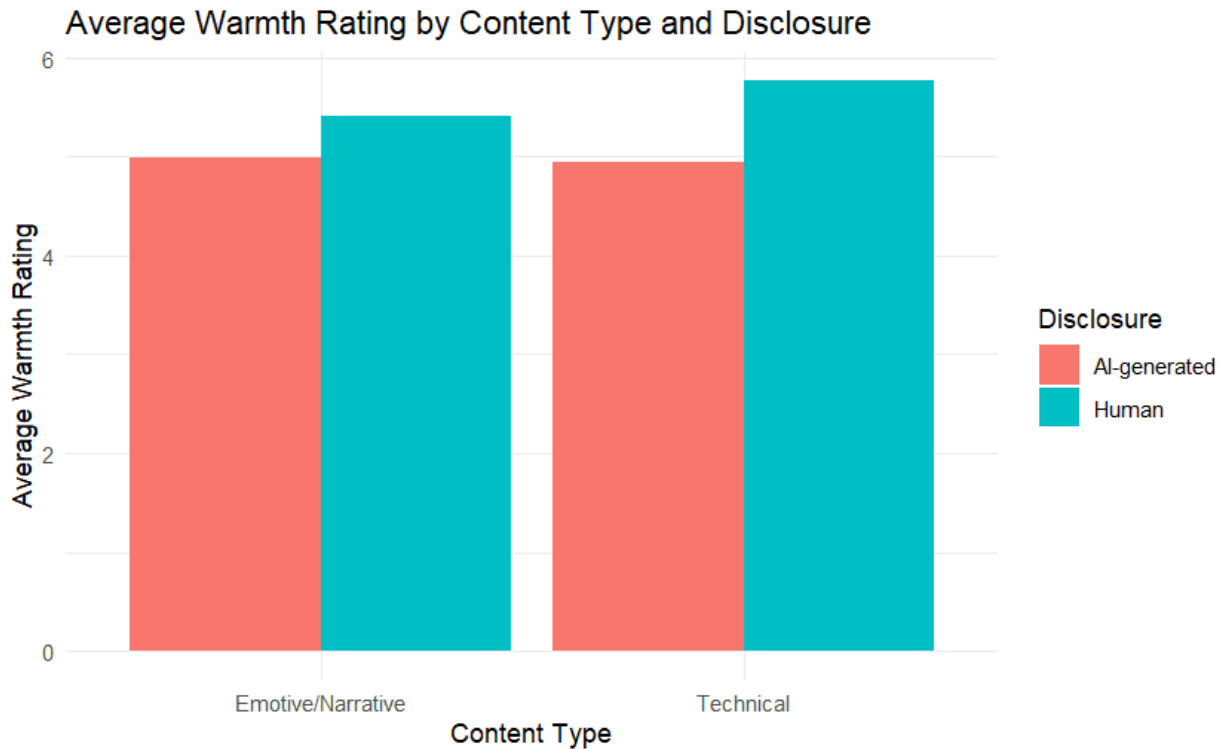
Results and Discussion

Manipulation Check: Usefulness. A two-way ANOVA was conducted to examine the effect of content type and disclosure on perceived usefulness. The main effect of content type was significant, $F(1, 198) = 56.67, p < .001, \eta^2 = .22$. Participants rated technical content ($M = 5.73, SD = 1.14$) as more useful than emotive/narrative content ($M = 4.33, SD = 1.48$). The main effect of disclosure was not significant, $F(1, 198) = 2.40, p = .123, \eta^2 = .01$. The interaction between content type and disclosure was not significant, $F(1, 198) = 0.73, p = .392, \eta^2 = 0$.

Manipulation Check: Storytelling. A two-way ANOVA was conducted to examine the effect of content type and disclosure on storytelling. The main effect of content type was significant, $F(1, 198) = 160.32, p < .001, \eta^2 = .44$. Participants rated storytelling differently depending on the content type, with Emotive/Narrative content being rated higher ($M = 6.16, SD = 0.92$) than Technical content ($M = 3.64, SD = 1.75$). The main effect of disclosure was not significant, $F(1, 198) = 2.58, p = .109, \eta^2 = .01$. The interaction between content type and disclosure was not significant, $F(1, 198) = 2.48, p = .117, \eta^2 = .012$.


Warmth. The two items measuring warmth were combined into an index ($r = 0.84$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on warmth. The main effect of content type was not significant, $F(1, 198) = 0.92, p = .338, \eta^2 = 0$. The main effect of disclosure was significant, $F(1, 198) = 13.06, p < .001, \eta^2 = .062$. Participants who were exposed to Human disclosure ($M = 5.58, SD = 1.06$) rated warmth higher than those exposed to AI-generated disclosure ($M = 4.97, SD = 1.34$). The interaction between content type and disclosure was not significant, $F(1, 198) = 1.43, p = .234, \eta^2 = 0$. See Figure 1.

Figure 1: Experiment 1 Brand Warmth Ratings




Competence. The two items measuring competence were combined into an index ($r = 0.89$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on competence. The main effect of content type was not significant, $F(1, 198) = 1.12, p = .292, \eta^2 = 0$. The main effect of disclosure was significant, $F(1, 198) = 11.11, p = .001, \eta^2 = .05$. Participants who were exposed to Human disclosure ($M = 5.75, SD = 1.08$) rated competence higher than those exposed to AI-generated disclosure ($M = 5.19, SD = 1.31$). The interaction between content type and disclosure was not significant, $F(1, 198) = 0.05, p = .819, \eta^2 = 0$. See Figure 2.

