

Louvain School of Management

How consumers respond to AI-generated advertising visuals: the moderating role of disclosure and product category

Authors: Alexia-Ilinca Mateiciuc and Elsa Vannetelbosch
Supervisor: Vincenzo Verardi
Academic year 2025-2026
Dissertation for the master in Business Management
Master subject and focus: Marketing, Consumer Behaviour, Generative Artificial Intelligence
Daytime schedule

Foreword

First of all, we would like to sincerely thank Vincenzo Verardi, our thesis supervisor, for his guidance and mentorship throughout the process of writing this thesis. He always provided us with quick and reactive responses, constructive criticism and valuable experience. We are grateful for this availability, support, and encouragements, which motivated us to give our best.

Secondly, we would like to thank those that participated in the realization of the thesis, and therefore contributed with the collection of responses to our questionnaire. Therefore, we are deeply grateful for our families, especially Ioana, Ana and Vincent, that offered us their precious advices and unconditional support during these last few months.

Declaration regarding AI tool usage

1. STATEMENTS ON YOUR USE OF GENERATIVE AI

Confirm the following statements:

Statements on responsible use of generative AI	Confirmation: Yes
The content of our master's thesis is our original input, even if we used generative AI to help prepare it.	☒
We master the theories, tools and methods that I/we use for the thesis.	☒
We verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc.	☒
We acknowledge that we master the final output and that we are able to explain it and answer questions about it.	☒

2. SIGNATURE

Sign your declaration:

	Author last name, first name(s)	Date	Signature
1	Mateiciuc, Alexia-Ilinca	27.05.2026	
2	Vannetelbosch, Elsa	27.05.2026	

Table of Contents

<i>I.</i>	<i>Introduction.....</i>	<i>1</i>
1.	Research Gap and Objectives.....	1
2.	Structure of the Thesis	3
<i>II.</i>	<i>Literature review</i>	<i>4</i>
1.	The Modern Advertising Context: The Integration and Impact of Artificial Intelligence	4
2.	Consumer Information Processing: How Consumers Perceive Advertisements	6
3.	The Core Conflict: The AI Penalty vs. Algorithm Appreciation	7
3.1.	The AI Penalty.....	7
3.2.	Algorithm Appreciation	7
3.3.	Key Dimensions of Consumer Evaluation	8
4.	The Moderating Role of Product Type: Hedonic vs. Utilitarian	9
5.	Disclosure Integrity: Truthful vs. Misleading	10
6.	Individual Differences and Exploratory Findings	11
7.	Strategic Trade-offs and the Legal Mandate.....	12
8.	Conclusion and Identification of Research Gaps.....	14
9.	Hypothesis Formulation	14
9.1.	Hypothesis 1: The Influence of Disclosure and the "AI Penalty"	14
9.2.	Hypothesis 2: The Moderating Role of Product Type.....	15
9.3.	Hypothesis 3: Deception, Revelation, and Behavioral Shifts	15
<i>III.</i>	<i>Methodology.....</i>	<i>16</i>
1.	Study type and background	16
1.1.	Type of study.....	16
1.2.	Independent & Dependent variables	16
1.3.	Research environment.....	17
1.4.	Group structure.....	18
1.5.	Sample size and characteristics.....	19
1.6.	Filler ads.....	19
1.7.	Tools used and justified choice	20
2.	Procedure of generation of images	20
3.	Questionnaire structure	22

3.1.	Onboarding and Ethics	22
3.2.	The Randomization Mechanism	23
3.3.	Groups	23
3.4.	Questions.....	24
3.5.	Label Check	26
3.6.	Utilitarian vs Hedonic perception of products Check.....	26
3.7.	Demographics	27
3.8.	GAAIS.....	27
3.9.	Frequency of purchase	28
IV.	<i>Analysis of Results</i>	29
1.	Final data gathered.....	29
2.	Preparation of the data gathered	29
3.	Stata analysis	30
3.1.	Data cleaning.....	30
3.2.	Identifying Outliers	30
3.3.	Sample Description	31
3.4.	Groups Description	33
3.5.	Scale Reliability	33
3.6.	Randomization Check	34
3.7.	Normality of the dependent variables	35
3.8.	Homogeneity of Variances (Homoscedasticity).....	36
4.	Tests of the hypothesis	38
4.1.	Hypothesis 1 (H1)	38
4.2.	Hypothesis 2 (H2)	40
4.3.	Hypothesis 3 (H3)	42
5.	Exploratory tests.....	45
5.1.	Hypothesis 1.....	45
5.2.	Hypothesis 2.....	48
5.3.	Hypothesis 3.....	49
V.	<i>Discussion</i>.....	52
1.	AI Disclosure and the Labeling Paradox (H1)	52
1.1.	Absence of a General AI Penalty (H1a).....	52
1.2.	The Label Scrutiny Effect: Human Attribution Backfire (H1b)	52
1.3.	The Human Effort Heuristic (H1c)	53
2.	Product Type and Hedonic Vulnerability (H2)	53
2.1.	Absence of an Interaction Effect (H2a).....	53

2.2.	Hedonic Products (H2b) and Utilitarian Products (H2c)	53
3.	The Cost of Deception: Effects of Revelation (H3)	54
3.1.	The “Caught in a Lie” Effect (H3a)	54
3.2.	The Inconsistency Effect (H3b)	54
4.	Exploratory Evidence: Consumer Heterogeneity	54
4.1.	Education level and Algorithm Acceptance	54
4.2.	Age-Based Divergence.....	54
4.3.	Gender	55
4.4.	Attitudes Toward AI (GAAIS)	55
VI.	Conclusion.....	56
1.	Summary of Key Findings	56
2.	Limitations and future research	56
VII.	Bibliography.....	59
VIII.	Appendices.....	67
1.	Stata Graphs	67
1.1.	Demographics	67
1.2.	Normality	68
1.3.	Outliers.....	72
2.	Generated Images (Nano Banana Pro, Gemini 3 Pro Image).....	76
2.1.	Juice.....	76
2.2.	Chocolate.....	77
2.3.	Cereals.....	82
2.4.	Shampoo.....	85
2.5.	Soap.....	86
2.6.	Toothpaste	87
2.7.	Canned Food	90
3.	G*Power Calculus	92
4.	Questionnaire link (Google Forms) + Google sheets	92
5.	Codes + Results Stata	95
5.1.	Data cleaning.....	95
5.2.	Outliers.....	96
5.3.	Sample Description	96
5.4.	Groups Description	97
5.5.	Scale Reliability (Cronbach’s Alpha).....	97
5.6.	Composite mean scores.....	100

5.7.	Randomization Check	100
5.8.	Normality of the dependent variables (Jarque-Bera): H1 + H2	101
5.9.	Normality of the dependent variables (Shapiro-Wilk): H3	103
5.10.	Homogeneity of variances (Homoscedasticity)	104
5.11.	Hypothesis 1	105
5.12.	Hypothesis 2	106
5.13.	Hypothesis 3A	107
5.14.	Hypothesis 3B	108
5.15.	Exploratory test - Hypothesis 1A	109
5.16.	Exploratory test – Hypothesis 1B	112
5.17.	Exploratory test – Hypothesis 1C	115
5.18.	Exploratory test – Hypothesis 2	117
5.19.	Exploratory test – Hypothesis 3A	120
5.20.	Exploratory test – Hypothesis 3B	121

I. Introduction

Recent technological advancements have driven a rapid expansion of generative artificial intelligence (GAI) tools for visual content creation, with advertising emerging as a key domain due to their potential to reduce time and production costs compared to traditional methods. This shift from high-effort, time-consuming human craftsmanship to scalable, AI-generated content represents a cost-benefit trade that is difficult to argue against.

An example of this tension can be seen in two recent Christmas campaigns, which illustrate this contrast clearly. On one hand, Intermarché's campaign relied on traditional human craftsmanship, puppetry and human labor, highlighting strong emotional storytelling (Mouriquand, 2025b). On the other hand, Coca-Cola focused on time and cost saving through a campaign mainly featuring AI-generated visuals (Mouriquand, 2025a). As a result, this sparked significant controversy in public discourse, as most viewers appreciated the human-made advertisement more, associating it with greater warmth, authenticity and emotional depth. This comparison raises broader questions about the trade-off between technological efficiency, and the damaging impact on consumers' perceptions of authenticity, effort and emotion.

A fundamental question therefore emerges. If exceptional technological efficiency reduces perceptions of authenticity, emotional resonance, and creative effort, can brands still treat that efficiency as a purely positive gain? The issue extends beyond aesthetics alone. GAI¹ systems now allow the creation of images and videos so realistic that audiences cannot distinguish them from human-made content anymore. Advertising therefore enters a new environment, one in which the origin of visual material remains invisible to the eye, but highly significant for consumer judgment, raising both strategic and ethical concerns.

1. Research Gap and Objectives

This rapidly evolving environment contributes to the emergence of a complex research domain that requires careful examination to ensure its optimal development. Existing literature has already explored generative AI from multiple perspectives, but several important gaps remain unresolved. Most prior studies concentrate either on AI-generated text or on consumers' ability to recognize synthetic visuals. Far less attention has been devoted to disclosure effects, specifically to the consequences that emerge when brands openly identify advertising content

¹ Generative artificial intelligence

as AI-generated. Therefore, this study addresses three dimensions that remain largely underexplored.

First, it analyzes how disclosure labels shape emotional reactions, perceived authenticity, brand trust, and purchase intention. Despite their central role in the effectiveness of advertising, researchers have rarely examined these factors together within the context of AI-origin disclosure.

Second, the study evaluates whether consumer reactions vary across product categories. It compares chocolate, a hedonic product associated with pleasure and craftsmanship, with toothpaste, a utilitarian product generally purchased out of practical necessity. Existing research offers a limited understanding of how sensitivity to AI-generated advertising varies across psychological characteristics and involvement levels of the advertised product.

Third, the study looks at how changes in disclosure labels over time affect consumers' reactions. In a world with increasing legal pressures to disclose AI use, existing research offers limited understanding of whether companies have an incentive to lie at the risk of being caught, or how consumers react to these label contradictions.

Table 1. Summary of Research Objectives

	Focus Area	Research Goal	Implications
Baseline	No Disclosure	Control Group	Observe initial reactions of consumers on an advertising picture: emotional responses, authenticity, brand trust, purchase intention.
Q1: Transparency	Truthful Disclosure	Does the "AI-Generated" label hurt, or help compared to the baseline?	When we tell the truth about AI, how does it change consumers' perception compared to having no label at all? Management Focus: Is it better for a brand to stay silent or be transparent?
Q2: Moderation	Product Category	To see if consumers have different sensitivities to AI-made visuals for one product compared to another.	Does the "AI Penalty" (from Q1) differ significantly between hedonic (chocolate) and utilitarian (toothpaste) products? Management Focus: Which product advertisements should be done with AI, and which not?

Q3: Deception	Within-Subject	To test if the perceptions of the same individual change because of a different label.	Does changing the label of an ad over time fundamentally alter how consumers feel about the ad? Specifically, does switching from a human to an AI label (<i>Human</i> → <i>AI</i>) trigger feeling of betrayal, or can switching from an AI to a human label (<i>AI</i> → <i>Human</i>) actually increase reactions positively?
--------------------------	----------------	--	--

Overall, many studies explore the topic of generative AI in marketing, but none combine all these dimensions in one single framework, especially in the era of Nano Banana Pro (Gemini 3 Pro Image), a tool that can generate high-quality visual content in a few seconds. The originality of this study lies precisely in this integration: disclosure effects, product-category moderation, and deception dynamics all operate within one experimental design.

Therefore, the research question that this study is seeking to answer to is: *“To what extent does the disclosure of AI-generated advertising images influence consumer responses, and how is this relationship moderated by the hedonic or utilitarian nature of the product?”*

2. Structure of the Thesis

To successfully address the research objectives, this paper is structured into six interconnected chapters. Chapter II establishes the theoretical foundation by reviewing the evolution of AI in advertising, examining consumer information-processing models, and developing the core research hypotheses. Following this, Chapter III details the experimental methodology, including the design of the stimuli, the structure of the questionnaire, and the analytical methods employed. The empirical data and statistical findings for each hypothesis are presented in Chapter IV, which are then contextualized and interpreted considering existing literature in Chapter V. Finally, Chapter VI concludes the study by synthesizing its main contributions, addressing its limitations, and proposing future directions for research.

II. Literature review

To establish a clear theoretical foundation for this study, the present literature review is structured as a progressive framework that moves from broad industry shifts to specific psychological mechanisms. The first part explores the evolutionary context of advertising in marketing, especially with the introduction of artificial intelligence. Then, it examines the foundational information-processing models that explain how consumers interpret explicit disclosure labels as extrinsic cues. Third, it navigates the core psychological conflict in consumer evaluations: the negative biases of the AI penalty against the positive bias of algorithm appreciation. Next, this literature review analyzes how reactions shift based on product classification and the truthfulness of a brand's disclosure over time. Finally, this section tackles the growing legal pressures and individual consumer differences that modern marketers increasingly face, exposing the thin line between consumer trust and corporate manipulation.

1. The Modern Advertising Context: The Integration and Impact of Artificial Intelligence

In the constant evolution of our modern world, few innovations have been as transformative and revolutionary as the emergence of artificial intelligence, defined as “the application of computer systems able to perform tasks or produce output normally requiring human intelligence, especially by applying machine learning techniques to large collections of data” (Oxford University Press, n.d.). Through the introduction of new ways of thinking, producing, and interacting with information, this technological advancement is changing every aspect of human activity, from communication, creativity, and commerce to healthcare and education. In addition to changing how people live, work, and make decisions, artificial intelligence (AI) is now redefining the boundaries separating human and machine capabilities, an effect increasingly affecting marketing activities.

Historically, advertising was defined as the "paid form of non-personal presentation and promotion of ideas, goods, or services by an identified sponsor" (Brooks, 1948). However, this definition has evolved into "a multi-channel business practice where companies pay to place their messaging or branding to promote sales and establish corporate culture" (American Marketing Association, 2013). Therefore, in today's digitalized environment, advertising increasingly relies on diverse online platforms and technologies that continuously reshape how firms interact with consumers.

Within this large context, this study concentrates on a specific category of marketing: AI-generated images in advertising, which represent one of the most persistent and main sources of psychological mistrust in the uses of AI in marketing (L. Zhang & Hur, 2025). In fact, images are essential in communication as they create emotions, attract attention, and influence customers' opinions and perceptions of authenticity. Therefore, generative AI allows marketers to create or enhance images in no time and at a lower cost than human-made ones, enhancing productivity and lowering costs for companies (Chaisatitkul et al., 2024). Thus, it is a strategic and relevant solution to rising costs, with market forecasts estimating a growth in AI-driven spending, from \$16 billion in 2021 to over \$107 billion by 2028 (L. Zhang & Hur, 2025).

Despite being seen as a “transformative game-changer” for marketing companies, this shift creates a deep and unprecedented psychological challenge for the consumer, rooted in the rapid achievement of hyper-realism (Cui et al., 2025). As businesses increasingly incorporate these images into their communication strategies, consumers are becoming concerned about their capacity to differentiate AI-generated content from human-created one, which can in turn affect their trust and purchase intentions (Bai et al., 2023). As good as the human eyes can be, GAI systems are endlessly developing and refining their methods and output, shortening the gap between generated visuals and human ones (Zhou et al., 2019). Frameworks such as the HYPE² benchmark demonstrate that human observers cannot distinguish anymore synthetic from real images, with a 50% indistinguishability threshold where AI images are classified as authentic (Zhou et al., 2019). Similarly, consumers are generally incapable of telling the origin of a piece of art (Bellaiche et al., 2023).