Figure 2: Experiment 1 Brand Competence Ratings

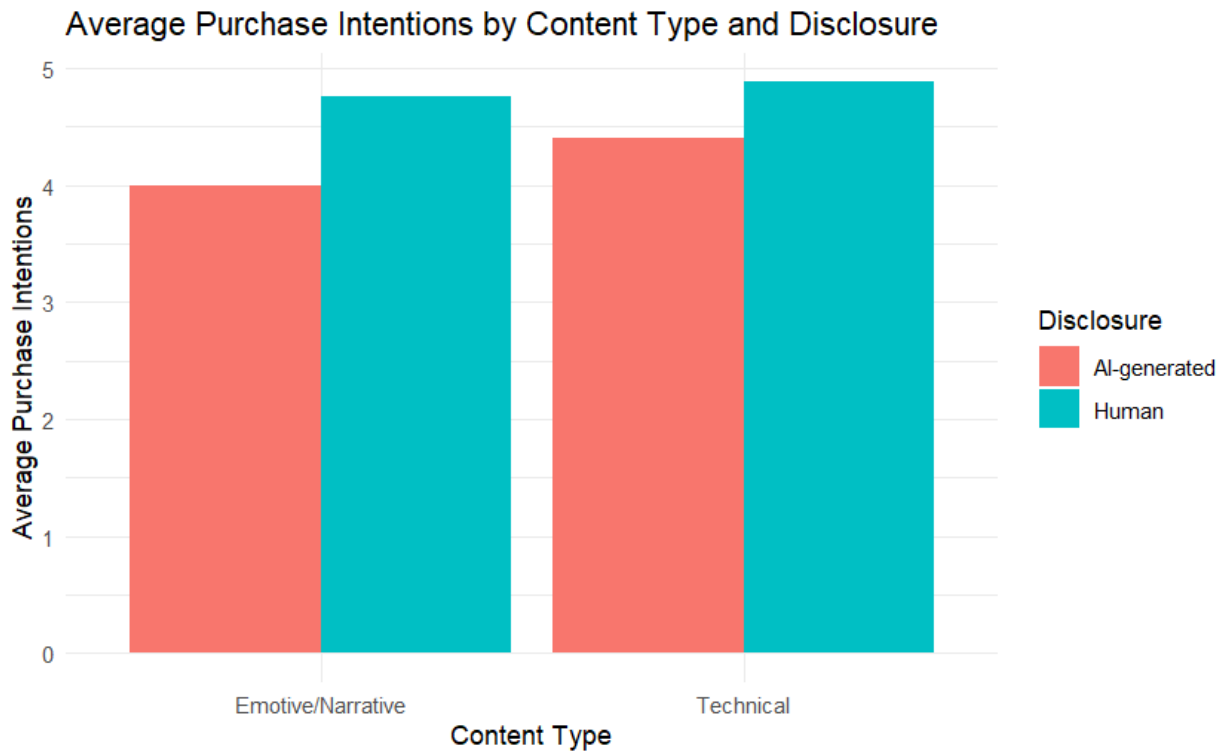


Brand Attitude. The items measuring brand attitude were combined into an index (Cronbach's $\alpha = 0.95$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on brand attitude. The main effect of content type was not significant, $F(1, 198) = 1.66, p = .198, \eta^2 = 0$. However, the main effect of disclosure was significant, $F(1, 198) = 4.64, p = .032, \eta^2 = .02$. Participants who were exposed to Human disclosure ($M = 2.75, SD = 0.70$) rated brand attitude higher than those exposed to AI-generated disclosure ($M = 2.53, SD = 0.73$). The interaction between content type and disclosure was not significant, $F(1, 198) = 0.99, p = .319, \eta^2 = 0$. See Figure 3.




Purchase Intentions. The two items measuring purchase intentions were combined into an index ($r = 0.86$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on purchase intentions. The main effect of content type was not significant, $F(1, 198) = 1.41, p = .237, \eta^2 = 0$. The main effect of disclosure was significant, $F(1, 198) = 7.72, p = .006, \eta^2 = .03$. Participants who were exposed to Human disclosure ($M = 4.82, SD = 1.43$) had higher purchase intentions than those exposed to AI-generated disclosure ($M = 4.20, SD = 1.72$). The interaction between content type and disclosure was not significant, $F(1, 198) = 0.37, p = .545, \eta^2 = 0$. See Figure 4.

Figure 4: Experiment 1 Purchase Intentions Ratings



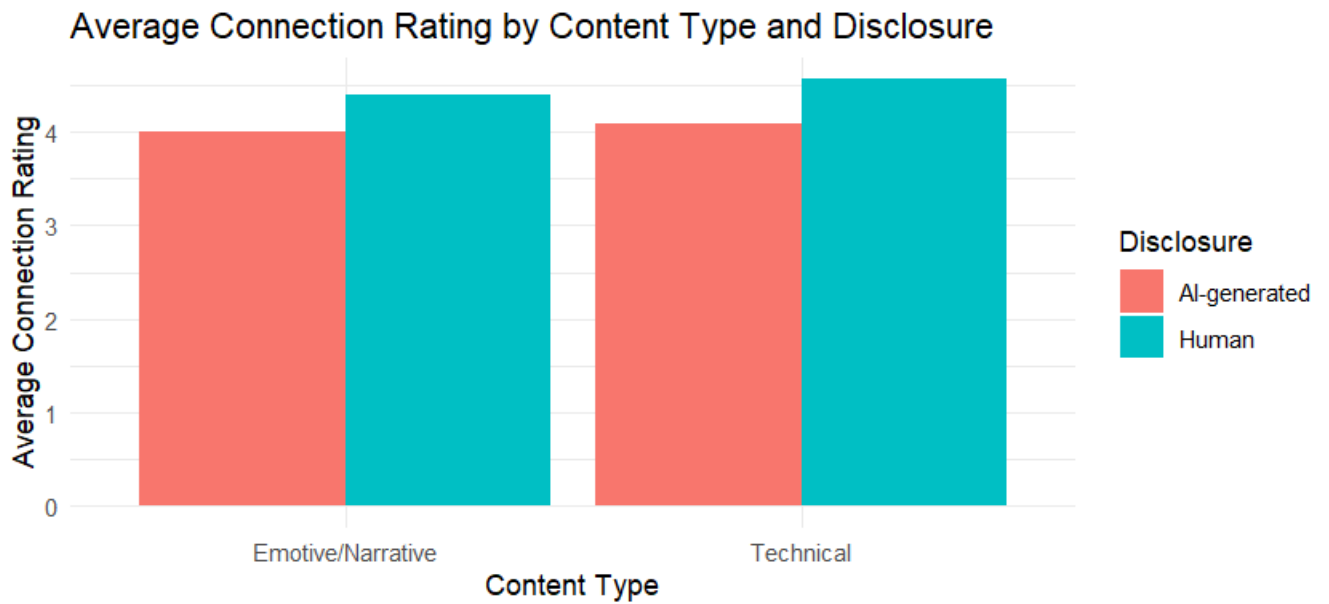
Credibility. The items measuring credibility were combined into an index (Cronbach's $\alpha = 0.91$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on credibility. The main effect of content type was significant, $F(1, 198) = 10.36, p = .002, \eta^2 = .05$. The narrative blog post rates ($M = 4.79, SD = 1.50$) higher on credibility compared to the technical post ($M = 5.36, SD = 1.08$). The main effect of disclosure was significant, $F(1, 198) = 10.97, p = .001, \eta^2 = .05$. Participants who were exposed to Human disclosure ($M = 5.37, SD = 1.15$) rated credibility higher than those exposed to AI-generated disclosure ($M = 4.78, SD = 1.44$). The interaction between content type and disclosure was not significant, $F(1, 198) = 0.55, p = .461, \eta^2 = 0$. See Figure 5.

Figure 5: Experiment 1 Brand Credibility Ratings



Connection. The two items measuring connection were combined into an index ($r = 0.86$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on connection. The main effect of content type was not significant, $F(1, 198) = 0.28, p = .597, \eta^2 = 0$. The main effect of disclosure was marginally significant, $F(1, 198) = 3.19, p = .075, \eta^2 = 0$. Participants who were exposed to Human disclosure ($M = 4.48, SD = 1.64$) rated connection higher than those exposed to AI-generated disclosure ($M = 4.05, SD = 1.76$). The interaction between content type and disclosure was not significant, $F(1, 198) = 0.03, p = .843, \eta^2 = 0$. See Figure 6.

Figure 6: Experiment 1 Brand Connection Ratings



Discussion

In Experiment 1, participants consistently rated human-disclosed content as warmer and more competent than AI-disclosed content. In addition, human-disclosed content led to higher purchase intentions and a more favorable brand attitude. Thus, H1 was supported. Content type (narrative vs. informational) did not significantly impact these perceptions, nor did it interact with the disclosure condition. Thus, H2a & H2b were not supported. These findings provide preliminary evidence that disclosure of human-authored content has more favorable impacts on various dimensions of brand perception.

Given AI is cutting-edge technology, it is possible that disclosure of AI use may have different effects for more high-tech brands. Thus, the next experiment tested these hypotheses in

a product category that is more technologically advanced than bed mattresses (cybersecurity). Experiment 2 also examined the effects on additional variables (e.g., perceived technological competence of the brand).

EXPERIMENT 2

The main objective of Experiment 2 was to provide another test of the hypotheses using a brand in a different product category. For this experiment, the product category I chose was in the cybersecurity/antivirus product space. The experiment centered on the fictional brand CyberGuardian, which offers antivirus software designed to protect users from digital threats. The goal of this experiment was to build on Experiment 1 and understand how AI disclosure and type of content may influence brand perception particularly in a more technology-focused setting. Given that a high-tech brand was used, the experiment also sought to examine perceptions of the brand's technological capabilities.

Methods

Three-hundred and nine U.S. participants were recruited from the Connect platform and were compensated US\$0.60 for a 4-minute study. Participants who failed the attention check were excluded. After exclusions, 287 participants were included in the analyses (MAge = 39.00, SD = 11.83; 152 Male).

Participants first provided informed consent and were then presented with instructions for completing the study. Participants were then randomly assigned to one of four conditions in a 2 (content type: narrative vs. informational) by 2 (authorship disclosure: AI-generated vs. human-written) between-participants design. They were then shown a blog post about the same topic but different content type and disclosure conditions (i.e., informational post & AI-generated, informational post & human written, narrative post & AI-generated, and narrative post & human written).

The blog post was posted by a fictional brand, with the fictional brand being called CyberGuardian, which offers antivirus software designed to protect users from a wide range of digital threats, including malware, phishing, and other cybersecurity risks. The content of the posts and the brands were carefully crafted to be representative of typical blog posts in their respective categories. In developing the blog posts used in this study, ChatGPT (an AI language model developed by OpenAI) was used in order to assist in generating the content. The AI was used in order to create initial drafts of both the narrative and informational blog post types, which I then manually reviewed and edited in order to ensure they aligned with the intended purpose of the experiment and maintained consistency in tone and style.

Study 2 used a high-tech brand in the cybersecurity/antivirus software category in order to investigate AI disclosure in a context where technological sophistication is crucial, as opposed to the category in Study 1 where technological sophistication was barely required. Given that cybersecurity is an industry where consumers expect the highest reliability solutions, the perceived competence of a brand becomes critical. This helps to build on Study 1 by further

examining whether AI-generated content is viewed as enhancing or detracting from a brand's perceived competence and technological capabilities. This product category also allowed me to examine how consumers perceive AI use in a field where innovation and technological competence are essential.

The title of the informational blog post (both for AI disclosure and human-writer disclosure) read the following: "Choosing the Right Antivirus Software: Essential Insights for Advanced Protection." This blog post provided readers with critical and up-to-date advice on selecting the best antivirus software, and offered detailed insights into various features of antivirus software they should pay attention to when shopping around for one. The same content was presented in both the AI and human conditions.

The title of the storytelling blog post (both for AI disclosure and human-writer disclosure) read the following: "Personal Story: 'The Antivirus Rescue'". The blog post tells the story of a user whose laptop was saved from a severe malware attack by CyberGuardian's antivirus software. The narrative highlights the user's panic, the efficient resolution provided by the software, and the relief and peace of mind experienced after the successful malware removal. The same content was presented in both the AI and human conditions.

In the narrative condition, the blog post was a story told from a first-person perspective, of a user who had recently experienced a significant cyberattack on their personal computer. The post detailed the story of the user's fear and uncertainty during the attack, and their subsequent decision to make use of CyberGuardian's antivirus software. The story emphasized the effectiveness of CyberGuardian in quickly identifying and neutralizing the threat, ultimately

restoring the user's peace of mind and protecting their valuable data. In the AI condition, participants were provided with a disclosure at the top of the blog post that read: “This CyberGuardian blog post, generated by our advanced AI system, illustrates how CyberGuardian antivirus can be your digital lifesaver when you need it the most.” In the human condition, the disclosure instead read: “This CyberGuardian blog post, written by our relieved customer Alex Johnson, illustrates how CyberGuardian antivirus can be your digital lifesaver when you need it the most.”

In the informational condition, the blog post was designed to provide readers with detailed and practical advice on how to choose the best antivirus software to protect against cyber threats, highlighting the most important features consumers should look for in order to safeguard their digital security. In the AI condition, participants were provided with a disclosure at the top of the blog post that read: “This CyberGuardian blog post was generated by our advanced AI system to provide you with the most critical and up to date advice. In this blog post, our AI delves into the complexities of cyber threats to guide you in selecting the most effective antivirus software.” In the human condition, the disclosure instead read: “This CyberGuardian blog post was written by our cybersecurity expert, Alex Johnson, to bring you the most critical and up-to-date advice. In this blog post, Alex delves into the complexities of cyber threats to guide you in selecting the most effective antivirus software.” (Appendix A provides the stimuli used in this study.)

The content was kept constant across posts which belonged to the same blog post type but differed in disclosure nature, in order to achieve consistency across our experiment

and its results.

Thus, the overall design of the experiment was a 2 (content type) x 2 (authorship disclosure) mixed design.

The same manipulation and attention checks used for experiment 1 were used for experiment 2, as well as the same measures for warmth, competence, and brand attitude.

Participants were also asked to report their perceptions of CyberGuardian's technological capabilities using a 7-point scale (1 = Strongly Disagree to 7 = Strongly Agree) with the following statements: "CyberGuardian is at the forefront of technology in its industry", "CyberGuardian's technological capabilities are highly sophisticated and advanced", "CyberGuardian is a technologically innovative company", and "CyberGuardian uses cutting-edge technology to develop its products".

Participants were also asked to report their acceptance of AI technology using a 7-point scale (1 = Strongly Disagree to 7 = Strongly Agree) with the following statements: "I fear artificial intelligence," "I trust artificial intelligence," "Artificial intelligence will destroy humankind," "Artificial intelligence will benefit humankind," and "Artificial intelligence will cause many job losses."

Finally, participants were offered an open ended section where they could share any comments they may have had about this blog post study, and were then asked standard demographics questions (age and gender).

Results and Discussion


Manipulation Check: Usefulness. A two-way ANOVA was conducted to examine the effect of content type and disclosure on usefulness. The main effect of content type was significant, $F(1, 283) = 43.10, p < .001, \eta^2 = .13$. The narrative blog post rates ($M = 4.97, SD = 1.37$) higher on usefulness compared to the technical post ($M = 5.89, SD = 1.13$). The main effect of disclosure was also significant, $F(1, 283) = 10.3, p = .001, \eta^2 = .03$. Participants who were exposed to Human disclosure ($M = 5.60, SD = 1.21$) rated usefulness higher than those exposed to AI-generated disclosure ($M = 5.24, SD = 1.43$). The interaction between content type and disclosure was not significant, $F(1, 283) = 1.73, p = .189, \eta^2 = 0$.

Manipulation Check: Storytelling. A two-way ANOVA was conducted to examine the effect of content type and disclosure on storytelling. The main effect of content type was significant, $F(1, 283) = 318.31, p < .001, \eta^2 = .52$. Participants rated storytelling differently depending on the content type, with Emotive/Narrative content being rated higher ($M = 6.14, SD = 0.92$) than Technical content ($M = 3.22, SD = 1.70$). The main effect of disclosure was also significant, $F(1, 283) = 3.63, p = .164, \eta^2 = 0$. Participants who were exposed to Human disclosure ($M = 4.99, SD = 1.96$) rated storytelling higher than those exposed to AI-generated disclosure ($M = 4.45, SD = 2$). The interaction between content type and disclosure was not statistically significant, $F(1, 283) = 0, p = .957, \eta^2 = 0$.

Warmth. The two items measuring warmth were combined into an index ($r = 0.92$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on warmth. The main effect of content type was not significant, $F(1, 283) = 0.93, p = .334, \eta^2 = 0$.


The main effect of disclosure was significant, $F(1, 283) = 7.16, p = .008, \eta^2 = .02$. Warmth ratings between participants exposed to human disclosure ($M = 4.75, SD = 1.40$) were higher than those exposed to AI-generated disclosure ($M = 4.27, SD = 1.53$). There was a marginally significant interaction between content type and disclosure, $F(1, 283) = 3.30, p = .070, \eta^2 = .01$. Specifically, for Technical content, participants exposed to human disclosure ($M = 4.83, SD = 0.18$) rated warmth higher than those exposed to AI disclosure ($M = 4.05, SD = 0.16$), $F(1, 283) = 3.97, p = .047$. For Narrative content, participants exposed to AI disclosure ($M = 4.53, SD = 0.18$) rated warmth lower than those exposed to human disclosure ($M = 4.68, SD = 0.16$), $F(1, 283) = 0.35, p = .553$. See Figure 7.