Therefore, if consumers can no longer determine the origin of an advertisement, they remain entirely unaware of the synthetic nature of the visual unless the brand explicitly chooses to disclose it. This ambiguity is particularly important as AI-generated images are often associated with misleading or false claims, encouraging consumers to question and verify the information they encounter online (Velásquez-Salamanca et al., 2025). Disclosure becomes then a key element in marketing research, shifting the focus towards how consumers perceive and react to AI-disclosure labels. Recent research confirms the powerful role labels can play in consumers' minds, much like the "Coca-Cola effect," where the simple presence of a brand

² Human eYe Perceptual Evaluation

name can increase reported enjoyment of a drink (Bellaiche et al., 2023). Consequently, this raises concerns about the modern vision of creativity, realism, and authenticity, all crucial elements for establishing customer confidence and credibility (Beverland & Farrelly, 2010).

2. Consumer Information Processing: How Consumers Perceive Advertisements

To understand consumer reactions, we must look at how individuals perceive information in an advertisement. The **Cue-Utilization theory** suggests that consumers assess product quality based on a range of product-related cues, which can be either intrinsic, such as stable physical features of the product or the image, or extrinsic, such as reviews or labels, meaning information that is not part of the product (Kakaria et al., 2024). Therefore, as consumers are navigating a hyper-realistic advertising world, they cannot rely anymore only on intrinsic cues to evaluate an ad or identify its origin. This means that an extrinsic feature like an AI-label can have much more influence on consumers' perception than any visual, intrinsic cue. Consequently, the disclosure label (extrinsic cue) becomes the main factor in evaluating an ad's message, credibility and influencing purchase decisions.

Once detected, this cue activates two primary psychological processes:

First, the **Persuasion Knowledge Model** is a defensive process explained in three steps:

- The agent-knowledge phase: consumers associate the creator of the ad as an agent with ulterior motives, and therefore start to doubt the brand's motives and the authenticity of the message (L. Zhang & Hur, 2025).
- The persuasion-knowledge phase: this doubt falls into suspicion, skepticism, and resistance that directly affect the ad's credibility and trust (Buder, 2024; L. Zhang & Hur, 2025).
- The algorithm aversion: consumers associate specific cues from the disclosure to a general aversion to algorithm-generated content, deepening the negative perception of the advertisement (Buder, 2024; L. Zhang & Hur, 2025).

Second, the **Stimulus-Organism-Response** explains that an AI label directly decreases source credibility, which then decreases ad likeability, and finally breaks the emotional connection between the viewer (the consumer) and the author of the ad (the GAI tool) which is no longer a human being (Tzara, 2025). Surprisingly, experiments have shown that this penalty exists even if the image quality is identical. For instance, Bellaiche et al. (2023) and Buder. (2024) found that, even if an advertising visual was manually created, and only labeled with "Human" or "AI", the AI label consistently dragged down liking, beauty and profundity scores.

3. The Core Conflict: The AI Penalty vs. Algorithm Appreciation

The literature reveals a division in how consumers evaluate AI-labeled content, characterized by a conflict between two psychological reactions.

3.1. The AI Penalty

The primary aspect of this conflict is a negative bias toward AI-generated content. This is explained by the **Human Effort Heuristic** theory, which quotes effort as a cue of quality of the picture (Kruger et al., 2004). In this framework, human-made is the norm as it was the only way until now, therefore pushing consumers to consider that the higher quality a product is, the highest effort might have been dedicated to it (Kruger et al., 2004; Schrift et al., 2016). Nowadays, with the introduction of GAI tools, consumers evaluate visuals more critically, as they perceive them to require less time and effort to produce (Kruger et al., 2004; Schrift et al., 2016). In fact, effort is here mistaken for competence, resulting in consumer's backlash on GAI advertisements, without any solid argument on the visual quality and value of the output, which may be identical as the human-made one (Kruger et al., 2004).

Similarly, an **Authenticity Gap** arises in consumers' minds when they encounter AI-generated content, which they often perceive as lacking "soul" and emotional depth, as it is produced by a machine (Loth et al., 2026). Therefore, when marketing messages written with an emotional tone are revealed to be AI-generated, positive word-of-mouth and brand loyalty tend to decline. The reason behind this stays in the moral disgust that consumers develop when they perceive inauthenticity (Gao et al., 2026).

Furthermore, the AI label leads to intense examination of it, leading to algorithm aversion and the **Uncanny Valley** effect. **Algorithm Aversion** refers to the tendency of individuals to quickly lose trust in a system after they observe an error (Zeng et al., 2026; Y. Zhang & Gosline, 2022). Thus, when a brand discloses that the image of their advertisement has been generated with AI, consumers scrutinize it for perceived "errors". If visual mismatches are detected, such as facial disproportions, it triggers the uncanny valley effect, which causes a sharp drop in the comfort of the individual who observes the ad (MacDorman et al., 2009; Seymour et al., 2017). Therefore, the observer perceives it as weird and unnatural instead of convincing, which destroys the emotional appeal of the ad (MacDorman et al., 2009).

3.2. Algorithm Appreciation

The secondary aspect of this conflict is a positive bias towards AI-Generated content. One theory standing for it is the **Algorithm Appreciation**, which occurs when consumers prefer or

trust AI more than humans. It generally happens in contexts demanding objectivity, analytical processing, and impartiality (Castelo et al., 2019; Kakaria et al., 2024). The reason behind this effect is the perception that AI is free from human bias and highly efficient (Kakaria et al., 2024). Empirical evidence suggests that in these specific contexts, consumers view the machine's lack of emotion as a strength and not a problem. It leads consumers to choose AI-driven options based on this perceived efficiency and therefore makes artificial intelligence the preferred source for impartial data processing and functional decision-making (Petrova Karina, 2026).

This conflict between punishing AI and endorsing AI creates an unpredictable landscape for advertisers, although current reactions and theories tend to negative. For instance, consumers may feel less manipulated by a transparent label saying “AI-made” but may still express lower levels of ad likeability because they attach importance to human effort.

3.3. Key Dimensions of Consumer Evaluation

Building on the previous literature, this research examines consumer reactions to AI-generated advertising through four central dimensions: emotional responses, trust, authenticity, and purchase intention. These dimensions emerge repeatedly in prior studies as the main mechanisms through which AI disclosure influences advertising effectiveness.

First, concerning **emotional responses**, consumers frequently perceive AI-generated marketing as emotionally flat, manipulative and insincere (Ahac, 2025). “The Uncanny Valley Effect” joins this view, further stating that discomfort arises when consumers sense an emotion is being synthetically simulated rather than relationally intended (MacDorman et al., 2009; Seymour et al., 2017).

Second, **trust** is influenced by conflicting dynamics of Algorithm Aversion and Algorithm Appreciation, depending on whether consumers interpret AI as deceptive or efficient. On one hand, trust sharply declines when consumers discover AI has simulated a human emotional tone, such as empathy or humor, and this especially without transparency (Castelo et al., 2019; MacDorman et al., 2009). On the other hand, consumers may actually trust AI more than humans because they view its lack of emotion as unbiased, objective and more efficient (Kakaria et al., 2024; Zeng et al., 2026)

Third, **authenticity** perceptions are linked to the perceived presence or absence of human effort and emotional sincerity. In fact, it was found that consumers often perceive AI-content as lacking soul, resulting in an Authenticity Gap that creates moral disgust and damages brand

loyalty (Gao et al., 2026; Loth et al., 2026). In addition, even polished AI ads are perceived as inauthentic, unless human co-creation is clearly present to preserve the "voice integrity" of the final message.

Finally, disclosing an AI label can create a "Taint Effect", where consumers punish the brand for perceived laziness, which directly decreases engagement, perceived ad value, and ultimately **purchase intentions** (Chen et al., 2022; Höltervennhoff et al., 2025). However, unlabeled AI-generated content may benefit from the Implied Truth Effect, and should result in higher perceived value and a greater likelihood of purchase (Kruger et al., 2004; Schrift et al., 2016; Sui, 2026).

Together, these dimensions form the theoretical basis of the proposed hypotheses and research model. However, prior literature also suggests that the effects of AI disclosure are not uniform across advertising contexts (Longoni & Cian, 2022; Shi & Jiang, 2026). In particular, consumer evaluations of AI-generated content may vary depending on the type of the product being advertised. Therefore, the distinction between hedonic and utilitarian products emerges as an important moderating factor.

4. The Moderating Role of Product Type: Hedonic vs. Utilitarian

The impact of AI disclosure on consumer behavior is heavily dependent on the advertised product category, as hedonic and utilitarian goods trigger different psychological pathways.

First, **hedonic products**, such as chocolate, fragrances or art, are generally consumed for pleasure or gratification and are enjoyment-oriented (Dhar & Wertenbroch, 2000). The value of this type of product comes from emotional attachment and perceived quality. Therefore, consumers seek the "human touch" when consuming a hedonic product because they associate "human-made" with uniqueness, quality and authenticity (Degirmenci et al., 2025; Fuchs et al., 2015).

Concretely, in the context of AI-generated advertising, hedonic products are vulnerable to the AI disclosure penalty because of the importance of the emotional part. In fact, when consumers evaluate an advertisement for a hedonic product (chocolate), they want something that is qualitative, visually appealing and to find the "human touch" (Longoni & Cian, 2022). Accordingly, disclosing that the image on the ad is AI-generated, disturbs the emotional aspect, and decreases perceived authenticity by replacing the human side with an algorithm (Degirmenci et al., 2025; Fuchs et al., 2015). In addition, it leads to a drop in consumer engagement and purchase intention because the affective connection is damaged (Longoni &

Cian, 2022). Moreover, hedonic products face even more consumer backlash when AI is used to simulate emotions, human warmth or empathy. The reason behind this is that consumers interpret this act as manipulative (Ahac, 2025).

Second, **utilitarian products**, such as toothpaste, are usually utilized out of necessity or for their functional aspect. Consequently, this type of product is evaluated based on performance metrics such as efficiency, durability, and cost-effectiveness (Okada, 2005; Voss et al., 2003). Therefore, consumers consider these products as essential and as potential practical solutions to a specific problem instead of emotional purchases. That being said, consumers may not value an advertisement as a factor when choosing a utilitarian product, no matter how nice its visual quality may be. In fact, they may base their product choice on technical cues such as ingredients and composition, rather than visual aspect (Longoni & Cian, 2022)

Concretely, in the context of AI-generated advertising for utilitarian products, AI disclosure is often seen as a positive or at least a neutral factor. In fact, the innovation boost provided by AI can offset in an effective way the reduction of human touch because consumers evaluate these goods based on practical benefits and not emotional ones (Shi & Jiang, 2026). Furthermore, consumers believe that AI is highly competent for functional tasks and, therefore, consider an AI label for a utilitarian product advertisement more as a sign of innovation and precision than a lack of emotion because the product's primary goal is utility (Longoni & Cian, 2022; Shi & Jiang, 2026). It is also important to mention that the "trust gap" is smaller because consumers prioritize whether the product works over who made the advertisement (Ahac, 2025).

5. Disclosure Integrity: Truthful vs. Misleading

The literature highlights a strategic trade-off between short-term consumer appeal and long-term brand trust regarding the disclosure of AI use in advertising. Several theories suggest that consumers react not only to AI itself, but also to the honesty, consistency, and transparency of the information communicated by the brand over time.

The **Implied Truth Effect** argues that when AI-generated advertising remains unlabeled, consumers tend to assume that the content is human-created, which positively influences evaluations of credibility, authenticity, and purchase intention (Kruger et al., 2004; Schrift et al., 2016). Because human creativity is still associated with effort, sincerity, and artistic value, unlabeled AI visuals can benefit from these positive assumptions (Aburass & Abu Rumman, 2024).

At the same time, **Deceptive Denial Theory** explains how concealing the true AI origin creates an information asymmetry in which the brand possesses information that consumers do not (Bimbo et al., 2020). This practice, described as “deceptive denial,” can generate strong negative reactions once the truth is revealed (Nyilasy & Gangadharbatla, 2025). In line with theories of perceived deception and trust violation, consumers who initially believe an image is human-made but later discover that it was AI-generated may perceive the brand’s behavior as intentionally misleading. This “caught in a lie” dynamic is expected to trigger stronger moral condemnation and a more severe decline in trust and authenticity perceptions than an honest AI disclosure from the beginning (Nyilasy & Gangadharbatla, 2025).

In addition, **Signaling Theory** suggests that consistency in brand communication functions as a signal of credibility and transparency (Spence, 1973). When consumers first receive a truthful “AI-generated” label but are later exposed to a contradictory “human-made” label for the same image, this inconsistency may undermine trust by creating uncertainty about the brand’s intentions. Conflicting source information can therefore be interpreted as manipulative or unreliable, reducing perceptions of transparency and weakening brand trust (Nyilasy & Gangadharbatla, 2025).

However, since prior research consistently shows that consumers attach greater emotional and symbolic value to human-made content, the later introduction of a “human-made” label could partially compensate for the negative effect of inconsistency. In this case, some consumers may still evaluate the brand more positively because they continue to value human creativity more highly than AI-generated production (Gao et al., 2026; Kruger et al., 2004; Loth et al., 2026; Schrift et al., 2016).

6. Individual Differences and Exploratory Findings

Beyond the advertisement itself, the consumer’s own background significantly alters how technological cues are processed.

Gender Differences: Women are found to be significantly more sensitive to relational discomfort and "emotional inauthenticity" in AI messaging (Ahac, 2025). Conversely, men tend to view AI communication through a technical lens, prioritizing functional clarity over the "warmth" of the source (Ahac, 2025).

Age Differences: A generational divide exists in AI perception. It was found that younger consumers (under 35) judge AI functionally and are more willing to accept it as part of the digital landscape, viewing it as efficient. Older consumers, however, often show a visceral

negative reaction, perceiving GAI as "cold" or "weird" and placing a much higher premium on traditional human authorship (Ahac, 2025).

AI Literacy and Attitudes (GAAIS³): Surprisingly, high AI literacy does not always soften reactions, as it can increase "interpretive vigilance" and skepticism regarding data privacy and emotional simulation (Ahac, 2025). Conversely, low-literacy users initially accept ads without further analysis, but experience severe "algorithm aversion" and disappointment when subsequent disclosure (Dietvorst et al., 2015). However, other studies found that consumers with high General Attitudes Toward AI (GAAIS) scores have significantly better attitudes towards AI-made advertisement, as they believe in the potential AI tools can have and reflect it into greater acceptability rates towards it. For these users, if an ad is well-designed, the automation is normalized and seen as "functional honesty" (Ahac, 2025; Bellaiche et al., 2023; Buder, 2024).

Professional and Cultural Context: Studies show that consumers studying or working in media, marketing, or tech show higher tolerance for AI automation as they care more about functionality, rather than authorship. Additionally, while not the primary focus of this study, it is noted that AI reception is not culturally neutral. For instance, it was found that collectivist or risk-avoidant cultures like South Korea or Sweden, tend to be more skeptical of perceived AI manipulation than convenience-driven markets like the USA (Ahac, 2025).

7. Strategic Trade-offs and the Legal Mandate

As generative AI becomes increasingly integrated into advertising, transparency is shifting from a strategic choice to a growing legal and ethical expectation. Regulatory frameworks worldwide are pushing towards mandatory disclosure of AI-generated content in order to reduce deception, misinformation, and copyright-related risks (Cui et al., 2025b; Milaninia & Vermynck, 2025). In Europe, the European Union AI Act already requires certain AI-generated or manipulated content to be disclosed when there is a risk of misleading users or impersonation (Piasecki et al., 2024). Similar initiatives are also emerging in the United States and China, while major platforms and companies such as Meta and YouTube have introduced AI-labeling

³ General Attitudes towards Artificial Intelligence Scale

policies to increase transparency and limit misinformation (Höltervennhoff et al., 2025; Milaninia & Vermynck, 2025; Wortel et al., 2024).