Figure 7: Experiment 2 Brand Warmth Ratings



Competence. The two items measuring competence were combined into an index ($r = 0.87$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on competence. The main effect of content type was not significant, $F(1, 283) = 0.14, p = .709, \eta^2 = 0$. The main effect of disclosure was significant, $F(1, 283) = 8.41, p = .004, \eta^2 = .02$. Competence ratings between participants exposed to human disclosure ($M = 5.99, SD = 0.96$) were higher than those exposed to AI-generated disclosure ($M = 5.65, SD = 1.05$). The interaction between content type and disclosure was not significant, $F(1, 283) = 0.67, p = .419, \eta^2 = 0$. See Figure 8.

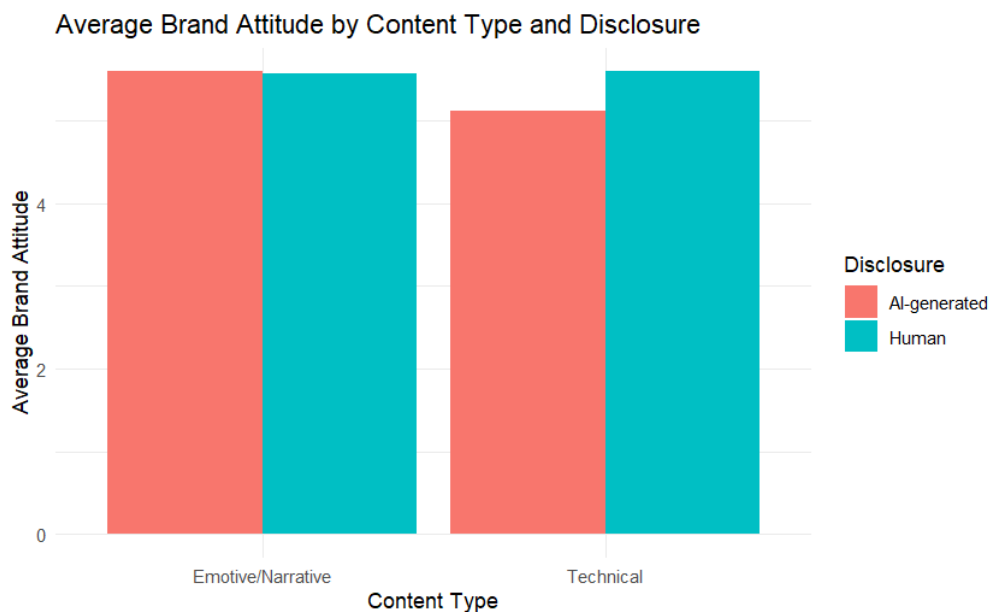
Figure 8: Experiment 2 Brand Competence Ratings



Brand Attitude. The items measuring brand attitude were combined into an index (Cronbach's $\alpha = 0.95$). A two-way ANOVA was conducted to examine the effect of content type

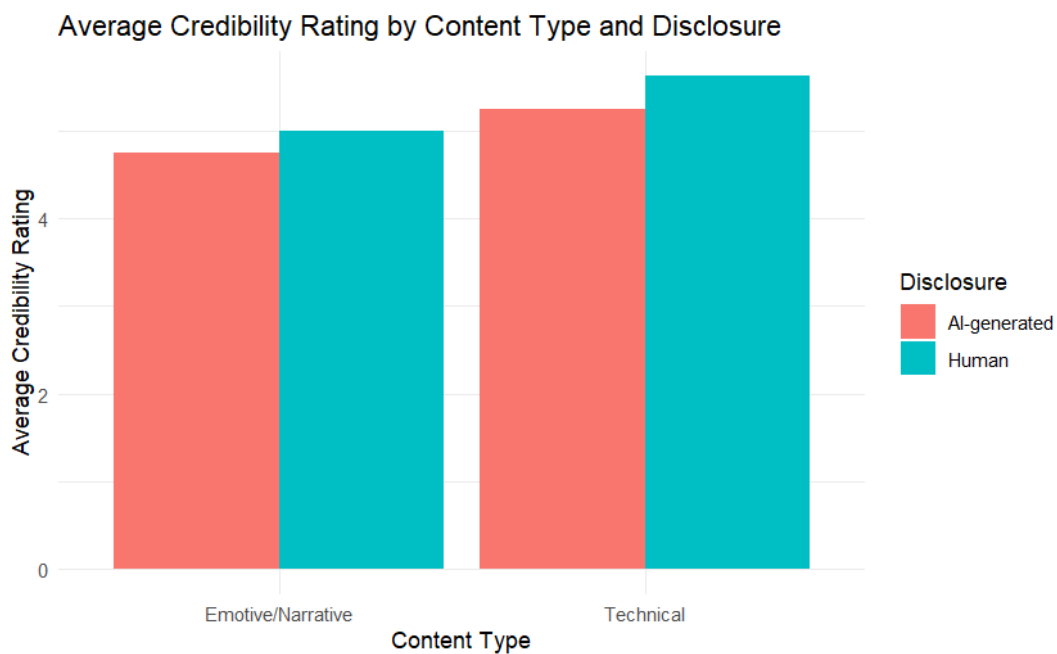
and disclosure on brand attitude. The main effect of content type was not significant, $F(1, 283) = 2.46, p = .118, \eta^2 = 0$. Similarly, the main effect of disclosure was not significant, $F(1, 283) = 2.55, p = .111, \eta^2 = 0$. However, there was a marginally significant interaction between content type and disclosure, $F(1, 283) = 3.18, p = .075, \eta^2 = .01$. Specifically, for Technical content, participants exposed to human disclosure ($M = 5.60, SD = 0.15$) rated brand attitude higher than those exposed to AI disclosure ($M = 5.12, SD = 0.13$), $F(1, 283) = 5.75, p = .017$. For Narrative content, participants exposed to AI disclosure ($M = 5.59, SD = 0.14$) rated brand attitude similarly to those exposed to human disclosure ($M = 5.57, SD = 0.13$), $F(1, 283) = 0.23, p = .880$. See Figure 9.