These developments place marketing departments in a difficult strategic position. Brands are therefore faced with a cost-benefit dilemma between maximizing immediate financial gains or sustaining long-term brand equity (Casenave & Klarmann, 2020). This is called a “double blind” trade-off: either label AI-made assets to sustain long-term brand equity, but pay the AI penalty (Swaminathan et al., 2022; Tzara, 2025), or hide it to protect short-term click-through rates and engagement, but risk exposing the company to reputational damage and consumer backlash if the synthetic nature of the content is later revealed (Mizik & Jacobson, 2007; Wortel et al., 2024). In this sense, while AI disclosure may destroy the perceived authenticity of an advertisement, a transparent labeling can simultaneously reinforce the integrity of the brand overall (Liu et al., 2023; Sansome et al., 2025). Therefore, embracing disclosure remains the more sustainable long-term strategy, as honest communication is consistently associated with stronger consumer trust and brand credibility (Cambier & Poncin, 2020; Chauhan, 2025; Montecchi et al., 2024; Yang & Battocchio, 2021). To mitigate this AI penalty, some researchers suggest over-communicating transparency using softer forms of disclosure, such as “co-created with AI” or “made with the help of AI,” which may attenuate consumer backlash (Tzara, 2025). Therefore, these limitations underscore the importance of critical thinking and media literacy in this increasingly AI-saturated digital environment.

However, some studies show potential “backfiring effects”, arguing that labels may not be a solution as they impact consumer reasoning without improving their ability to identify misleading content (Höltervennhoff et al., 2025). Even if labels were mandatory, concerns still remain about the reliability and legitimacy of the labeling process itself, particularly when labels rely on automated detection systems rather than human verification (Höltervennhoff et al., 2025). Moreover, even if everything is labeled and detectors are perfectly accurate, labels will lose their informational value as everything is flagged.

Looking forward, the literature suggests two possible futures for AI in advertising. First, like Photoshop, which was initially criticized, but become the industry’s standard, the “shock” of AI may wear off and GAI imagery will become normalized (Koning & Voorveld, 2025). Second, a permanent divide may emerge in the market, where “Human-Made” becomes a premium, high-value competitive advantage, while “AI-Made” becomes the factory-standard basic option (Koning & Voorveld, 2025; Wortel et al., 2024). Ultimately, the reception of AI

will depend on current usage trends. If GAI continues to be associated with bad-quality scams, algorithm aversion will increase, but if brands use these tools to increase creativity and innovation, the AI aversion will decrease (Wortel et al., 2024).

8. Conclusion and Identification of Research Gaps

While GAI impacts on advertising are increasingly discussed in marketing and consumer research, the specific impact of AI-generated imagery on consumers' trust, authenticity perceptions, and purchase intention remains mitigated. Therefore, while emerging studies cover either AI disclosure effects or the human detection of AI content, this thesis contributes to the current literature by integrating these three dimensions in one study:

- Current research often fails to account for the customer's complex, nuanced opinions and feelings toward AI-generated ads beyond simple binary reactions.
- Little is known about whether consumers' perception of AI-generated advertisements varies depending on the type of product being advertised.
- Few articles analyze customer behavioral changes specifically before and after knowing the true source of the advertised visuals.

9. Hypothesis Formulation

9.1. Hypothesis 1: The Influence of Disclosure and the "AI Penalty"

Drawing from the Persuasion Knowledge Model and Algorithm Aversion Theory, the literature points towards a negative reaction in the presence of an AI-label, compared to an undisclosed advertisement that may benefit from the "Implied Truth Effect". Moreover, the only presence of a label could diminish consumer perception on the advertisement. According to the Stimulus-Organism-Response framework, the label acts as a stimulus that reduces source credibility, thereby breaking the emotional connection with the advertisement. However, to observe this, the visual quality of the image must be good, indistinguishable for the human eyes to detect any algorithmic cues.

H1a – The AI Penalty: Truthful disclosure of AI generation, compared to a baseline of no disclosure, will negatively influence consumer responses, resulting in significantly lower scores for emotional response, perceived authenticity, brand trust, and purchase intention.

H1b – The Label Scrutiny Effect: The simple presence of any label, including a "Human-made" one will result in worse consumer attitudes compared to a no-disclosure baseline, as the act of labeling itself alerts consumers of the potential manipulative intent.

H1c – The Illusion of Effort: According to the Human Effort Heuristic, consumers equate high labor with high quality. Therefore, labeling an AI-generated image as "Human-generated" will result in higher authenticity, brand trust, emotional responses and purchase intention compared to AI-disclosure.

9.2. Hypothesis 2: The Moderating Role of Product Type

The literature suggests a core conflict between the "authenticity penalty" and "algorithm appreciation," which is moderated whether the product is hedonic (pleasure-oriented) or utilitarian (function-oriented).

H2a – The Category Interaction: The reduction in consumer scores caused by an AI label will be significantly greater for the hedonic product category than for the utilitarian category.

H2b– The Hedonic Vulnerability: For hedonic products (e.g., chocolate), truthful AI disclosure will result in a significant decrease in perceived authenticity, emotional responses and purchase intention compared to no disclosure.

H2c – The Utilitarian Buffer: For utilitarian products (e.g., microwaves), there will be no difference in consumer responses between no disclosure and AI disclosure, as consumers prioritize functional efficiency over the "human touch."

9.3. Hypothesis 3: Deception, Revelation, and Behavioral Shifts

This section utilizes a within-subjects approach to test the theories of deceptive denial and information asymmetry. To test the literature on "AI-Washing" and deceptive denial, this study assesses the impact of truthful AI disclosure, followed by misleading human disclosure, compared to misleading human disclosure, followed by the revelation of the true "AI" origin.

H3a – The "Caught in a Lie" Penalty: For consumers initially exposed to a misleading "Human" label, the subsequent revelation of the true "AI" origin will result in a drastic decrease in emotional responses, brand trust and authenticity scores, triggered by perceived betrayal and moral condemnation.

H3b – The Human Premium: For consumers initially exposed to a truthful "AI" label, the subsequent presentation of a conflicting "Human" label for the same image will increase brand trust scores, as the positive premium of human authorship is bigger than the possible source inconsistency penalty.

III. Methodology

To translate the theoretical tensions identified in the literature into empirical data, this section outlines the quantitative research methodology, designed to capture both immediate consumer impressions and subsequent behavioral shifts. Therefore, the first part of this section consists of the general methods and objective of this paper, covering the main components of our approach, followed by the quantitative phase done through a survey designed to identify consumers' behavior and opinions.

1. Study type and background

1.1. Type of study

The method used was a quantitative mixed experimental design combining between-subjects and within-subjects components. This experimental approach was chosen because it allows for the examination of the effect of one (or several) variable(s) on another by manipulating one or more factors while holding others constant.

1.2. Independent & Dependent variables

Independent variables are characteristics of the physical or social environment that are manipulated by researchers to examine their impact (Pleyers, 2025). In this research, two independent variables were manipulated: the product category (two levels: hedonic vs utilitarian) and the disclosure label (Three levels: no disclosure, truthful AI disclosure label ("Made with AI") and misleading disclosure ("Human-made")).

Four dependent variables were used to evaluate the AI penalty:

- Emotional responses (initial participants' reactions to the stimuli),
- Trust (perceived reliability of the advertisement),
- Authenticity (the genuineness of the visual content),
- Purchase intention (the likelihood of the participant to buy the products).

Table 2. Independent Variables vs Dependent Variables

Independent variables (IV)	Dependent variables (DV)
• Product category (2 levels): • Hedonic/Chocolate • Utilitarian/Toothpaste • Disclosure label (3 levels): • No label • AI-made • Human-made	 • Emotional responses • Trust • Authenticity • Purchase intention

Therefore, the research model was defined by following the OLS model (Ordinary Least Squares model). It is the most common method used in linear regression to estimate the relationship between a dependent variable and one or more independent variables.

$$Y_i = \beta_0 + \beta_1(PC_i) + \beta_2(DL_{AI,i}) + \beta_3(DL_{Human,i}) + \beta_4(PC_i \times DL_{AI,i}) + \beta_5(PC_i \times DL_{Human,i}) + \epsilon_i$$

Where:

- Y_i represents the dependent variables. Specifically, it includes:
 - Y_1 = Emotional response
 - Y_2 = Trust
 - Y_3 = Authenticity
 - Y_4 = Purchase Intention
- β_0 (Intercept) represents the predicted value of Y_i when all predictors are at their reference levels.
- β_1 to β_5 represent the coefficients estimating the effect of each predictor X on Y_i .
- PC_i (X_1) represents the independent variable: Product category (Hedonic vs Utilitarian).
- $DL_{AI,i}$ and $DL_{Human,i}$ (X_2) represent the independent variables: Disclosure label (No-disclosure, Truthful AI disclosure $\rightarrow DL_{AI,i} = 1$, Misleading Human disclosure $\rightarrow DL_{Human,i} = 1$).
- $PC_i \times DL_{AI,i}$ and $PC_i \times DL_{Human,i}$ represent the interaction effects between the product type and the label type.
- ϵ represents the error term (accounting for the noise and extraneous variables).

1.3. Research environment

The study was conducted as an online experiment. Participants completed a questionnaire and responded to various questions using their smartphones, iPads or computers. This format was chosen because lab or field experiments necessitate a significant amount of funds. In fact, an online experiment is the most cost-effective option, often requiring minimal or zero financial resources. For this study, no external funding was needed, as data collection relied on the digital distribution of the questionnaire. Moreover, an online experiment allows for a large sample of participants to be gathered quickly and more efficiently, because it is more accessible to individuals (Pleyers, 2025).

However, it is more difficult to control for extraneous variables. In fact, it was complicated to determine if the participants were truly engaged, honest, and reliable, or to account for the environmental conditions in which they responded. Therefore, participants could have been in a crowded room, on the subway, on a train, in the street, or in a park. Consequently, not all participants were in the same environment at the moment of their response.

1.4. Group structure

The primary structure of the study consisted of a 2x3 **between-subjects** factorial design. This resulted in 6 groups, with participants randomly assigned to one of them. Randomization is the process of ensuring that each participant has the same chance of being assigned to any specific group (Suresh, 2011). It allows for the comparison of the different groups, eliminates selection bias, and balances the groups. Furthermore, it forms the basis for statistical tests by ensuring the initial equivalence of the treatment groups.

Table 3. Structure of the groups (2 x 3)

	No Disclosure (control group)	Truthful AI Disclosure	Misleading Disclosure
Hedonic (Chocolate)	Group 1	Group 2	Group 3
Utilitarian (Toothpaste)	Group 4	Group 5	Group 6

Groups 1, 2, and 3 were presented with the chocolate advertisement, while Groups 4, 5, and 6 evaluated the toothpaste advertisement. This division was implemented to test the hypotheses regarding hedonic versus utilitarian product categories.

Groups 1 and 4 represented the control groups. Participants in these groups saw the target ad (= the advertisement with the chocolate or toothpaste) with no disclosure label. Therefore, this part was essential because it would show how much people naturally liked the advertisement when they didn't know how it was made (baseline) and also was essential to test the hypotheses.

Groups 2 and 5 were designed to test truthful AI disclosure by informing participants, via a label, that the advertisement was indeed AI-generated. Conversely, Groups 3 and 6 tested misleading disclosure by stating that the advertisement was human-made when it was, in fact, AI-generated.

However, it is important to mention that a secondary **within-subjects** component in the labelled conditions was added. It meant that the participants that were assigned to the AI (groups 2 and 5) or Human (groups 3 and 6) disclosure groups evaluated an additional version of the same AI-generated advertisement presented with the opposite disclosure label. This within-person comparison would test whether evaluations change solely due to the label for the same person.

Table 4. Structure Target Ads

Group	First Target Ad Seen	Second Target Ad Seen
G2	AI label	Same ad with Human label
G3	Human label	Same ad with AI label
G5	AI label	Same ad with Human label
G6	Human label	Same ad with AI label

1.5. Sample size and characteristics

To determine a sufficient sample size, an *a priori* power analysis was conducted using the G*Power, a statistical software used to determine the required sample size for hypothesis testing. By setting an expected effect size of 0.15, a Type I error rate (α) of 5%, and a statistical power of 80%, with around 19 predictors, the analysis indicated that a minimum sample size of 153 participants was required to achieve sufficient power. However, to provide a small buffer and account for the expected 10% of invalid responses, as well as possible outliers that may arise, a larger sample was targeted. Moreover, as the sample was randomized into six groups, at least 30 relevant participants per group were needed. Therefore, it was decided to aim for 200 responses.

Furthermore, as this was an exploratory and non-representative study of a population, a convenience sample was used. In fact, convenience sampling involves selecting participants that are readily accessible (Golzar, 2022). In this case, the questionnaire was initially shared within personal networks and distributed further through a snowball effect as participants shared it with others.

1.6. Filler ads

Five filler ads were used so that respondents in Groups 2, 3, 5, and 6 would not only see the two target ads with different labels. In fact, participants could have guessed the study's purpose if they had seen the same advertisement twice, once with the “Made with AI” label and after,

directly with the “Human-made” label. Therefore, one filler ad figured before the advertisement of interest, three in-between and one after the second advertisement of interest.

For the five filler ads, five distinct products were strategically selected: juice, cereal, shampoo, soap, and canned food. Specifically, two of these products are considered hedonic (juice and cereal), two utilitarian (shampoo and soap), and one neutral (canned food). Therefore, using two products for each category was intended to allow for more nuanced observations during the data analysis if considered relevant.

Finally, these filler ads were also included in the control groups' questionnaires to ensure that the survey length did not differ significantly, thereby standardizing respondent fatigue across all conditions.

1.7. Tools used and justified choice

While platforms such as Qualtrics and SurveyMonkey were initially considered, their free version didn't allow image downloads or unlimited commands needed to create the full survey. Therefore, despite its simplicity, Google Forms was finally chosen to deliver the questionnaire. It was the most accessible option that provided all the features required to complete the research for free, such as downloading images and maintaining the entire survey in one single form. Moreover, the platform's familiar interface ensured good usability, minimizing potential navigation errors for participants. It was also chosen because it allowed the study to comply with data protection standards by ensuring complete anonymity for all respondents. The only real downside was the lack of a built-in randomization tool, but it was partially solved using a specific filter question and branching logic, which will be explained in the section 3.2 of this chapter.

2. Procedure of generation of images

The advertisements were made 100% with artificial intelligence, using Gemini's tool for creating photos “Nano Banana Pro (Gemini 3 Pro Image)”. Initially, the strategy was to simply ask the tool to create an advertisement for one of the specific products. However, the results were not immediately successful because the advertisements looked too AI-generated, and not realistic enough. For example, a systematic problem encountered was that the text within the images was not well written (some letters were missing, some words were oblique, etc.).

Therefore, after the first version of the picture was generated, additional instructions were given to the AI chatbot, such as “change the background”, “improve the quality of the label on the product”, “remove the green plants”, etc. However, the more changes that were requested, the

less effective and lower quality the advertisement became. Consequently, it was decided to create a more consistent, detailed prompt, so that fewer rectifications would be needed. In fact, creating the right prompt proved to be the most important factor in order to achieve a good and natural advertisement.

To continue, one picture per product was needed at the end. However, various pictures were generated for each of the six products. Therefore, at least three picture options per product were produced, such as: one with a human hand present; one with a human face present; one with just the product and a background; one with a different background; and one with a different brand name.

Afterwards, the visuals were reviewed by an industry specialist, Ioana Mateiciuc, currently working as Marketing Manager at Chilas B.V. (Chilas Lasers), a photonic-based high tech product company based in Enschede, the Netherlands. She directed detailed comments, as well as evaluated each product's images using a three-point scale, where 1 was associated with the most AI-looking, and 3 represented the more human-made aesthetic. This exploratory pre-selection process was conducted in order to choose the most qualitative and least "AI-looking" image for each product, focusing especially on the two target products: the chocolate and the toothpaste.