Figure 9: Experiment 2 Brand Attitude Ratings



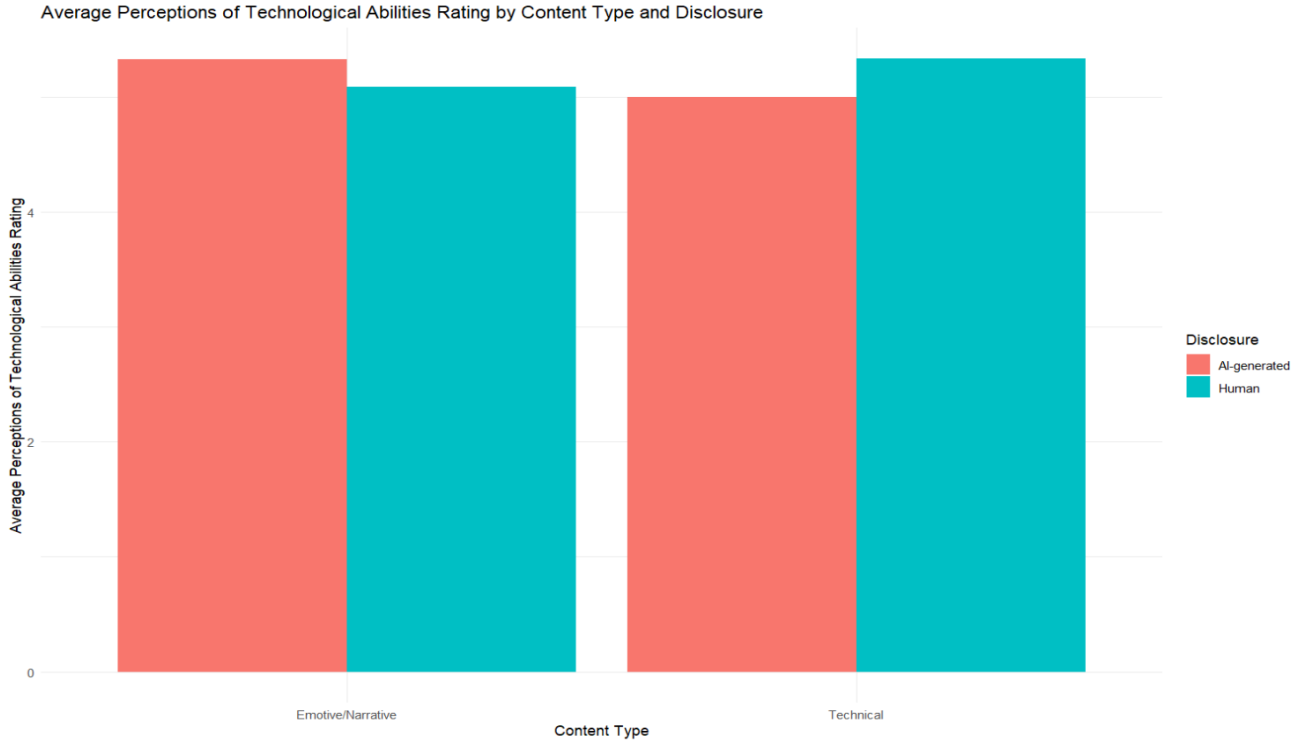
Credibility. The items measuring credibility were combined into an index (Cronbach's $\alpha = 0.91$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on credibility. The main effect of content type was significant, $F(1, 283) = 13.46, p < .001, \eta^2 = .04$. Participants rated credibility differently depending on the content type, with Emotive/Narrative content being rated lower ($M = 4.89, SD = 1.50$) than Technical content ($M = 5.41, SD = 0.99$). The main effect of disclosure was also significant, $F(1, 283) = 4.34, p = .038, \eta^2 = .01$. Participants who were exposed to Human disclosure ($M = 5.27, SD = 1.37$) rated credibility higher than those exposed to AI-generated disclosure ($M = 5.02, SD = 1.23$). There was no significant interaction between content type and disclosure, $F(1, 283) = 0.20, p = .655, \eta^2 = 0$. See Figure 10.

Figure 10: Experiment 2 Brand Credibility Ratings



Perceptions of Technological Capabilities. The items measuring Perceptions of Technological Capabilities were combined into an index (Cronbach's $\alpha = 0.94$). A two-way ANOVA was conducted to examine the effect of content type and disclosure on the perception of tech. The main effect of content type was not significant, $F(1, 283) = 0.10, p = .749, \eta^2 = 0$. The main effect of disclosure was also not significant, $F(1, 283) = 0.16, p = .683, \eta^2 = 0$. However, there was a significant interaction between content type and disclosure, $F(1, 283) = 4.95, p = .027, \eta^2 = .01$. Specifically, for Technical content, participants exposed to human disclosure ($M = 5.33, SD = 0.13$) rated perceptions of technological capabilities higher than those exposed to AI disclosure ($M = 4.99, SD = 0.12$), $F(1, 283) = 3.31, p = .070$. For Narrative content, participants exposed to AI disclosure ($M = 5.32, SD = 0.13$) did not perceive technological capabilities significantly differently than those exposed to human disclosure ($M = 5.09, SD = 0.12$), $F(1, 283) = 1.77, p = .184$. See Figure 11.

Figure 11: Experiment 2 Brand Technological Abilities Ratings



Discussion

In Experiment 2, participants once again rated human-written content as warmer and more competent than AI-generated content. However, unlike Experiment 1, content type and disclosure interacted to affect perceived warmth, brand attitude, and perceptions of technological capabilities. The moderation effect of content type suggests that the nature of the content (narrative vs. informational) plays a role in how AI disclosure influences consumer perceptions. However, the moderation did not operate as expected. Specifically, for the key dimension of warmth, AI disclosure decreased perceived brand warmth more for informational content rather than narrative content, contrary to what was hypothesized. I discuss this point further in the General Discussion. AI Acceptance was not further explored.

GENERAL DISCUSSION

This thesis explored how AI-generated versus human-written blog post content influences consumer perceptions of brand warmth and competence across two experiments. In Experiment 1, participants perceived content disclosed as human-written to be warmer and more competent. The type of content (narrative vs. informational) did not significantly alter perceptions of warmth or competence. Human-written content was also perceived as more credible, improved brand attitude, increased purchase intentions, and fostered a stronger sense of connection with consumers. Content type did not moderate these outcomes, reinforcing the importance of human authorship in shaping positive consumer perceptions.

Experiment 2 further examined these effects. Content disclosed as human-written was again perceived as warmer and more competent than content disclosed as AI-generated. Contrary to experiment 1, content type did moderate the effect of disclosure for warmth. However, the moderation did not operate as expected, with AI disclosure leading to *lower* brand warmth for the informational blog than for the narrative blog. A moderation of content type also occurred for brand attitude and technological capabilities, with similar patterns as for warmth.

Overall, both Experiment 1 and 2 suggest that human authorship is generally preferred, but initial evidence suggests this effect may be moderated under some conditions.

Limitations and Future Directions for Research

One notable difference between Study 1 (SerenitySnooze) and Study 2 (CyberGuardian) is the presence of interaction effects between content type and disclosure in Study 2 for warmth and brand attitude. This interaction was not observed in Study 1, where content type did not significantly moderate any outcomes. Study 2 also observed a significant moderation effect for content type on perceptions of technological competence based on the authorship disclosure. The stronger moderation effects found in Study 2 may be explained by the difference in product categories examined. In study 1, a product category more related to personal wellness was examined (bed mattresses), while in study 2, a product category much more reliant on technology and innovation was examined (antivirus software). These differences in results suggest that the impact of AI-disclosure versus human-disclosure may vary depending on the product category chosen, where certain attributes become more pronounced in technical or expertise-driven products.

In Study 2, the moderation of content type for AI disclosure did not align with my original hypothesis, where I expected AI disclosure to more negatively affect brand warmth in emotive/narrative content than in informational content. Instead, Study 2 results showed that AI disclosure actually had a more negative impact on informational content than it did emotive/narrative content. One possible explanation is that in a highly technical product category like antivirus software, consumers inherently expect informational content to be as accurate as possible, and supported by real life human expertise. When such content about a sensitive and important topic such as cybersecurity is disclosed as AI-generated, this may raise concerns about the depth of understanding and the trustworthiness of the information being presented to the consumer, which may ultimately cause a more pronounced decline in perceived warmth. By contrast, consumers may not have strong expectations for warmth when it comes to narrative content in a highly technical product category like antivirus software.

There are a few limitations my studies faced, which future research could potentially address. First of all, the experiments focused only on two types of content: narrative and informational. Even though these two content types are widespread in marketing communications, they do not represent the full spectrum of content types available in this field. A significant limitation of the study was that no moderating effect of content type on the relationship between AI disclosure and consumer perceptions of brand warmth and competence was found. This could suggest that content type does not significantly alter how people perceive disclosures, or it may indicate that the study did not capture the right content types that would moderate these effects. Future research should explore a broader range of content types beyond

solely blog posts, such as social media posts and video content, to determine if certain formats of content are more susceptible to variations in disclosure.

Second, this research focused on niche product categories, antivirus software and mattresses, which is not an accurate representation of the many different industries in which AI-generated content can leave a mark. Consumer perceptions of brand warmth and competence may not be uniform across all product categories - instead, they may vary significantly across different product categories. Future studies should examine a wider range of products and industries in order to better understand how AI-generated content is perceived in different contexts.

Third, the studies used a uniform framing of disclosure, where content was presented as either AI-generated or human-written. However, the way disclosure was set up in these studies does not represent all possibilities, and the way in which AI disclosure is communicated to consumers could significantly impact their perceptions of brand warmth and competence. Different framings of AI disclosure, whether via alternate wording (example: highlighting the collaborative role of AI along with expert human oversight) can alter consumer responses. Hence, future research should investigate how different disclosure strategies might influence consumer perceptions of brand warmth and competence. By doing so, future research can potentially uncover ways to mitigate negative biases against AI-generated content.

Fourth, the research did not account for the possible impact of social and cultural norms on consumer perceptions of AI-generated content. Cultural attitudes towards technology, as well

as societal values, can vary widely across different regions and demographic groups. These factors may influence how consumers perceive a brand's warmth and competence when told that the content they just read was AI-generated. Thus, future research should explore how social and cultural contexts influence responses to AI disclosure. To do so, future research could potentially conduct cross-cultural studies to identify differences in AI acceptance across different regions and demographic groups, and analyze how these differences impact brand perceptions. This future research could come up with managerial implications especially useful for global companies with global marketing strategies, in order to make sure that these marketing strategies are specifically tailored to the needs of relative consumer bases.

Fifth, while the brands included in this thesis's experiments were fictional, a strategy implemented in order to control for pre-existing biases towards the brand that participants may have otherwise had, future research could investigate how consumers' familiarity with real and well-known brands they already interact with in their day to day lives influences their perceptions of AI-generated content. Future research may find that consumers might react differently to AI disclosures from brands they already recognize and trust, as compared to unfamiliar or new brands they do not recognize and trust.

Finally, while the choice of a mattress brand for Study 1 and a high-tech cybersecurity brand for Study 2 provided valuable insights into consumer perceptions of AI-generated content across different contexts, the use of these particular brands and product categories also introduces certain limitations that future research should address. The selection of a mattress brand in Study 1 and a Cybersecurity brand in Study 2 focus on products that may lend

themselves more to the informational blog post type rather than the narrative blog post type. It is possible this may have impacted perceptions of brand warmth and competence. Future research should consider looking into other product categories where the emotional connection is much more crucial for the brand-consumer relationship, as these categories lend themselves much more to a storytelling approach.

Practical Contributions

My research also offers some important managerial implications for marketing professionals. First of all, my findings indicate that human-written content is perceived as warmer and possibly at times more competent compared to AI-generated content. Brands should consider the potential impacts of AI authorship on perceptions of brand warmth and competence. In the context of this thesis, the context is content marketing.

Second, findings show that consumers may not always perceive AI-generated content as authentic as human-written content. However, the efficiency of AI for content marketing purposes is clear. For this reason, brands should consider adopting a balanced approach that offers the best of both worlds for consumers and the business alike, where they would still use AI, but only do so along with human oversight to maintain authenticity. For example, AI could be used for initial content generation, only to be followed by expert human revision and refinement in later steps in order to ensure that the content meets the brand's content quality requirements and aligns with their values and those of their customers. By adopting such a hybrid approach, brands could possibly benefit from content efficiency offered by AI, without compromising as much on consumer perceptions of warmth and competence.

Third, while integrating AI use into content creation offers a plethora of advantages in terms of efficiency and scalability, it also raises notable ethical considerations that managers should take into consideration. As brands increasingly rely on AI to generate marketing content, this raises the ethical question of transparency with their consumers. It can be argued that consumers have the right to know whether the content they are engaging with was partially or entirely generated by AI, or partially or entirely generated by a human being. Lack of transparent or truthful disclosure may constitute deceptive behavior, which is not only unethical, but can also potentially damage the brand's reputation and consumer trust. Thus, the question of accurate and truthful disclosure deserves careful consideration.

THEORETICAL CONTRIBUTIONS

This thesis makes several key theoretical contributions to existing literature on AI in marketing, especially the niche field of content marketing, and specifically when it comes to consumer perceptions of AI-generated content and its impact on brand perception.

First, this research helps fill a crucial gap in existing literature by examining how AI disclosure affects consumer perceptions of brand warmth and competence. While previous studies have extensively examined AI's broader implications in the greater field of marketing, there has been very limited insight on how AI disclosure influences these specific brand attributes, especially in the field of content marketing.

Second, this thesis hypothesizes and finds preliminary evidence that content type (narrative vs. informational) may moderate the impact of generative AI use on brand perceptions.

Given my findings which suggest that the negative effects of AI disclosure on perceptions of brand warmth are more pronounced in informational content as compared to narrative content, this adds to existing literature by suggesting that the context in which AI-generated content is presented plays an important role in shaping consumer perceptions. By identifying content type as a potential moderator, this thesis opens the door for a more detailed understanding, in future research, of the different contexts in which AI-generated content influences consumer attitudes.

CONCLUSION

In summary, this thesis contributes to our understanding of how AI-generated content and its disclosure impact consumer perceptions of brand warmth and competence. Across two experiments, the findings reveal that AI disclosure generally leads to lower perceptions of warmth and competence, with content type potentially playing an important moderating role. While further research is required to explore these effects in more depth, this research offers valuable insights into the strategic considerations that brands should consider making when integrating AI into their content marketing strategies.

REFERENCES:

- Allioui, H., & Mourdi, Y. (2023). Unleashing the potential of AI: Investigating cutting-edge technologies that are transforming businesses. *International Journal of Computer Engineering and Data Science (IJCEDS)*, 3(2), 1-12.
- 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.
- 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, 103790.
- Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809-825.
- Chu, H., & Liu, S. (2023). Can AI tell good stories? Narrative transportation and persuasion with ChatGPT.
- 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.
- Escalas, J. E. (2004). Narrative processing: Building consumer connections to brands. *Journal of Consumer Psychology*, 14(1-2), 168-180.

Fisher, W. R. (1984). Narration as a human communication paradigm: The case of public moral argument. *Communication Monographs*, 51(1), 1-22.

Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77-83.

George, A. S., & George, A. H. (2023). A review of ChatGPT AI's impact on several business sectors. *Partners Universal International Innovation Journal*, 1(1), 9-23.

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627-660.

Gonçalves, A. R., Pinto, D. C., Rita, P., & Pires, T. (2023). Artificial intelligence and its ethical implications for marketing. *Emerging Science Journal*, 7(2), 313-327.

Green, M. C., & Brock, T. C. (2000). The role of transportation in the persuasiveness of public narratives. *Journal of Personality and Social Psychology*, 79(5), 701.

Kervyn, N., Fiske, S. T., & Malone, C. (2022). Social perception of brands: Warmth and competence define images of both brands and social groups. *Consumer Psychology Review*, 5(1), 51-68.

Kim, J., Giroux, M., & Lee, J. C. (2021). When do you trust AI? The effect of number presentation detail on consumer trust and acceptance of AI recommendations. *Psychology & Marketing*, 38(7), 1140-1155.