While these initial ratings made by the specialist were heavily considered for the final selection, a decision was also made to only use advertisements featuring the product itself, without any human faces. In fact, this choice was made to avoid biases directed towards a person rather than the product. The objective was to ensure that no one rated an advertisement higher on any aspect simply because, for example, an attractive person was present in it. The final step was to use Google Photos to erase the Gemini logo from the pictures and use Canva to add the labels "Made with AI" and "Human-made". Finally, this process of generating the final images took approximately 10 hours.

Therefore, this was the created final prompt that generated the final advertisements for the products (see images in appendix 2):

Prompt: Can you create an advertising picture for a chocolate brand? It's for a master thesis on the subject of labeling changes in consumers attitudes, so the objective is for the image to look 0 percent made by AI, like non detectable by the human eye.

Therefore, high quality and everything. But the thing it's that it's an advertisement, so it should mention the brand, the brand logo, a slogan, a real-looking packaging... like existing ads from well-known brands. I don't want just the picture for the ad but the final version of a real ad, with a fictional logo, brand and messages...

After the picture was generated: I need you to improve the quality of the label to look more human made. I don't like the text quality, it's too AI.

3. Questionnaire structure

3.1. Onboarding and Ethics

At the beginning of the questionnaire, a cover story was used to minimize respondent bias and prevent demand characteristics, where participants might alter their responses if they were aware of the true purpose of the research (McCambridge and al., 2012). Therefore, respondents were informed that the study's objective was simply to explore how individuals perceive and interact with different types of advertisements. Moreover, the filler ads helped to reinforce the credibility of the cover story because it makes it more legitimate by seeing multiple advertisements and not just the ones about chocolate and toothpaste.

To continue, in accordance with the General Data Protection Regulation (GDPR), the first section of the questionnaire explicitly addressed the protection of personal data. Participants were provided with a clear statement regarding data anonymity and the strict usage for academic purposes, ensuring full ethical compliance.

The language of the questionnaire chosen was English. The decision was made to ensure consistency in the interpretation of the different scales that were coming all from academic papers in English and were already validated. Furthermore, using a single language facilitated the data analysis phase because it eliminated some potential nuances lost in translation or some discrepancies, allowing to increase the internal validity of the findings of the research.

The introductory section also informed participants that the survey would require approximately 10 to 15 minutes to complete. This allowed potential respondents to assess their availability before starting, thereby ensuring that just those who had sufficient time took the survey and that they did not respond in a rush.

Finally, the first two lines of the introductory section clarified that although Google Forms might technically display 46 pages due to the complex randomization logic and multiple

experimental branches, each individual would only see approximately 10 pages. This transparency was essential to prevent respondent dropout caused by the number of pages.

3.2. The Randomization Mechanism

The questionnaire was structured into six distinct experimental paths, ensuring that each participant was exposed to only one of the six specific groups. Therefore, in order to implement randomization, a question was used just after the introductory section, asking to select a number between one and six. Each number was systematically linked to a corresponding experimental condition. For example, if a participant selected “1”, they were redirected to Group 1; if they chose “2”, they were redirected to Group 2, and so on for all six options.

3.3. Groups

As the control groups, Groups 1 and 4 were exposed to six advertisements without any disclosure labels. This provided a baseline measurement of consumer reactions in the absence of AI or human-made claims. The only variable separating the two was the product type, with Group 1 seeing the chocolate ad and Group 4 seeing the toothpaste ad.

The other groups were shown seven advertisements. The order of the ads and whether or not they had a label were chosen strategically. The sequence began with the filler ad “Juice” with no disclosure label for the 4 groups. This was in order to establish a cognitive baseline and to serve as a control point to set a “normal” browsing expectation, such as they were really just rating an advertisement online. Therefore, by starting with a neutral and unlabeled product, it ensured that the first interaction was not influenced by AI skepticism or labeling awareness.

For the second ad, the participants encountered the first target ad (chocolate for Group 2 and 3 and toothpaste for Group 5 and 6). For Groups 2 and 5, they evaluated the ad with the label “Made with AI” and for Groups 3 and 6, with the label “Human-made”. It was the first time that they were faced with a label, therefore, it represents their first impression/reaction to the source claim without any prior comparison.

For the third ad, the participants viewed the filler ad “Cereals”. This ad had a different label than the previous ad. In fact, if Ad 2 was labeled “Made with AI”, Ad 3 is labeled “Human-made” (and vice versa). This switch of label is made in order to prevent the respondents from assuming the entire survey is only about AI. Moreover, the purpose of the third ad is to normalize the presence of labels on the advertisements and mask the research intent.

For the fourth ad, the participants encountered the filler ad “Shampoo” without any disclosure label. It is in order to prevent automaticity. In fact, this forces the respondents to re-engage cognitively and ensures that they do not click random answers based on a perceived pattern rather than the actual content of the ad.

For the fifth ad, the participants viewed the filler ad “Soap” with the same label as the second advertisement. This choice was in order to normalize the upcoming label change.

For the sixth ad, the participants encountered the second target ad. It means the same ad as the second one, but with the opposite disclosure label. The purpose is to measure the label shift for the same person. In fact, this choice isolates the label as the only changing variable and therefore allows to test the hypothesis if the respondents will rate higher facing a “Human-made” label or lower facing a “Made with AI” label in comparison to the second ad.

For the last ad, the participants viewed the filler ad “canned food” without any label. This serves as a neutral buffer to make them continue believing in the cover story. The purpose of this last ad is also in order that the target ad is not the last ad they have seen before answering the General Attitudes towards Artificial Intelligence (GAAIS).

Table 5. Order of the different advertisements

Order	Product	Group 1 & 4	Group 2 & 5	Group 3 & 6
1	Juice	No Label	No Label	No Label
2	Chocolate/Toothpaste	No Label	"Made with AI"	"Human-made"
3	Cereals	No Label	"Human-made"	"Made with AI"
4	Shampoo	No Label	No Label	No Label
5	Soap	No Label	"Made with AI"	"Human-made"
6	Chocolate/Toothpaste	No presence of this ad	"Human-made"	"Made with AI"
7	Canned Food	No Label	No Label	No Label

3.4. Questions

3.4.1. Emotional Responses

Emotional responses were measured using two questions. The first one was a three-item, 7-point Likert scale going from strongly disagree to strongly agree (Buder et al., 2024). Participants rated the target advertisement(s) on three dimensions: visually appealing, creative and likable. However, for the filler ads, the respondents only had to rate them on the first two

dimensions in order to decrease the number of questions, so that the questionnaire would not be too long.

The second question was a two-item, 7-point semantic differential scale where 1 was the negative trait and 7 was the positive trait. The first item was an evaluative semantic differential scale (unpleasant/pleasant), used to measure the overall attitude of the respondents towards the advertisements (Kuhn, 2025). The second item was arousal semantic differential scale (boring/captivating), chosen to evaluate the level of excitement or arousal associated with the advertisements (Dal Palù, 2017). This question was asked for all advertisements. The semantic differential scale was used instead of a Likert Scale in order to force the participants to choose between two opposite qualities.

3.4.2. *Trust*

Ad Credibility was observed using a four-item, 7-point semantic differential scale. The first two items (unconvincing/convincing and untrustworthy/trustworthy) were adapted from the validated scale by MacKenzie and Lutz (1989). The two other items were typical semantic differential scales used in marketing research: low quality/high quality, unrealistic/realistic setup (Edwards, 2025). These questions were reserved only for the target advertisement(s), as a deliberate design choice to maintain the cover story. In fact, not questioning the trust of every filler ad prevented participants from scrutinizing the labels too closely and from asking themselves too many questions about the true nature of the research and objectives.

3.4.3. *Authenticity*

To measure authenticity, a single question was asked using a 7-point Likert scale with 3 items. The first item was “I can tell that the brand put a lot of effort into creating this ad” and was adapted from Napoli et al. (2014). The second item was “This ad is consistent with what I expect from a professional brand” that was adapted from Napoli et al. (2014). The third item was “This advertisement feels like a genuine and authentic representation of the brand” that was adapted from Morhart et al. (2015). This question was asked only for the target advertisement(s) because authenticity is not a factor that is generally asked in order to just rate some general advertisements. Moreover, for the Groups 2, 3, 5 and 6, another item was added for this question in order to measure the opinion of the respondents about the visual-label incongruity. This was represented by the following item: “There is a conflict between what I see and what I am told about this ad” (K. John et al., 2019). It was therefore not asked for Groups 1 and 4 because their advertisements had no disclosure label.

3.4.4. Purchase Intention

Purchase intention was measured using a 7-point Likert scale ranging from strongly disagree to strongly agree adapted from Gazley et al. (2015). The participants rated the advertisements on three dimensions: “If I were to buy this product, my likelihood of purchasing the product in the advertisement would be high”, “If I were to buy this type of product, I would consider the product in the advertisement” and “If I were to buy this type of product, I would be willing to buy the product in the advertisement”. However, for the filler ads, only the two last items were asked in order to diminish the length of the questionnaire.

Table 6. Number of questions per experimental condition

Number of questions per	Groups 1&4	Groups 2, 3, 5, 6
Filler ad (5x)	6	6
Target ad (1x or 2x)	15	16
Total	45	62

3.5. Label Check

To verify that participants had actively processed and recalled the disclosure labels of the two target advertisements, a memory check was performed at four intervals (following the “Juice” and “Soap” fillers and both target advertisements). This check was presented on the next page after each of these four advertisements to prevent respondents from revisiting the image and checking the label. Participants were asked: “To ensure you are paying attention to the visual details of the previous advertisement (product in question), which of the following elements were present in the advertisement you just saw?” Then, you had the following options: product price, brand logo, a “Made with AI” or “Human-made” label (depending on the disclosure label of the advertisement), a website URL and none of the above. By embedding the label within a list of other visual elements, the study masked that the label is one of the most important elements of the study and that the respondents guessed the objectives of the research. Consequently, if the participants did not tick the box with the label, it would mean that they did not see it, and thus their answers would be unreliable. However, this check does not ensure that participants really believed the label’s truthfulness.

3.6. Utilitarian vs Hedonic perception of products Check

To test the second hypothesis, it was necessary to verify that participants’ perceptions of the product categories aligned with theoretical definitions of hedonic and utilitarian goods.

Specifically, the questionnaire needed to confirm whether chocolate was perceived as a hedonic product and toothpaste as a utilitarian one. Therefore, following the experimental section, respondents were asked to rate some of the products that were present in the advertisements (the product in general and not these specific advertisements) using a Checkbox grid. On the vertical axis, four products were listed (two hedonic and two utilitarian). For Groups 1, 2 and 3, these included “Chocolate”, “Shampoo”, “Cereals” and “Soap”; and for Groups 4, 5 and 6, the list mentioned “Juice”, “Toothpaste”, “Cereals” and “Shampoo”. The horizontal axis presented six adjectives: three utilitarian (effective, functional and necessary) and three hedonic (exciting, delightful and enjoyable). The participants were allowed to tick multiple boxes for each product based on their personal perceptions.

3.7. Demographics

To facilitate a detailed analysis of participant diversity, the demographics section was divided into four variables: gender, age, education level and profession.

The gender question utilized a multiple-choice format with the options: Female, Male, Other, Prefer not to say. In fact, the “Other” category was intentionally selected over “Non-binary” to remain inclusive of the increasingly diverse range of gender identities (such as agender, pangender, genderqueer, etc.). The age question was divided into 6 categories (18 - 24, 25 - 34, 35 - 44, 45 - 54, 55 - 64 and 65+). This allows for easier and more efficient cross-generational comparisons.

Regarding the education level, the question was the following: “What is your highest education level achieved?” The options included primary education, secondary education, post-secondary/vocational education (such as internship/professional training), professional bachelor, university bachelor, master or higher and other. This broad number of propositions was designed to ensure the survey remained accessible to all backgrounds rather than focusing exclusively on academic paths. Finally, the profession question was presented in an open format and was asked to identify the potential differences and similarities across different sectors.

3.8. GAAIS

To assess if each participant had positive or negative sentiments of AI in general and to see if their responses are correlated or not with their opinion about AI, six items were selected from the General Attitudes Towards Artificial Intelligence Scale (GAAIS) (Schepman & Rodway, 2020). Specifically, only three positive and three negative items were utilized in order to

balance the scale while minimizing survey fatigue and respondent cognitive load. A 5-point Likert Scale was used to provide a neutral midpoint, allowing respondents to express a balanced view. Moreover, the items were presented in an alternating valence sequence (see the table under) rather than being grouped by sentiment. This strategy was implemented to mitigate agreement bias (= agreeing regardless of the content) and therefore to check if the respondents continued to pay attention to what they were reading. This question was asked after the demographics section, in order that the participants discovered only at that moment that the study had something to do with AI.

Table 7. Order of the GAAIS items

Order of items	Positive/Negative
Artificial Intelligence can provide new economic opportunities for this country.	POSITIVE
I am impressed by what Artificial Intelligence can do.	POSITIVE
I find Artificial Intelligence sinister.	NEGATIVE
There are many beneficial applications of Artificial Intelligence.	POSITIVE
Organizations use Artificial Intelligence unethically.	NEGATIVE
I think Artificial Intelligence is dangerous.	NEGATIVE

3.9. Frequency of purchase

The final question of the survey addressed purchase frequency for the respective target products. Specifically, participants in Groups 1, 2 and 3 were asked: “How often do you buy chocolate?”, with options including “Daily”, “2-3 times a week”, “Once a week”, “2-3 times a month”, “Once a month”, “Less than once a month”, “Never” and “Other”. Similarly, participants in Groups 4, 5 and 6, were asked: “How often do you buy toothpaste per year?”, with a numerical scale from one to eight and an open field for alternative figures.

These questions served to identify potential confounding variables. By capturing the respondents' level of product category involvement, the study could determine if a negative ad evaluation was a result of the experimental label or simply due to a lack of personal usage of the product category. For instance, if a participant does not eat chocolate at all, then maybe it could be the reason why he/she has rated low the advertisement about chocolate.

IV. Analysis of Results

1. Final data gathered

The questionnaire was conducted over a period of three weeks, from the 30th of March 2026 until the 17th of April 2026, resulting in a total of 187 responses. While this number did not fully meet the initial target of 200 responses, only 21 responses were identified as invalid based on the attention check (label check). Therefore, the final usable sample consisted of 166 respondents, which still exceeded the minimum required sample size of 153 participants determined through the a priori power analysis.

2. Preparation of the data gathered

The data was first exported into an Excel file to prepare it before transferring it to Stata. This step was completed using Google Sheets, which was easier to connect with the questionnaire created through Google Forms.

Therefore, an “id” column was first created to assign a unique number to each respondent, ranging from 1 to 187. Additional variables were then generated to identify the product type and label type associated with each group. Respondents exposed to the chocolate product (groups 1, 2, and 3) were assigned the value 0 for the variable “product_type”, while those exposed to the toothpaste product (groups 4, 5, and 6) were assigned the value 1. Similarly, the variable “label_type” was created to distinguish between the first exposure to ads without a label, with a truthful “AI-label”, and with a misleading “human-label” claim.

To make the dataset analyzable in Stata, all Likert-scale responses were transformed into numerical values. The 7-point Likert scales were converted from 1 (“Strongly disagree”) to 7 (“Strongly agree”), while the 5-point GAAIS scales were converted from 1 to 5 using the same logic. In addition, all columns were renamed with short and clear variable names, without spaces or special characters, to simplify the analysis process in Stata.

Next, as the randomization was done manually, each group initially had its own section with repeated variables. To facilitate the analysis, responses from the six groups were regrouped into unified variables for the target ads (chocolate and toothpaste). New variables were therefore created to combine the responses related to the same ad, label, and question across groups.

Control variables were also generated. A variable named “passed_check” was created to indicate whether respondents correctly identified the labels shown in the target ads. Likewise, a “product_type_check” variable was created to verify whether respondents associated

chocolate with hedonistic adjectives and toothpaste with utilitarian adjectives. Respondents who matched the expected product type received a value of 1, while the others received a value of 0.