- Kose, U., & Sert, S. (2016). Intelligent content marketing with artificial intelligence. In *International Conference of Scientific Cooperation for Future* (No. 837-43).
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90-103.
- Madhavan, P., & Wiegmann, D. A. (2007). Effects of information source, pedigree, and reliability on operator interaction with decision support systems. *Human Factors*, 49(5), 773-785.
- Palumbo, S., & Edelman, D. (2023). What smart companies know about integrating AI. *Harvard Business Review*, 101(7-8), 116-125.
- Paschen, J., Kietzmann, J., & Kietzmann, T. C. (2019). Artificial intelligence (AI) and its implications for market knowledge in B2B marketing. *Journal of Business & Industrial Marketing*, 34(7), 1410-1419.
- Pulizzi, J. (2012). The rise of storytelling as the new marketing. *Publishing Research Quarterly*, 28(2), 116-123.
- Roggeveen, A. L., & Sethuraman, R. (2020). Customer-interfacing retail technologies in 2020 & beyond: An integrative framework and research directions. *Journal of Retailing*, 96(3), 299-309.

Salvagno, M., Taccone, F. S., & Gerli, A. G. (2023). Can artificial intelligence help for scientific writing?. *Critical Care*, 27(1), 75.

Van Esch, P., & Stewart Black, J. (2021). Artificial intelligence (AI): Revolutionizing digital marketing. *Australasian Marketing Journal*, 29(3), 199-203.

Wang, K., Wang, E. T. G., & Farn, C. K. (2009). Influence of web advertising strategies, consumer goal-directedness, and consumer involvement on web advertising effectiveness. *International Journal of Electronic Commerce*, 13(4), 67-96.

APPENDIX A: EXPERIMENTAL STIMULI

Choosing the Right Mattress: Your Guide to a Better Night's Sleep

Disclosure: *This SerenitySnooze blog post was generated by our specialized AI system to ensure you receive the most up-to-date and well-rounded advice. Our AI blends the latest sleep research with practical insights to guide you in choosing the perfect mattress.*

Welcome to SerenitySnooze, where we use cutting-edge AI to revolutionize your sleep experience! Harnessing advanced algorithms and the latest sleep science, our AI system offers you tailored advice to find the perfect mattress for superb sleep!

Dive into Mattress Types: It's like choosing your favorite dessert. Memory foam mattresses hug your body just right. Prefer something firmer? Go for innerspring mattresses. Or, get the best of both with a hybrid mattress.

Find Your Perfect Firmness: Think about how you sleep. If you're a side sleeper, a softer mattress is like a gentle hug for your hips and shoulders. Back sleeper? A medium-firm mattress supports your spine just right. And for stomach sleepers, a firm mattress keeps your spine happy.

Get Even Weight Distribution: Imagine a mattress that fits you perfectly. Memory foam does that, spreading your weight evenly. Latex is also a champ for support. But innerspring mattresses? They're firm, though they might miss a spot or two.

Stay Cool All Night: If you get too warm at night, look for mattresses with cooling gels and breathable materials. It's like having a built-in AC for your bed.

Smart Shopping with Trials and Warranties: It's like test driving a car. Pick brands that let you try the mattress and offer good warranties. It's about making sure you love your mattress for years.

Keep these AI-crafted tips in mind, and you're well on your way to discovering your ideal mattress. Better sleep? With our AI's guidance, it's within your reach!

We hope you found this SerenitySnooze blog post from our specialized AI system helpful! We're always excited to share insights that can enhance your sleep quality. Thanks for reading!

Choosing the Right Mattress: Your Guide to a Better Night's Sleep

Disclosure: *This SerenitySnooze blog post was written by our mattress specialist Miles Lindsay to ensure you receive the most up-to-date and well-rounded advice. Miles blends the latest sleep research with practical insights to guide you in choosing the perfect mattress.*

At SerenitySnooze, we're all about transforming your sleep into something amazing. Miles Lindsay, our top-notch mattress whiz, is here to guide you through picking the best mattress for awesome sleep!

Dive into Mattress Types: It's like choosing your favorite dessert. Memory foam mattresses hug your body just right. Prefer something firmer? Go for innerspring mattresses. Or, get the best of both with a hybrid mattress.

Find Your Perfect Firmness: Think about how you sleep. If you're a side sleeper, a softer mattress is like a gentle hug for your hips and shoulders. Back sleeper? A medium-firm mattress supports your spine just right. And for stomach sleepers, a firm mattress keeps your spine happy.

Get Even Weight Distribution: Imagine a mattress that fits you perfectly. Memory foam does that, spreading your weight evenly. Latex is also a champ for support. But innerspring mattresses? They're firm, though they might miss a spot or two.

Stay Cool All Night: If you get too warm at night, look for mattresses with cooling gels and breathable materials. It's like having a built-in AC for your bed.

Smart Shopping with Trials and Warranties: It's like test driving a car. Pick brands that let you try the mattress and offer good warranties. It's about making sure you love your mattress for years.

Keep these tips from Miles in mind, and you're on track to find that dream mattress. Better sleep? It's just a mattress away!

We hope you found this SerenitySnooze blog post from our mattress specialist Miles Lindsay helpful! We're always excited to share insights that can enhance your sleep quality. Thanks for reading!

Personal Story: "The Mattress That Changed Everything"

Disclosure: **This SerenitySnooze blog post, written by our specialized AI system, highlights how finding the perfect mattress for your needs can significantly improve daily life and well-being.**

So, I've got to tell you about this new mattress I bought from SerenitySnooze.

For months, I'd been waking up feeling like I've been in a boxing match with Mike Tyson. Turns out my old mattress was basically run to the ground.

I finally decided to take the plunge and invest in something decent, because you know, my back can't keep up!

I didn't go for one of those high-end mattresses that have you paying an arm and a leg. Instead, I found this mid-range one online with great reviews – it ticked my boxes as far as comfort, support, all that jazz goes.

On arrival, I was like a kid checking on gifts under the Christmas tree! I unpacked the box, wrestled it a bit onto my bed frame, cut away the plastic, left it to expand, and left it alone for the day.

First night on it? Unreal. It was like sleeping on cloud nine.

I didn't realize how bad my old mattress was until I woke up the next morning feeling like I've actually slept! No aches, no pains, and for the first time in ages, I got up feeling energized even before my alarm went off.

I've even noticed changes ever since: better mood, increased energy, and heck, even my coffee started tasting better! It's wild how much of a difference in your life a good mattress can make.

Now whenever someone complains to me about their sleep, I'll go: "Dude, check your mattress!"

We hope you found this SerenitySnooze blog post from our specialized AI system helpful! We're always excited to share insights that can enhance your sleep quality. Thanks for reading!

Personal Story: "The Mattress That Changed Everything"

Disclosure: **This SerenitySnooze blog post, written by our satisfied customer Miles Lindsay, highlights how finding the perfect mattress for your needs can significantly improve daily life and well-being.**

So, I've got to tell you about this new mattress I bought from SerenitySnooze.

For months, I'd been waking up feeling like I've been in a boxing match with Mike Tyson. Turns out my old mattress was basically run to the ground.

I finally decided to take the plunge and invest in something decent, because you know, my back can't keep up!

I didn't go for one of those high-end mattresses that have you paying an arm and a leg. Instead, I found this mid-range one online with great reviews – it ticked my boxes as far as comfort, support, all that jazz goes.

On arrival, I was like a kid checking on gifts under the Christmas tree! I unpacked the box, wrestled it a bit onto my bed frame, cut away the plastic, left it to expand, and left it alone for the day.

First night on it? Unreal. It was like sleeping on cloud nine.

I didn't realize how bad my old mattress was until I woke up the next morning feeling like I've actually slept! No aches, no pains, and for the first time in ages, I got up feeling energized even before my alarm went off.

I've even noticed changes ever since: better mood, increased energy, and heck, even my coffee started tasting better! It's wild how much of a difference in your life a good mattress can make.

Now whenever someone complains to me about their sleep, I'll go: "Dude, check your mattress!"

We hope you found this SerenitySnooze blog post from our satisfied customer Miles Lindsay helpful! We're always excited to share insights that can enhance your sleep quality. Thanks for reading!