Finally, demographic variables such as gender, age, profession, education level, and purchase frequency were transformed into numerical categories to facilitate statistical analysis in Stata. For instance, professions were sorted to 5 different categories: students (undergraduate students, master students, PhD students), academia and research (university professors, lecturers, researchers), STEM⁴ and healthcare (engineers, IT, doctors, pharmacists, architects, physicists), business, social sciences and humanities (marketing, economics, law, HR, communication, finance, teachers, administratives, government workers), and other. Once these steps were completed, the dataset was then transferred to Stata to be analyzed.

3. Stata analysis

3.1. Data cleaning

To analyze the data gathered, it's essential to perform a data cleaning beforehand, meaning the process of preparing it by performing different manipulations such as fixing errors, renaming variables, transforming variables and dropping variables (Map and Data Library, 2017). First, using the passed_check column, 21 observations were deleted (reducing the count from 187 to 166). Then, the meanings of the different numerical values in the demographic, product_type and label_type columns were defined in Stata. This ensured that Stata interpreted and gave the data correctly. For instance, in the product_type column, 0 was defined as chocolate and 1 as toothpaste. Afterwards, missing data was managed by converting zeros (from missing values for questions not asked in some groups) into points, so that the means could be calculated correctly. Finally, the values of the conflict_visual_label column were reversed-coded because, for this question, "Strongly Agree" was seen negatively rather than positively (1=7, 2=6, 3=5, 4=4, 5=3, 6=2, 7=1) and also the values of the negative items of the GAAIS columns were reversed-coded (1=5, 2=4, 3=3, 4=2, 5=1).

3.2. Identifying Outliers

In order to identify outliers, box plots for each variable (target ad 1 and target ad 2) across all six experimental conditions were generated. In fact, it was necessary to check if some

⁴ Science Technology Engineering and Mathematics

participants gave almost the same answer throughout the questionnaire and therefore did not pay attention to the questions at all. To continue, outliers were assessed by inspecting the box plots to see if there were dots.

Then, the identification (ID) of the participants who corresponded to the dots was found, and their questionnaires were checked in order to verify which ones had a strong opinion and which ones were just answering without reading. Results identified 24 outliers (dots) across all variables and conditions, but only two of them (57 and 73) were considered problematic; therefore, their responses were deleted. To confirm the impact of these two outliers, a Hausman test was conducted, with H_0 (the null hypothesis) stating that the two models (with and without the removal of the two outliers) produce the same estimates. The results revealed a significant difference, with a $\chi^2(19)=51.05$ and $p>\chi^2=0.0001$. Therefore, H_0 was rejected, confirming that removing the two outliers was necessary to avoid bias of the model and ensure reliable coefficients. Consequently, the sample size became 164 instead of 166 ($N=164$). (see appendix 1.3)

Table 8. IDs of the outliers for each variable

Dependent Variable	Target advertisement	IDs + Group number
Emotional Responses	Target ad 1	73 and 184 (Group 6)
	Target ad 2	57 (Group 5)
Trust	Target ad 1	178 and 169 (Group 1), 61 (Group 4), 57 and 177 (Group 5), 149 and 73 (Group 6)
	Target ad 2	57 (Group 5)
Authenticity	Target ad 1	27 and 63 (Group 4), 57, 22, 103 and 177 (Group 5), 111 and 73 (Group 6)
	Target ad 2	57 (Group 5)
Purchase Intention	Target ad 1	169 and 178 (Group 1)
	Target ad 2	/

3.3. Sample Description

The final sample consisted of $N=164$ participants. In terms of gender distribution, 54.27% of participants identified as female ($n=89$), 44.51% as male ($n=73$), and 1.22% preferred not to say ($n=2$). In terms of profession, the majority of the sample were persons working in domains

of Business, Social Sciences and Humanities (35.37%) and students (31.71%). In terms of age, the two most common categories were 25-34 (35.98%), and 18-24 (33.54%). Finally, the majority of the sample had a master's degree or higher (56.71%).

Table 9. Sample Demographic Characteristics

Demographics	Frequency (n)	Percentage (%)	Cumulative (%)
Gender			
Male	73	44.51 %	44.51 %
Female	89	54.27 %	98.78 %
Prefer Not To Say	2	1.22 %	100 %
Profession			
Students	52	31.71 %	31.71 %
Academia and Researchers	24	14.63 %	46.34 %
STEM and Healthcare	19	11.59 %	57.93 %
Business, Social Sciences and Humanities	58	35.37 %	93.29 %
Other	11	6.71 %	100 %
Age			
Between 18-24 years old	55	33.54 %	33.54 %
Between 25-34 years old	59	35.98 %	69.51 %
Between 35-44 years old	13	7.93 %	77.44 %
Between 45-54 years old	12	7.32 %	84.76 %
Between 55-64 years old	25	15.24 %	100 %
Education level			
Secondary Education	8	4.88 %	4.88 %
Post-Secondary Education	2	1.22 %	6.10 %
Short-type / professional bachelor	14	8.54 %	14.63 %
University bachelor	47	28.66 %	43.29 %
Master or higher	93	56.71 %	100 %

3.4. Groups Description

The experimental groups are well-equilibrated across both product and label conditions. In fact, the sample is almost perfectly split by product type, with 83 participants for the chocolate product and 81 participants for the toothpaste product (N=164). To continue, regarding the label type, the “No Label” and the “Made with AI” label groups are almost identical in size (n=58 for the first condition and n=57 for the second), and the “Human-made” label group is a little bit smaller but still comparable (n=49).

Table 10. Number of participants per experimental condition

	No Label	Made with AI	Human-made	Total
Chocolate	27	31	25	83
Toothpaste	31	26	24	81
Total	58	57	49	164

3.5. Scale Reliability

In order to test the reliability of the scales and be able to create an average of the scores for the different dependent variables, Cronbach's Alpha (α) was calculated. In fact, Alpha allows to measure the internal consistency of a test or scale. It is expressed as a number between 0 and 1. Internal consistency serves to see if all the items in a test measure the same concept or construct and therefore if the items within the test are correlated or not (Tavakol & Dennick, 2011). In order to ensure the scale is reliable, the Cronbach's Alpha must be superior to 0.7 (even better to be above 0.8) and under 0.95. In fact, under 0.7, it means that the items might be confusing and therefore averaging them might be a problem. To continue, above 0.95, it means that the items are too similar, too redundant.

This was done for the four variables of interest (emotional responses, trust, authenticity, purchase intention), for both target advertisements. However, the item measuring visual-label incongruity was excluded because it did not fit with the "authenticity" construct. Specifically, its inclusion caused the Cronbach's alpha to drop to 0.7523 for t1 and 0.7162 for t2; removing the item increased these values to 0.8450 and 0.8933, respectively.

The results found that every scale was reliable (the smaller alpha obtained was 0.8153) because it was not possible to delete one item without decreasing the value of the alpha. Moreover, the results of the Item-rest correlation were always superior to 0.30 meaning that people who have scored high on one of the items for a variable, have also scored high on the rest of the items.

However, it is important to mention that the alpha's found for purchase intention (for t1 and t2) were both very close to 0.95 meaning that the items could be a little bit redundant but as it did not go above it, it was decided to keep the 3 items to calculate the purchase intention.

Table 11. Results scale reliability test (Cronbach's Alpha)

Variable	Advertisement	Number of items	Cronbach's alpha
Emotional Responses	Target ad 1	5	0.8654
	Target ad 2	5	0.9119
Trust	Target ad 1	4	0.8153
	Target ad 2	4	0.8979
Authenticity	Target ad 1	3	0.8450
	Target ad 2	3	0.8933
Purchase Intention	Target ad 1	3	0.9264
	Target ad 2	3	0.9450

Following this, the composite mean scores for each variable were composed using the command "egen", which allows the generation of new variables. These are the columns that have been used to test the different hypotheses.

3.6. Randomization Check

In order to verify if the randomization process across the six groups had worked, a series of Chi-square tests of independence were realized for the four demographics variables (gender, age, education level and profession).

Results indicated no significant differences between the experimental conditions in terms of gender, age, education level and profession, suggesting that the groups were well-balanced. In fact, all the p-values⁵ found were higher than 0.05. For example, for gender it was obtained a p-value of 0.506 (50.6%) meaning that it is a high probability that this repartition happened naturally. Indeed, in statistics, it starts becoming worrying if that probability drops below 5% (0.05).

Table 12. Results of the randomization check (Chi-square)

⁵ The threshold of 5% will be taken into account for all the results as it is the norm.

Demographic Variable	Pearson chi2	P-value
Gender	9.2745	0.506
Age	27.5849	0.120
Education level	19.1563	0.512
Profession	23.4693	0.266

3.7. Normality of the dependent variables

Next, in order to see if the dependent variables were following a normal distribution, the Jarque-Bera test was applied for each of them and for both target advertisements. In fact, it is a statistical tool that checks that the data is normally distributed, this being the null hypothesis (H_0) (Islam et al., 2022). To interpret the results of this test, it is required to look at the p-value. In fact, a large p-value indicates the data is normally distributed and a low p-value ($p < 0.05$) indicates that it isn't and that the null hypothesis is rejected. In addition to the test, a Q-Q plot (quantile-quantile plot) for each variable was generated in order to see visually if it was normally distributed or not (Islam et al., 2022). It was done by observing if the points followed more or less the reference line (see appendix 1.2).

Visual inspection of the Q-Q plots and the Jarque-Bera tests confirmed that all the dependent variables were normally distributed ($p > 0.05$).

Table 13. Results normality of dependent variables test (Jarque-Bera)

Dependent Variable	Target advertisement	P-value
Emotional Responses	Target ad 1	0.5723
	Target ad 2	0.8052
Trust	Target ad 1	0.0588
	Target ad 2	0.6368
Authenticity	Target ad 1	0.4555
	Target ad 2	0.6750
Purchase Intention	Target ad 1	0.1199
	Target ad 2	0.5601

For Hypothesis 3, which requires a paired t-test, the normality condition applies specifically to the difference between the variables (Target ad 1 - Target ad 2). Concretely, for H3a we had to

check that these differences follow a normal distribution for groups 3 and 6, and for H3b, for groups 2 and 5. To do this, the Shapiro-Wilk test was applied. This specific test was chosen because it has higher statistical power than the Jarque-Bera test, especially for sample sizes of around 50 (Islam et al., 2022). Similar to the previous check, the null hypothesis (H0) assumes that the data is normally distributed and a p-value above 0.05 means the data follows a normal distribution.

However, unlike the previous dependent variables, the results here were mixed. The Shapiro-Wilk tests showed that only a few differences were actually normally distributed: Emotional Responses for groups 2 & 5 ($p = 0.076$), Trust for groups 3 & 6 ($p = 0.055$), and Purchase Intention for groups 2 & 5 ($p = 0.257$). Because of these violations, we relied on the Central Limit Theory which justifies the use of t-test for our sample size. In addition, to ensure extra precaution, the non-parametric Wilcoxon Signed-Rank test will be used alongside the t-tests to verify the reliability of the results.

Table 14. Results normality test for H3

Dependent Variable	Groups	P-value
Emotional Responses Target ad 1 - Target ad 2	Groups 3 & 6	0.00281
	Groups 2 & 5	0.07607
Trust Target ad 1 - Target ad 2	Groups 3 & 6	0.05566
	Groups 2 & 5	0.00031
Authenticity Target ad 1 - Target ad 2	Groups 3 & 6	0.00083
	Groups 2 & 5	0.00031
Purchase Intention Target ad 1 - Target ad 2	Groups 3 & 6	0.00043
	Groups 2 & 5	0.25749

3.8. Homogeneity of Variances (Homoscedasticity)

The assumption of homogeneity of variances was assessed using Levene's test. In fact, it is one of the most common statistical tests used to determine whether two or more groups have equal variances (Islam et al., 2022). Stata provides three options for Levene's test results:

- W0, which is the test statistic of Levene's test centered at the mean
- W50, which is the test statistic centered at the median
- W10, which is the test statistic centered at a 10% trimmed mean

The p-value of W0 has been used because all the dependent variables were normally distributed. For the assumption to be met, a p-value above the threshold of 5% (0.05) is needed because it tests the null hypothesis that the population variances are equal (homoscedasticity).

3.8.1. For Hypothesis 1

Levene's test was performed to check if the variances of each dependent variable (emotional responses, trust, authenticity and purchase intention) for target ad 1, at each level of the independent variable (label type) are the same or not. The results showed, for all four dependent variables for target ad 1, a p-value higher than 0.05, meaning that the assumption of homoscedasticity was met and therefore that the variances were not different (equals).

Table 15. Results homoscedasticity test for H1 (Levene)

Dependent Variable	W0 (Levene)	df (degree of freedom)	P-value
Emotional Responses	0.23381466	(2, 161)	0.7917768
Trust	0.87077353	(2, 161)	0.42058965
Authenticity	0.02142231	(2, 161)	0.9788083
Purchase Intention	1.4307433	(2, 161)	0.24215493

3.8.2. For Hypothesis 2

The assumption of homoscedasticity for the 2x2 factorial design (interaction between product type (chocolate and toothpaste) and label type (No Disclosure and AI Disclosure)) was assessed. Levene's test across the four experimental cells was non-significant (p-value > 0.05), confirming that the variances were equal and therefore that the interaction effects could be interpreted reliably.

Table 16. Results homoscedasticity test for H2 (Levene)

Dependent Variable	W0 (Levene)	df (degree of freedom)	P-value
Emotional Responses	1.2521833	(3, 111)	0.29441955
Trust	1.4728442	(3, 111)	0.22585524
Authenticity	1.0737156	(3, 111)	0.36331923

Purchase Intention	1.6199285	(3, 111)	0.18881971
--------------------	-----------	----------	------------

For H3a and H3b, the traditional Levene's test for homogeneity of variances between independent groups was not applicable as these hypotheses utilized within-subjects (repeated measures) designs.

4. Tests of the hypothesis

4.1. Hypothesis 1 (H1)

In order to test hypothesis H1, a Two-Way ANOVA was conducted for each dependent variable for target ad 1. Since the assumptions of normality for the dependent variables were met (see previous Jarque–Bera test), as well as the assumption of homogeneity of variances (see previous homoscedasticity, Levene's test), a Two-Way ANOVA was considered appropriate (Islam et al., 2022). This test is used to compare group means when the independent variable includes more than two levels (categories) (Islam et al., 2022). While Hypothesis 1 focuses on the difference between two label type groups, a Two-Way ANOVA was preferred over a One-Way ANOVA because it allows for the examination of the main effect of the Label Disclosure while simultaneously accounting for the product type (Hedonic vs. Utilitarian). In fact, including the variable product type in the model allows it to control its influence, reducing the error variance (the "noise") and also increasing the statistical power in order to detect the true effect of the AI label.

Then, it was followed by a pairwise comparison if the p-value was significant (post hoc test). In fact, if the ANOVA test is significant, it only indicates that not all group means are equal and does not mention which group-means are significantly different. Therefore, the use of the post hoc test is essential in order to identify them.

The results of the Two-Way ANOVA showed that, only for the variable authenticity, there was a main effect for label type with a p-value of 0.002 (p-value < 0.05).

Table 17. Results Two-Way ANOVA H1 (row label-type)

Dependent Variable	P-value (ANOVA)	Pairwise comparison	Result	Confirming hypothesis?
Emotional Responses	0.5754	No	Not significant.	No, H0 cannot be rejected.
Trust	0.5767	No	Not significant.	No, H0 cannot be rejected.

Authenticity	0.002	Yes	Significant	See pairwise comparison.
Purchase Intention	0.1189	No	Not significant.	No, H0 cannot be rejected.

4.1.1. Hypothesis H1a: Control vs. Truthful AI

Because Hypothesis 1a predicts a specific negative direction (an "AI Penalty"), the null hypothesis and the alternative hypothesis were defined as a directional test:

H0 = Disclosing an advertisement as AI-generated (Groups 2 & 5) has no significant negative effect on [Dependent Variable] compared to no disclosure (Groups 1 & 4). $\Rightarrow \mu_{AI} \geq \mu_{Control}$

Ha = Disclosing an advertisement as AI-generated (Groups 2 & 5) leads to significantly lower mean scores for [Dependent Variable] compared to no disclosure (Groups 1 & 4). $\Rightarrow \mu_{AI} < \mu_{Control}$ ⁶

As a result, the pairwise comparison for authenticity (using predictive margins) did not confirm that there were significant differences in authenticity between the AI label and No label conditions (p-value = 0.455). Therefore, the null hypothesis cannot be rejected.

4.1.2. Hypothesis 1b: Control vs. Human label

Since this hypothesis is suggesting that the bare presence of a label, even a human disclosure one could reduce the dependent variable, the following hypothesis were designed:
H0 (Null hypothesis) = The presence of a human-generated label does not reduce authenticity scores compared to the control condition or is equal/higher. $\Rightarrow \mu_{Human} \geq \mu_{Control}$

Ha (Alternative Hypothesis) = The presence of a human-generated label leads to lower authenticity scores compared to the control condition. $\Rightarrow \mu_{Human} < \mu_{Control}$

As a result of the pairwise comparison for authenticity, a statistically significant difference between conditions was found (p = 0.007), with the human-label producing lower scores than the control, no-disclosure one (mean difference = 0.6692). Therefore, the null hypothesis is rejected, in favor of the H1b, indicating that the simple presence of a label, even a human one, is associated with a reduction in perceived authenticity compared to the control condition.

⁶ μ represents the mean for that group.

4.1.3. Hypothesis 1c: (Truthful) AI label vs. (Misleading) Human label

Hypothesis 1c is focused on the first ad exposure, comparing the truthful AI disclosure (Groups 2 & 5) and the misleading human disclosure (Groups 3 & 6). Therefore, the hypotheses were defined as follow:

H0 (Null Hypothesis) = Truthfully disclosing an advertisement as AI-generated (Groups 2 & 5) has no significant negative effect on [Dependent Variable] compared to misleading human-generated disclosure (Groups 3 & 6). => $\mu_{\text{TruthfulAI}} \geq \mu_{\text{MisleadingHuman}}$

Ha (Alternative Hypothesis) = Truthfully disclosing an advertisement as AI-generated (Groups 2 & 5) leads to significantly lower mean scores for [Dependent Variable] compared to misleading human-generated disclosure (Groups 3 & 6). => $\mu_{\text{TruthfulAI}} < \mu_{\text{MisleadingHuman}}$

As a result, the predictive margins comparison for authenticity resulted in a statistically significant difference in the expected direction ($p = 0.001$), with higher mean scores observed in the misleading human condition compared to the truthful AI disclosure condition. Therefore, the null hypothesis is rejected in favor of the alternative hypothesis, clearly stating that perceived authenticity AI-label < perceived authenticity Human-label.

Table 18. Results of the Pairwise comparison H1 (a + b + c), Authenticity

Authenticity	Beta Coefficient / Contrast	P-value	Result	Confirming hypothesis?
H1a (Null Hypothesis): AI $\geq$ Control	- 0.1763384	0.455	Not significant.	No, H0 cannot be rejected : AI > Control
H1b (Null Hypothesis): Human $\geq$ Control	0.6692084	0.007	Significant.	H0 rejected, Ha approved: Human < Control
H1c (Null Hypothesis): AI $\geq$ Human	0.8455468	0.001	Significant.	H0 rejected, Ha approved: AI < Human

4.2. Hypothesis 2 (H2)

4.2.1. H2a: The Interaction Effect: Product Type and Label Type

In order to test the hypothesis H2a, a Two-Way ANOVA was used. This analysis specifically examined the interaction between label type (No Disclosure vs. Truthful AI) and product type

(Chocolate vs. Toothpaste) (Islam et al., 2022). This time the row of the interaction between product type and label type was taken into account.

The results found no significant interaction effect between product type and label type on any of the dependent variables. The lack of a significant interaction indicates that the impact of AI disclosure compared to no disclosure did not significantly differ between hedonic and utilitarian products. In addition, to address the structural model formula, a complementary linear regression analysis was performed, isolating the interaction beta coefficients for each dependent variable.

Table 19. Results Two-Way ANOVA H2a (row label-type + product-type)

Dependent Variable	Partial SS	df	F	P-value (Prob > F)	Interaction Beta	Result
Emotional Responses	0.52006786	1	0.43	0.5134	0.2698	Not significant
Trust	0.2939613	1	0.24	0.6269	0.2029	Not significant
Authenticity	5.8063142	1	3.66	0.0585	0.9016	Not significant
Purchase Intention	1.6299009	1	0.97	0.3261	0.4777	Not significant

4.2.2. H2b + H2c: Product Type Moderating Effects

Despite the non-significant interaction, a simple effects analysis (using the margins command) was performed to examine the specific patterns within each product category (chocolate and toothpaste). The results showed that the difference between the "No Label" group and the "AI Label" group for chocolate was only significant for purchase intention (H2a). However, because the interaction effect between product type and label type was not significant, the observed simple effect for chocolate should be interpreted with caution, as it does not constitute robust evidence of a difference between AI Disclosure and No-Disclosure across product categories. For toothpaste (H2b) no significant results were found.

Table 20. Results simple effects analysis

Dependent Variable	Product Type	Beta / dy/dx	Standard Error	t	P-value (P > t)	Results	Confirming hypothesis ?
--------------------	--------------	--------------	----------------	---	-----------------	---------	-------------------------

Emotional Responses	Chocolate	-0.1001195	0.289557	-0.35	0.730	Not Significant	No, H0 cannot be rejected.
	Toothpaste	0.1697271	0.292518	0.58	0.563	Not Significant	No, H0 cannot be rejected.
Trust	Chocolate	-0.1364994	0.2928186	-0.47	0.642	Not Significant	No, H0 cannot be rejected.
	Toothpaste	0.0663772	0.295813	0.22	0.823	Not Significant	No, H0 cannot be rejected.
Authenticity	Chocolate	-0.6216647	0.3317673	-1.87	0.064	Not Significant	No, H0 cannot be rejected.
	Toothpaste	0.2799835	0.3351601	0.84	0.405	Not Significant	No, H0 cannot be rejected.
Purchase Intention	Chocolate	-0.6762246	0.3407331	-1.98	0.05	Significant	H0 rejected, Ha approved.
	Toothpaste	-0.1985112	0.3442175	-0.58	0.565	Not Significant	No, H0 cannot be rejected.

4.3. Hypothesis 3 (H3)

This hypothesis looks at the within-subjects part of the study, focusing on how participants' responses change from the first exposure of the advertisement to the second exposure of the same advertisement but with the opposite label. To test H3, a paired-samples t-test was conducted for each dependent variable, comparing target 1 and target 2.

This test is appropriate here because it compares two related measurements taken from the same participants, meaning it evaluates whether the average difference between the two conditions is significantly different from zero (Islam et al., 2022). In this case, it allows us to directly test whether responses shift from the first exposure to the second exposure. Because hypotheses 3a and 3b point towards a directional decrease/increase, unilateral (one-tailed) paired t-tests were employed, isolating the alternative hypothesis to determine significance.

However, as the normality assumption of the difference between the two measurements (first exposure and second exposure) was not met for most of the dependent variables (see previous Shapiro-Wilk test for normality), the sufficient size of the sample ($H3a = 49$ and $H3b = 57$) allows us to apply the Central Limit Theorem, ensuring that the sample distribution of the mean is approximately normal (Islam et al., 2022). While using a paired t-test remains robust here, to further validate these findings, a non-parametric Wilcoxon signed-rank test was also conducted as a robustness check (see appendix 5.13 and 5.14 for further details). Therefore, as the results for both Student's paired t-test and Wilcoxon were always in accordance with each other, the t-test result is the one figuring in this analysis.

4.3.1. Hypothesis 3a: The "Caught in a Lie" Penalty

The objective is to determine whether revealing the AI label after an initial Human label (Groups 3 & 6) leads to a significant drop in the dependent variables, consistent with a "caught in a lie" effect driven by perceived deception and consumer backlash. Therefore, the hypotheses are defined as follows:

H_0 (Null Hypothesis) = The shift from a Human-labeled first exposure to an AI-labeled second exposure does not lead to a significant decrease in [Dependent Variable] $\Rightarrow \mu_{t1_Human} \leq \mu_{t2_AI}$

H_a (Alternative Hypothesis) = The shift from a Human-labeled first exposure to an AI-labeled second exposure leads to a significant decrease in [Dependent Variable]. $\Rightarrow \mu_{t1_Human} > \mu_{t2_AI}$

Table 21. Results t-test for H3a

Dependent Variable t1 - Dependent variable t2	Mean	Std.Dev	t-value	One-tailed P-value (Prob T > t)	Result	Confirming hypothesis?
Emotional Responses	0.4653061	0.7856926	4.1456 (48 degrees of freedom)	0.0001	Significant	H_0 rejected, H_a approved.
Trust	0.5306122	0.9569274	3.8815 (48 degrees of freedom)	0.0002	Significant	H_0 rejected, H_a approved.

Authenticity	0.829932	1.196295	4.8563 (48 degrees of freedom)	0.0000	Significant	H0 rejected, Ha approved.
Purchase Intention	0.5102041	1.04106	3.4306 (48 degrees of freedom)	0.0006	Significant	H0 rejected, Ha approved.

As a result, the t-test showed a statistically significant difference in all four dependent variables (all p-values < 0.0006) in the expected direction, as the difference in means of the t1 and t2 for all four variables is positive. Therefore, the null hypothesis is rejected in favor of the alternative hypothesis, indicating that the shift from a first exposure to a Human-labeled ad, to a subsequent exposure to an AI-labeled ad leads to a significant decrease in emotional responses, trust, authenticity and purchase intentions.

4.3.2. H3b: The Human Premium Shift

The objective is to determine whether switching from an AI Label to a subsequent Human label (Groups 2 & 5) leads to a significant increase in the dependent variables, according to the idea that consumers value human authorship more than AI, creating a positive shift when the human origin is revealed. Therefore, the hypotheses are defined as follows:

H0 (Null Hypothesis) = The shift from an AI-labeled first exposure to a Human-labeled second exposure does not lead to a significant increase in [Dependent Variable] => $\mu_{t1_AI} \geq \mu_{t2_Human}$

Ha (Alternative Hypothesis) = The shift from an AI-labeled first exposure to a Human-labeled second exposure leads to a significant increase in [Dependent Variable]. => $\mu_{t1_AI} < \mu_{t2_Human}$

Table 22. Results t-test for H3b

Dependent Variable t1 - Dependent variable t2	Mean	Std. Dev	t-value	One-tailed P-value (Prob T < t)	Result	Confirming hypothesis?
Emotional Responses	0.2035088	0.7116291	2.1591 (56 degrees of freedom)	0.9821 (-> Prob T > t = 0.0176)	Not Significant (Opposite effect)	No, H0 cannot be rejected.

Trust	0.0263158	0.9690342	0.2050 (56 degrees of freedom)	0.5809	Not Significant	No, H0 cannot be rejected.
Authenticity	-0.2807017	0.9713954	-2.1817 (56 degrees of freedom)	0.0167	Significant	H0 rejected, Ha approved.
Purchase Intention	- 0.26331579	0.9144711	-2.1726 (56 degrees of freedom)	0.0170	Significant	H0 rejected, Ha approved.

As a result, the unilateral t-test showed a statistically significant difference in the expected direction for authenticity and purchase intention (p-value = 0.0167 for authenticity, and p-value = 0.0170 for purchase intention). As the difference in means is negative (for authenticity = -0.2807017, and for purchase intention = -0.26331579), it proves that the scores significantly increased with the second exposure. Therefore, the null hypothesis is rejected in favor of the alternative hypothesis for these two variables, indicating that the shift from a first exposure to an AI-labeled ad, to a subsequent exposure to a Human-labeled ad leads to a significant increase in authenticity and purchase intention, but not in emotional responses and trust.

5. Exploratory tests

5.1. Hypothesis 1

This exploratory section investigates the moderating role of consumer characteristics. In fact, by analyzing the demographic factors (age, gender, education level, profession) alongside behavioral and psychological variables (purchase frequency and the General Attitudes towards Artificial Intelligence Scale (GAAIS)), it was intended to see if some specific consumer segments were more impacted depending on the label type. To do that, a multiple regression was realized, and the homoscedasticity of the residuals was verified by using the Breusch-Pagan test for each hypothesis.

5.1.1. H1a

- Homoscedasticity test

The results showed that all p-values were above 0.05 meaning that we cannot reject the null hypothesis and that therefore the spread of the residuals is constant.

Table 23. Results homoscedasticity of the residuals test for H1a (Breusch-Pagan)

Dependent Variable	Chi(2)	P-value
Emotional Responses	0.3	0.5846
Trust	0.03	0.8613
Authenticity	1.73	0.1880
Purchase Intention	0.85	0.3556

- Regression

Four exploratory multiple regressions were conducted to observe if some factors moderated the impact on emotional responses, trust, authenticity and purchase intention of the AI label. Emotional responses revealed a significant (at 5%) moderating effect of education. In fact, it was found that holding a master's degree or higher significantly altered positively the emotional responses, when faced with an AI label compared to the baseline ($\beta = 2.31$, $p = 0.046$). While no other results were significant at the 5% level, three additional effects were observed when applying a 10% significance threshold (see table below).

Table 24. Results regression for H1a

Dependent Variable	Significant Category	Coefficient	P-value
Emotional Responses	Education level => Master or Higher	2.305697	0.046** ⁷
Emotional Responses	Profession => Business, Social Sciences, etc.	-1.055306	0.095* ⁸
Trust	Age => Between 45 and 54	-2.744333	0.053*
Purchase Intention	Education Level => Master or Higher	2.330389	0.074*

5.1.2. H1b

- Homoscedasticity test

The results showed that all p-values were higher than 0.05, meaning that we cannot reject H0.

Table 25. Results homoscedasticity of the residuals test for H1b (Breusch-Pagan)

Dependent Variable	Chi(2)	P-value
Emotional Responses	0.24	0.6226

⁷ Significant at the threshold of 5%

⁸ Significant at the threshold of 10%

Trust	0	0.9689
Authenticity	1.53	0.2158
Purchase Intention	0.19	0.6621

- Regression

Four exploratory multiple regressions were conducted to observe if some factors moderated the impact on emotional responses, trust, authenticity and purchase intention when seeing no label in comparison to a Human label. The only significant result at 5% found was for the variable “Emotional Responses”. In fact, the frequency of purchase positively impacted the emotional responses when a human label was disclosed instead of no disclosure ($\beta = 2.46$, $p = 0.010$). Two additional effects were observed when applying a 10% significance threshold (see table below).

Table 26. Results regression for H1b

Dependent Variable	Significant Category	Coefficient	P-value
Emotional Responses	Frequency of Purchase	2.46063	0.010**
Trust	Profession => Other	2.541049	0.099*
Trust	Frequency of Purchase	1.793182	0.080*

5.1.3. H1c

- Homoscedasticity test

The results found that all p-values were above the threshold of 5% meaning that the null hypothesis cannot be rejected.

Table 27. Results homoscedasticity of the residuals test for H1c

Dependent Variable	Chi(2)	P-value
Emotional Responses	0.21	0.6496
Trust	0	0.9728
Authenticity	1.55	0.2126
Purchase Intention	0.03	0.8732

- Regression

Four exploratory multiple regressions were conducted to observe if some factors moderated the impact on emotional responses, trust, authenticity and purchase intention when seeing the

AI label in comparison to the Human label. The only significant result at 5% found was for the variable “Trust”. In fact, the frequency of purchase positively impacted the trust when a human label was disclosed instead of an AI label ($\beta = 2.40$, $p = 0.021$). Two additional effects were observed when applying a 10% significance threshold (see table below).