Choosing the Right Antivirus Software: Essential Insights for Advanced Protection

Disclosure: *This CyberGuardian blog post was generated by our advanced AI system to provide you with the most critical and up to date advice. In this blog post, our AI delves into the complexities of cyber threats to guide you in selecting the most effective antivirus software.*

Welcome to CyberGuardian, where our cutting-edge AI is dedicated to fortifying your digital security! Utilizing sophisticated algorithms and deep cybersecurity knowledge, our AI system offers you personalized recommendations to choose the best antivirus defense for your needs.

Advanced Threat Detection: Opt for antivirus software that goes beyond basic malware scanning. Look for heuristic analysis, behavior-based detection, and AI-driven threat intelligence to identify and neutralize sophisticated attacks.

Comprehensive Malware Removal: Ensure your antivirus can thoroughly eradicate complex malware, including rootkits, spyware, and ransomware. It should provide a detailed cleanup process without leaving any traces that could lead to reinfection.

Low System Impact with High Protection: Balance is key. Seek software that offers robust protection without significant performance degradation. Consider solutions with lightweight agents and cloud-based scanning.

Proactive Exploit Prevention: Choose antivirus software with exploit protection capabilities. It should shield vulnerable software and prevent zero-day attacks by monitoring and blocking suspicious activity in real time.

Multi-Layered Phishing Defense: Phishing is a prevalent threat. Your antivirus should include advanced anti-phishing features, such as real-time link analysis, web reputation scoring, and integration with email clients for comprehensive protection.

Independent Testing and Certification: Look for antivirus software that has undergone rigorous independent testing and received certifications from reputable organizations like AV-TEST or AV-Comparatives, confirming its effectiveness against advanced threats.

Armed with these AI-driven insights, you're equipped to choose antivirus software that not only meets your needs but also provides advanced protection in an ever-evolving threat landscape. In the digital realm, staying ahead of threats is paramount!

We hope this CyberGuardian blog post, generated by our advanced AI system, has provided you with valuable information for making an informed decision about your antivirus software. We're dedicated to sharing expertise that enhances your digital security. Thanks for reading!

Choosing the Right Antivirus Software: Essential Insights for Advanced Protection

Disclosure: *This CyberGuardian blog post was written by our cybersecurity expert, Alex Johnson, to bring you the most critical and up-to-date advice. In this blog post, Alex delves into the complexities of cyber threats to guide you in selecting the most effective antivirus software.*

At CyberGuardian, we prioritize your digital security. Alex Johnson, our cybersecurity expert, is here to help you navigate the nuances of antivirus software and ensure you choose the best defense for your needs!

Advanced Threat Detection: Opt for antivirus software that goes beyond basic malware scanning. Look for heuristic analysis, behavior-based detection, and AI-driven threat intelligence to identify and neutralize sophisticated attacks.

Comprehensive Malware Removal: Ensure your antivirus can thoroughly eradicate complex malware, including rootkits, spyware, and ransomware. It should provide a detailed cleanup process without leaving any traces that could lead to reinfection.

Low System Impact with High Protection: Balance is key. Seek software that offers robust protection without significant performance degradation. Consider solutions with lightweight agents and cloud-based scanning.

Proactive Exploit Prevention: Choose antivirus software with exploit protection capabilities. It should shield vulnerable software and prevent zero-day attacks by monitoring and blocking suspicious activity in real time.

Multi-Layered Phishing Defense: Phishing is a prevalent threat. Your antivirus should include advanced anti-phishing features, such as real-time link analysis, web reputation scoring, and integration with email clients for comprehensive protection.

Independent Testing and Certification: Look for antivirus software that has undergone rigorous independent testing and received certifications from reputable organizations like AV-TEST or AV-Comparatives, confirming its effectiveness against advanced threats.

Armed with these insights from Alex, you're equipped to choose antivirus software that not only meets your needs but also provides advanced protection in an ever-evolving threat landscape. In the digital realm, staying ahead of threats is paramount!

We hope this CyberGuardian blog post from our cybersecurity expert, Alex Johnson, has provided you with valuable information for making an informed decision about your antivirus software. We're dedicated to sharing expertise that enhances your digital security. Thanks for reading!

Personal Story: "The Antivirus Rescue"

Disclosure: *This CyberGuardian blog post, generated by our advanced AI system, illustrates how CyberGuardian antivirus can be your digital lifesaver when you need it the most.*

There I was, thinking I had all my digital bases covered. Little did I know, a seemingly innocent download was about to turn my digital world upside down.

My laptop started acting up, running slower than ever, and strange pop-ups began appearing out of nowhere. It hit me – I had malware.

At first, I panicked. Years of documents, photos, and personal projects were at risk. The antivirus I had was basic and clearly not up to the task. That's when I remembered reading about CyberGuardian's antivirus software, hailed for its advanced malware detection and removal capabilities.

On a separate, clean device, I purchased CyberGuardian and created a bootable USB drive with the antivirus installer, following the instructions provided on their website. The software sprang to life, scanning my system with precision.

Watching the progress bar became my beacon of hope. CyberGuardian's antivirus didn't disappoint. It identified and quarantined the malware that had taken my laptop hostage.

But it didn't stop there; it repaired the damage done to my system files, bringing my laptop back from the brink. The whole process was smoother and quicker than I could have hoped for.

Relief washed over me as I saw my laptop return to its normal speed, free from pop-ups and lag and security threats. CyberGuardian had not only saved my digital life but also restored my peace of mind.

From that day forward, I knew I had a reliable guardian in the digital realm. Now, I make sure to tell everyone about the importance of having a strong antivirus solution.

CyberGuardian didn't just clean up the mess – it prevented future invasions, keeping me and my digital world safe.

We hope this CyberGuardian blog post from our advanced AI system has highlighted the critical role of effective antivirus software in combating malware. Protecting your digital environment is imperative, and with CyberGuardian, you're in safe hands. Thanks for reading!

Personal Story: "The Antivirus Rescue"

Disclosure: *This CyberGuardian blog post, written by our relieved customer Alex Johnson, illustrates how CyberGuardian antivirus can be your digital lifesaver when you need it the most.*

There I was, thinking I had all my digital bases covered. Little did I know, a seemingly innocent download was about to turn my digital world upside down.

My laptop started acting up, running slower than ever, and strange pop-ups began appearing out of nowhere. It hit me – I had malware.

At first, I panicked. Years of documents, photos, and personal projects were at risk. The antivirus I had was basic and clearly not up to the task. That's when I remembered reading about CyberGuardian's antivirus software, hailed for its advanced malware detection and removal capabilities.

On a separate, clean device, I purchased CyberGuardian and created a bootable USB drive with the antivirus installer, following the instructions provided on their website. The software sprang to life, scanning my system with precision.

Watching the progress bar became my beacon of hope. CyberGuardian's antivirus didn't disappoint. It identified and quarantined the malware that had taken my laptop hostage.

But it didn't stop there; it repaired the damage done to my system files, bringing my laptop back from the brink. The whole process was smoother and quicker than I could have hoped for.

Relief washed over me as I saw my laptop return to its normal speed, free from pop-ups and lag and security threats. CyberGuardian had not only saved my digital life but also restored my peace of mind.

From that day forward, I knew I had a reliable guardian in the digital realm. Now, I make sure to tell everyone about the importance of having a strong antivirus solution.

CyberGuardian didn't just clean up the mess – it prevented future invasions, keeping me and my digital world safe.

We hope this CyberGuardian blog post from our customer Alex Johnson has highlighted the critical role of effective antivirus software in combating malware. Protecting your digital environment is imperative, and with CyberGuardian, you're in safe hands. Thanks for reading!