Table 28. Results regression for H1c

Dependent Variable	Significant Category	Coefficient	P-value
Emotional Responses	Age => Between 45 and 54	2.621836	0.066*
Trust	Frequency of Purchase	2.399839	0.021**
Purchase Intention	Age => Between 25 and 34	1.929164	0.091*

5.2. Hypothesis 2

This exploratory section investigates the moderating role of consumer characteristics. In fact, it was intended to see if some specific consumer segments were impacted depending on the interaction between the label type (No Disclosure vs AI) and the product type (Chocolate vs Toothpaste). To do that, a multiple regression was realized, and the homoscedasticity of the residuals was verified by using the Breusch-Pagan test for each hypothesis.

- Homoscedasticity test

The results showed that all p-values were higher than 0.05, meaning that we cannot reject H0.

Table 29. Results homoscedasticity of the residuals test for H2

Dependent Variable	Chi(2)	P-value
Emotional Responses	0.32	0.5705
Trust	0.03	0.8735
Authenticity	2.06	0.1510
Purchase Intention	0.91	0.3410

- Regression

Four exploratory multiple regressions were conducted. The only significant result at 5% found was for the variable “Emotional Responses”. In fact, having a Master or Higher impacted positively the interaction ($\beta = 2.32$, $p = 0.047$). Two additional effects were observed when applying a 10% significance threshold (see table below).

Table 30. Results regression for H2

Dependent Variable	Significant Category	Coefficient	P-value
Emotional Responses	Education Level => Master or Higher	2.322321	0.047**
Trust	Age => Between 45 and 54	-2.754545	0.054*
Purchase Intention	Education Level => Master or Higher	2.305747	0.08*

5.3. Hypothesis 3

This exploratory section investigates the moderating role of consumer characteristics. The goal is to determine if specific consumer segments were influenced differently and changed their perceptions when rating the second target advertisement.

To do that, a multiple regression was realized, and the homoscedasticity of the residuals was verified by using the Breusch-Pagan test for each hypothesis. When homoscedasticity was not found, `vce(robust)`⁹ was aggregated to maintain a heteroscedastic assumption.

5.3.1. H3a

- Homoscedasticity test

The results showed that no p-values were higher than 0.05 meaning that the null hypothesis can be rejected for all variables and therefore that the spread of errors is not constant.

Table 31. Results homoscedasticity of the residuals test for H3a

Dependent Variable	Chi(2)	P-value
Emotional Responses	4.91	0.0266
Trust	3.91	0.0479
Authenticity	6.54	0.0105
Purchase Intention	5.81	0.0159

- Regression

Four exploratory multiple regressions were conducted to observe if some factors moderated the impact on the dependent variables when seeing the second target ad with an AI label. It was added `vce(robust)` to these regressions as heteroscedasticity of the residuals was found. For the variable “Emotional Responses”, no significant results were found, as every p-value was higher

⁹ It represents the command in stata in order to correct heteroscedasticity.

than 0.05. For the variable “Trust”, two significant results were found. In fact, having a University Bachelor or a Master or Higher impacted their trust positively ($\beta=1,40$ and $p=0.001$ for University Bachelor and $\beta=1.22$ and $p=0.021$ for Master or Higher). For the variable “authenticity”, one significant result was found, and it is that having a university bachelor also impacted the authenticity positively ($\beta=1.16$ and $p=0.048$). For the variable “purchase intention”, three significant and positive results were found related to the profession (for Academia and Researchers, STEM and Healthcare and Business, Social Sciences and Humanities). Moreover, one additional effect was observed when applying a 10% significance threshold (see table below).

Table 32. Results regression for H3a

Dependent Variable	Significant Category	Coefficient	P-value
Trust	Education Level => University Bachelor	1.397735	0.001*** ¹⁰
Trust	Education Level => Master or Higher	1.214463	0.021**
Authenticity	Education Level => University Bachelor	1.155746	0.048**
Purchase Intention	Age => Between 25 and 34	-1.281133	0.022**
Purchase Intention	Profession => Academia and Researchers	2.142563	0.003***
Purchase Intention	Profession => STEM and Healthcare	1.376495	0.081*
Purchase Intention	Profession => Business, Social Sciences, etc.	1.327828	0.014**

5.3.2. H3b

- Homoscedasticity test

The results found that the p-values for emotional responses, trust and purchase intention were above 0.05 meaning that we cannot reject the null hypothesis for these variables. However, for authenticity, its p-value was smaller than 0.05 meaning that H0 is rejected for this variable.

Table 33. Results homoscedasticity of the residuals test for H3b

Dependent Variable	Chi(2)	P-value
Emotional Responses	2.80	0.0941

¹⁰ Significant at the threshold of 1%

Trust	2.08	0.1490
Authenticity	15.99	0.0001
Purchase Intention	2.49	0.1144

- Regression

Four exploratory multiple regressions were conducted to observe if some factors moderated the impact on the dependent variables when seeing the second target ad with a Human label. For the multiple regression of authenticity, $vce(robust)$ was aggregated, as homoscedasticity was not found. For the variables “Emotional Responses” and “Trust, no significant results were found, as every p-value was higher than 0.05. For the variable “authenticity”, two significant results were found. The first one is that women rated the advertisement with a Human label more positively than men ($\beta=0.74$ and $p=0.033$). The second is that the persons aged between 45 and 54 rated the ad with a Human Label less great than the persons aged between 18 and 24 ($\beta= -2.46$ and $p=0.000$). For the variable “purchase intention”, one significant and negative result was found related to the general attitudes towards artificial intelligence ($\beta= -0.46$ and $p=0.038$). Moreover, three additional effects were observed when applying a 10% significance threshold (see table below).

Table 34. Results regression for H3b

Dependent Variable	Significant Category	Coefficient	P-value
Trust	Education Level => Master or Higher	1.465282	0.070*
Authenticity	Gender => Female	0.7352448	0.033**
Authenticity	Age => Between 45 and 54	-2.461573	0.000***
Purchase Intention	Age => Between 25 and 34	0.6605406	0.087*
Purchase Intention	Age => Between 35 and 44	1.124527	0.051*
Purchase Intention	GAAIS	-0.4564266	0.038**

V. Discussion

This research examined how AI disclosure labels in advertisements shape consumer responses (emotional responses, trust, perceived authenticity and purchase intention). Furthermore, this study tested whether AI disclosure always activates a negative reaction or not, basing this assumption on Persuasion Knowledge Model, algorithm aversion theory, and the Stimulus–Organism–Response Model framework. The results showed that consumer responses to AI-generated advertising were not universally negative and that they varied according to the specific category of response (emotions, trust, authenticity, purchase intention), the product type, and whether the participant was exposed to the AI label for the first target advertisement or the second.

1. AI Disclosure and the Labeling Paradox (H1)

1.1. Absence of a General AI Penalty (H1a)

No significant difference was found between the AI disclosure condition and the No-Disclosure baseline for the four dependent variables. This result goes against the algorithm aversion theory and other earlier studies. In fact, it challenges the assumption that when the utilization of AI is revealed, it automatically harms consumer reactions and evaluations meaning that for this specific sample and in these advertising settings, algorithm aversion does not happen directly.

One possible explanation could be the quality of the stimulus. In fact, if the AI-generated images appear professionally designed and even visually attractive, consumers may judge the advertisement based on the outcome “Does the ad looks good in general?” instead of just paying attention to the source. This explanation aligns with the idea of functional honesty, which states that if the quality is high, the transparency about AI involvement may not interfere in the judgment (remain neutral).

Another possible explanation is that the participants in the control groups could have guessed that the advertisements were made with AI, therefore judging it negatively. Another reason behind this result could be related to the characteristics of the participants and therefore to this specific sample.

1.2. The Label Scrutiny Effect: Human Attribution Backfire (H1b)

Between the No-disclosure condition and the Human disclosure condition, a significant difference was found which is that a human label significantly reduces the perceived authenticity. The reason behind this can be related to the Persuasion Knowledge Model, as explicit persuasive signals arise questions by increasing the vigilance and the skepticism of

consumers. As consumers directly assume that a non-disclosed advertisement is human-made, the label might feel unnecessary. As a result, the consumer questions the presence of the label and, consequently, reduce their perceived authenticity of the advertisement.

1.3. The Human Effort Heuristic (H1c)

Between the AI disclosure condition and the Human disclosure one, a significant difference in terms of authenticity was found. In fact, participants rated the human-labeled advertisement as much more authentic than the AI-labeled one. This finding might be explained by the fact that consumers tend to associate human effort with more emotional depth and higher quality, which is corroborated by the Human Effort Heuristic theory.

To conclude this point, it is important to mention that this preference for human-made coexists with the absence of an AI penalty found in H1a. This means that participants do not necessarily reject AI or rate lower when it is AI-made but they still favor human authorship. Therefore, H1 reveals divergent effects where AI disclosure does not directly generate negative reactions, but the presence of a human-label does.

2. Product Type and Hedonic Vulnerability (H2)

2.1. Absence of an Interaction Effect (H2a)

No significant interaction effect was found between the product type and the label type, meaning that H2a was not supported. This finding suggests that participants processed the product category and the label type independently rather than combined. In fact, it could mean that consumer's psychological reaction to a warning or informational label is universal. Therefore, whether the label is placed on an essential product, a luxury item, or a hedonic product, consumers process and judge the label the same way. However, another possible explanation could be a divergence in the perceived quality or attractiveness of the product images. For instance, if the hedonic product image was evaluated less positively than the utilitarian one, this may have reduced potential differences between conditions.

2.2. Hedonic Products (H2b) and Utilitarian Products (H2c)

Looking at the results of the simple effects analysis for chocolate, AI disclosure significantly reduced purchase intention. However, this effect was not accompanied by significant differences in emotional response, trust, or perceived authenticity, meaning that the reasons behind this result remains unclear and should be explored in future research. No significant results were found for toothpaste and, therefore H2c was not supported.

3. The Cost of Deception: Effects of Revelation (H3)

3.1. The “Caught in a Lie” Effect (H3a)

H3a was strongly supported across all four dependent variables. In fact, when participants first saw a Human-made label and later, the same ad with an AI-made label, their emotional responses, trust, authenticity, and purchase intention decreased sharply. This effect can be reflected by the Deceptive Denial Theory and the Persuasion Knowledge Model. In fact, consumers did not react negatively because of AI itself but because they felt betrayed by the misleading disclosure. Therefore, their expectations were violated, which triggered moral and emotional rejection.

3.2. The Inconsistency Effect (H3b)

For the reverse sequence, weaker effects were found. When participants first saw an AI-made label and later a Human-made label for the same advertisement, perceptions of authenticity and purchase intention significantly increased. This effect may be explained by the Human Effort Heuristic, which suggests that consumers tend to value content more positively when they believe human effort was involved.

However, emotional responses showed the opposite pattern. One possible explanation is that participants perceived the switch from an AI-made label to a Human-made label as inconsistent. This could have created a “Caught in a Lie” effect, where consumers react negatively because the communication feels less coherent and transparent.

4. Exploratory Evidence: Consumer Heterogeneity

The results of the exploratory regressions revealed several significant moderating effects associated with demographic characteristics.

4.1. Education level and Algorithm Acceptance

For education level, it was found that having a master’s degree or more increases acceptance of AI-generated content, especially for two of the four dependent variables: emotional responses and purchase intention. This effect is supported by the idea of functional literacy, where more educated consumers view AI as a production tool that optimize content creation (Ahac, 2025)

4.2. Age-Based Divergence

Older participants (particularly those aged between 45–54) have stronger negative reactions to AI disclosure, especially in terms of trust. This aligns with the frameworks of interpretive

vigilance, where older people rely more on human reasoning strategies (Ahac, 2025). In fact, they evaluate information critically instead of simply accepting AI-generated content. On the contrary, younger participants (specifically between 25-34) have more positive reactions towards AI, as they may see it as an efficient tool.

4.3. Gender

A significant difference in terms of gender has been found in H3b in terms of authenticity. In fact, women tend to have more positive responses to Human-made labels following AI exposure than men. In fact, one potential explanation could be that women fear that if they use AI, it will be considered as cheating and, therefore, that they will experience an even higher “competence penalty”. As a result, women may react more positively to signals emphasizing human involvement.

4.4. Attitudes Toward AI (GAAIS)

Participants who expressed more favourable general attitudes towards AI demonstrated a significant reduction in terms of purchase intention when transitioning from an AI-labelled to a human-labelled advertisement (H3b). These results suggest that consumers with a greater familiarity with AI place a higher value on technological honesty than on human-made elements. Consequently, the revelation of a misleading disclosure is likely to be perceived negatively.

VI. Conclusion

This thesis has investigated consumer reactions to the disclosure of AI-generated advertisements. By studying both chocolate and toothpaste visuals, the project analyzed how product categories modify these specific consumer responses. The study used an experimental approach and gathered data from 164 participants.

1. Summary of Key Findings

AI-labelled advertisements did not trigger negative consumers' responses. When evaluated in isolation, disclosure left trust, emotions, and purchase intentions completely unchanged. However, a human label reduced perceived authenticity when compared to no label. This unexpected effect suggests that explicitly claiming human origin activates the persuasion knowledge model and increases skepticism. However, when comparing human disclosure and AI disclosure, participants preferred human-made advertisements. This effect could be explained by the heuristics based on perceived human effort. To continue, an interaction effect between product categories and label type has not proven significant results.

Misleading labeling caused significant effects. In fact, revealing AI origin after claiming human authorship caused trust, authenticity, and purchase intention to decrease. This confirmed the "caught in a lie" effect. Conversely, the reverse sequence produced weaker effects. As a matter of fact, switching from an AI to a human label caused a positive effect for the human one in terms of authenticity and purchase intention.

2. Limitations and future research

The first limitation concerns the randomization process. It is demonstrated that respondents frequently choose midpoint options, therefore showing that letting the respondents manually choose a number between 1-6 can distort results, as consumers avoid extreme categories. However, even though this is true, this study demonstrated that the results were robust enough with the manual randomization, as shown by the independence chi-square test on demographics showing no significative difference between the six groups.

The second limitation is related to the stimuli. As only one advertisement per product type (i.e. one chocolate ad and one toothpaste ad) was shown to participants, the responses can be biased by the characteristics of the images themselves (such as the lighting, the color preference, the aesthetics), therefore limiting the observable interaction effect of the product category or just image attributes (Judd et al., 2012). However, due to budget and time constraints, we had to accept this potential bias to ensure that participants will complete the whole questionnaire

attentively, especially as they didn't have any compensation. Further research could use multiple product categories with various images for each condition, therefore using random effects to control image-specific variances and ensure true generalizability.

The third limitation is that the within-subjects part of this study might be biased (Hypothesis 3). In fact, because participants were exposed to two identical ads but with different labels, they could have guessed the true purpose of the study. This creates a demand effect, where participants might unconsciously align their behavior to match what they think is expected of them (i.e. punishing AI because they feel they should prefer human-made ads, or trying to reduce the penalty (Shimp et al., 1991). However, to limit the damage on the internal validity of the study, multiple mechanisms were implemented. First, we used filler ads (three in-between, and one before and after the target ads) to distract participants and stop them from comparing the images too easily. Second, they were told to look at the advertisements like they would in real life, with a short exposure of 10-15 seconds per image to get more natural reactions. Third, even if participants guessed that the study was about labels, they could not guess the exact interaction being tested, which ensures the validity of the data. For future research, further studies could use implicit attitude tests or eye-tracking metrics rather than self-reported ones to completely avoid these conscious biases. Another option would be to increase the number of filler ads. While this would make the survey longer, offering a good compensation would keep participants motivated and help successfully hide the objectives of the study.

The fourth limitation is the fact that a convenience sample was used, making it unrepresentative of the wider population and not easily generalizable, therefore compromising the external validity of the experiment (Wild et al., 2022). For instance, a large majority of the respondents here are young (specifically, 55 aged 18–24 and 59 aged 25–34) and highly educated (47 having a university bachelor's degree and 93 a master's or higher). As these variables were included as moderators for the exploratory hypotheses, the analysis required comparing groups that are heavily unequal, which explains why the interaction effects were not significant. However, demographics do matter, but our study could not prove their impact, as the sample is not representative of the public and instead mirrors an academic environment of students. Therefore, future research operating without the strict budget, time, and resource constraints could address these gaps by using a representative sample over a longer timeframe to allow the generalizability of the findings to the wider population.

The fifth limitation involves the selection and validation process of the visual stimuli used in the experiment. While this pre-selection of images was an insightful step to do, the fact that it relied on a 1-to-3 grading scale evaluated by a one single digital marketing professional introduces a risk of subjectivity or selection bias. In a larger-scale study, this limitation could be addressed by employing a larger, diverse panel of multiple independent creative experts, ensuring a highly robust and completely unbiased stimulus validation.

The sixth limitation is that strict internal validity is difficult to achieve in an online environment because controlling extraneous variables is more challenging than in field or laboratory experiments. For example, it was difficult to verify the participants' engagement, honesty and reliability. Furthermore, it was almost impossible to control the environmental conditions in which the participants responded, as they completed the survey on their personal devices in a variety of uncontrolled settings (e.g. at home, at work, on public transport, or in the street). Additionally, as the survey appeared differently on phones than on computers, screen size and load times could have subtly influenced how participants interacted with the survey. Furthermore, participants who responded in noisy environments may have reflected less deeply than those in calm environments. For future research, the survey could be realized in a lab environment instead of an online environment as it will largely improve its internal validity by controlling for extraneous variables.

The seventh limitation is that the stimuli were static images, meaning the results cannot be applied to dynamic media, such as videos. Future research could examine advertisements in video format.

The last limitation is the probability of participant contamination at the moment of the distribution of the online questionnaire. In fact, because of the sample of convenience and the snowball effect, a risk exists that respondents who forwarded the survey link to peers discussed the underlying nature of the experiment beforehand. This could have biased their responses.

VII. Bibliography

- Aburass, S., & Abu Rumman, M. (2024). Authenticity in authorship: The Writer's Integrity framework for verifying human-generated text. *Ethics and Information Technology*, 26(3), 62. <https://doi.org/10.1007/s10676-024-09797-z>
- Ahac, M. (2025). *Consumer Interpretations of AI-generated Marketing: Trust, Authenticity and Emotional Tone* [Master Degree Project in Business Administration, University of Skövde]. <https://www.diva-portal.org/smash/get/diva2:1986003/FULLTEXT01.pdf>
- American Marketing Association. (n.d.). Marketing vs. Advertising. *American Marketing Association*. Retrieved 27 May 2026, from <https://www.ama.org/marketing-vs-advertising/>
- Bai, L., Huang, D., Liu, X., Lu, Z., Ouyang, W., Qu, J., & Wu, C. (2023). Seeing is not always believing: Benchmarking Human and Model Perception of AI-Generated Images. *Advances in Neural Information Processing Systems* 36, 25435–25447. <https://doi.org/10.52202/075280-1105>
- Bellaiche, L., Shahi, R., Turpin, M. H., Ragnhildstveit, A., Sprockett, S., Barr, N., Christensen, A., & Seli, P. (2023). Humans versus AI: Whether and why we prefer human-created compared to AI-created artwork. *Cognitive Research: Principles and Implications*, 8(1), 42. <https://doi.org/10.1186/s41235-023-00499-6>
- Beverland, M. B., & Farrelly, F. J. (2010). The Quest for Authenticity in Consumption: Consumers' Purposive Choice of Authentic Cues to Shape Experienced Outcomes. *Journal of Consumer Research*, 36(5), 838–856. <https://doi.org/10.1086/615047>
- Bimbo, F., Roselli, L., Carlucci, D., & De Gennaro, B. C. (2020). Consumer Misuse of Country-of-Origin Label: Insights from the Italian Extra-Virgin Olive Oil Market. *Nutrients*, 12(7), 2150. <https://doi.org/10.3390/nu12072150>
- Brooks, J. B. (1948). The Organization of the Selling Function in a Department Store. *Journal of Marketing*, 13(2), 189–194. <https://doi.org/10.1177/002224294801300206>
- Buder, F. (2024). Transparency Without Trust. *NIM INSIGHTS Research Magazin*, 7. https://www.nim.org/fileadmin/PUBLIC/3_NIM_Publikationen/NIM_INSIGHTS_Research_Magazin/VOL_7/Downloads/9_Transparency_Without_Trust.pdf
- Buder, F., & Unfried, M. (n.d.). *AI-labeling rule presents a challenge: Research shows that consumers become more skeptical and less engaged when they know content is AI-generated.*

- Cambier, F., & Poncin, I. (2020). Inferring brand integrity from marketing communications: The effects of brand transparency signals in a consumer empowerment context. *Journal of Business Research*, 109, 260–270. <https://doi.org/10.1016/j.jbusres.2019.11.060>
- Casenave, E., & Klarmann, M. (2020). The accountability paradox: How holding marketers accountable hinders alignment with short-term marketing goals. *Journal of Business Research*, 112, 95–108. <https://doi.org/10.1016/j.jbusres.2020.02.047>
- 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>
- Chaisatitkul, A., Luangngamkhum, K., Noulpum, K., & Kerdvibulvech, C. (2024). The power of AI in marketing: Enhancing efficiency and improving customer perception through AI-generated storyboards. *International Journal of Information Technology*, 16(1), 137–144. <https://doi.org/10.1007/s41870-023-01661-5>
- Chauhan, P. S. (2025). *Effect Of Brand Transparency On Consumer Trust And Purchase Intention*. 15(2). <https://rjpn.org/ijcspub/papers/IJCSP25B1124.pdf>
- Chen, H., Chan-Olmsted, S., Kim, J., & Mayor Sanabria, I. (2022). Consumers' perception on artificial intelligence applications in marketing communication. *Qualitative Market Research: An International Journal*, 25(1), 125–142. <https://doi.org/10.1108/QMR-03-2021-0040>
- Claudy, M. C., Aquino, K., & Graso, M. (2022). Artificial Intelligence Can't Be Charmed: The Effects of Impartiality on Laypeople's Algorithmic Preferences. *Frontiers in Psychology*, 13, 898027. <https://doi.org/10.3389/fpsyg.2022.898027>
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd edition, pp. 1–17). Lawrence Erlbaum Associates.
- Cui, W., Liu, M. J., & Yuan, R. (2025a). Exploring the Integration of Generative AI in Advertising Agencies: A Co-Creative Process Model for Human–AI Collaboration. *Journal of Advertising Research*, 65(2), 167–189. <https://doi.org/10.1080/00218499.2024.2445362>
- Dal Palù, D., Buiatti, E., Puglisi, G. E., Houix, O., Susini, P., De Giorgi, C., & Astolfi, A. (2017). The use of semantic differential scales in listening tests: A comparison between context and laboratory test conditions for the rolling sounds of office chairs. *Applied Acoustics*, 127, 270–283. <https://doi.org/10.1016/j.apacoust.2017.06.016>

- Degirmenci, T., Mathmann, F., & Mortimer, G. (2025). Made With Love: Examining Consumer Engagement With Handmade Versus Machine-Made Production Cues. *Journal of Consumer Behaviour*, 24(2), 982–1002. <https://doi.org/10.1002/cb.2455>
- Dhar, R., & Wertenbroch, K. (2000). Consumer Choice between Hedonic and Utilitarian Goods. *Journal of Marketing Research*, 37(1), 60–71. <https://doi.org/10.1509/jmkr.37.1.60.18718>
- 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>
- Edwards, K. (2024, June 25). What is Semantic Differential Scale, and How to Use It? *SurveyLab*. <https://www.surveylab.com/blog/semantic-differential-scale/>
- Farjam, S., & Hongyi, X. (2015). Reviewing the Concept of Brand Equity and Evaluating Consumer-Based Brand Equity (CBBE) Models. *The International Journal of Management Science and Business Administration*, 1(8), 14–29. <https://doi.org/10.18775/ijmsba.1849-5664-5419.2014.18.1002>
- Frank, B., Herbas-Torrico, B., & Schvaneveldt, S. J. (2021). The AI-extended consumer: Technology, consumer, country differences in the formation of demand for AI-empowered consumer products. *Technological Forecasting and Social Change*, 172, 121018. <https://doi.org/10.1016/j.techfore.2021.121018>
- Fuchs, C., Schreier, M., & Van Osselaer, S. M. J. (2015). The Handmade Effect: What's Love Got to Do with It? *Journal of Marketing*, 79(2), 98–110. <https://doi.org/10.1509/jm.14.0018>
- Gao, X., Li, W., & Zhao, Y. (2026). Value-dependent and empathy-mediated: How artificial intelligence-generated marketing content influences customer engagement, and when to disclose its origin. *Frontiers in Psychology*, 16, 1701085. <https://doi.org/10.3389/fpsyg.2025.1701085>
- General Data Protection Regulation (GDPR) – Legal Text*. (n.d.). General Data Protection Regulation (GDPR). Retrieved 27 May 2026, from <https://gdpr-info.eu/>
- Golzar, J., Noor, S., & Tajik, O. (2022). Convenience Sampling. *International Journal of Education Language Studies*, 1(2). <https://doi.org/10.22034/ijels.2022.162981>

Höltervennhoff, S., Ricker, J., Raphael, M. M., Schwedes, C., Weil, R., Fischer, A., Holz, T., Schönherr, L., & Fahl, S. (2025). *Security Benefits and Side Effects of Labeling AI-Generated Images* (Version 2). arXiv. <https://doi.org/10.48550/ARXIV.2505.22845>

Islam, M., Kabir, R., & Nisha, M. (2022). *Data Analysis with Stata A Comprehensive Guide for Data Analysis and Interpretation of Outputs First Edition Stata Version 13*.

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. <https://doi.org/10.1037/a0028347>

Kakaria, S., Simonetti, A., & Bigne, E. (2024). Interaction between extrinsic and intrinsic online review cues: Perspectives from cue utilization theory. *Electronic Commerce Research*, 24(4), 2469–2497. <https://doi.org/10.1007/s10660-022-09665-2>

Koning, B., & Voorveld, H. A. M. (2025). Disclaimer! This Content Is AI-Generated: How AI-Disclosures Influence Trust in Advertisements and Organizations. *Journal of Interactive Advertising*, 25(3), 240–253. <https://doi.org/10.1080/15252019.2025.2554149>

Kruger, J., Wirtz, D., Van Boven, L., & Altermatt, T. W. (2004). The effort heuristic. *Journal of Experimental Social Psychology*, 40(1), 91–98. [https://doi.org/10.1016/S0022-1031\(03\)00065-9](https://doi.org/10.1016/S0022-1031(03)00065-9)

Kuhn, G. (2023, May 4). Semantic Differential Scales Explained + Examples. *Drive Research*. <https://www.driveresearch.com/market-research-company-blog/what-is-a-semantic-differential-scale/>

Lim, S., & Schmälzle, R. (2024). The effect of source disclosure on evaluation of AI-generated messages. *Computers in Human Behavior: Artificial Humans*, 2(1), 100058. <https://doi.org/10.1016/j.chbah.2024.100058>

Liu, Y., Heinberg, M., Huang, X., & Eisingerich, A. B. (2023). Building a competitive advantage based on transparency: When and why does transparency matter for corporate social responsibility? *Business Horizons*, 66(4), 517–527. <https://doi.org/10.1016/j.bushor.2022.10.004>

- Longoni, C., & Cian, L. (2022). Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The “Word-of-Machine” Effect. *Journal of Marketing*, 86(1), 91–108. <https://doi.org/10.1177/0022242920957347>
- Loth, A., Kappes, M., & Pahl, M.-O. (2026). *Can Humans Tell? A Dual-Axis Study of Human Perception of LLM-Generated News* (Version 1). arXiv. <https://doi.org/10.48550/ARXIV.2604.03755>
- MacDorman, K. F., Green, R. D., Ho, C.-C., & Koch, C. T. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25(3), 695–710. <https://doi.org/10.1016/j.chb.2008.12.026>
- MacKenzie, S. B., & Lutz, R. J. (1989). An Empirical Examination of the Structural Antecedents of Attitude toward the Ad in an Advertising Pretesting Context. *Journal of Marketing*, 53(2), 48–65. <https://doi.org/10.1177/002224298905300204>
- Map & Data Library. (2017). *Cleaning Data in Stata*. <https://mdl.library.utoronto.ca/technology/tutorials/cleaning-data-stata>
- McCambridge, J., De Bruin, M., & Witton, J. (2012). The Effects of Demand Characteristics on Research Participant Behaviours in Non-Laboratory Settings: A Systematic Review. *PLoS ONE*, 7(6), e39116. <https://doi.org/10.1371/journal.pone.0039116>
- Milaninia, N., & Vermynck, E.-C. (2025). *Comparative Analysis of Transatlantic AI Regulatory Frameworks: Jurisdictional Convergence and Divergence in U.S.-EU Compliance Architectures*. https://www.kslaw.com/attachments/000/012/927/original/Comparative_Analysis_of_Transatlantic_AI_Regulatory_FrameworksJurisdictional_Convergence_and_Divergence_in_U.S.-EU_Compliance_Architectures.pdf?1753281742
- Mizik, N., & Jacobson, R. (2007). Myopic Marketing Management: Evidence of the Phenomenon and Its Long-Term Performance Consequences in the SEO Context. *Marketing Science*, 26(3), 361–379. <https://doi.org/10.1287/mksc.1060.0261>
- Montecchi, M., Plangger, K., West, D., & De Ruyter, K. (2024). Perceived brand transparency: A conceptualization and measurement scale. *Psychology & Marketing*, 41(10), 2274–2297. <https://doi.org/10.1002/mar.22048>

- Morhart, F., Malär, L., Guèvremont, A., Girardin, F., & Grohmann, B. (2015). Brand authenticity: An integrative framework and measurement scale. *Journal of Consumer Psychology*, 25(2), 200–218. <https://doi.org/10.1016/j.jcps.2014.11.006>
- Mouriquand, D. (2025a, May 11). *Coca-Cola's AI-generated Christmas ad sparks backlash (again)*. Euronews. <https://www.euronews.com/culture/2025/11/05/real-magic-coca-colas-ai-generated-christmas-ad-sparks-widespread-backlash-again>
- Mouriquand, D. (2025b, December 15). *Real magic: Why is this French Christmas ad going viral worldwide?* Euronews. <https://www.euronews.com/culture/2025/12/15/real-magic-why-is-french-supermarket-intermarches-christmas-advert-going-viral-worldwide>
- Napoli, J., Dickinson, S. J., Beverland, M. B., & Farrelly, F. (2014). Measuring consumer-based brand authenticity. *Journal of Business Research*, 67(6), 1090–1098. <https://doi.org/10.1016/j.jbusres.2013.06.001>
- Nyilasy, G., & Gangadharbatla, H. (2025). *AI-washing: The Asymmetric Effects of Its Two Types on Consumer Moral Judgments* (Version 1). arXiv. <https://doi.org/10.48550/ARXIV.2507.04352>
- Okada, E. M. (2005). Justification Effects on Consumer Choice of Hedonic and Utilitarian Goods. *Journal of Marketing Research*, 42(1), 43–53. <https://doi.org/10.1509/jmkr.42.1.43.56889>
- Oxford University Press. (n.d.). Artificial intelligence. In *Oxford English Dictionary*. Oxford University Press. <https://www.oed.com>
- Petrova Karina. (2026, April 11). *Artificial intelligence makes consumers more impatient*. <https://www.psypost.org/artificial-intelligence-makes-consumers-more-impatient/>
- Piasecki, S., Morosoli, S., Helberger, N., & Naudts, L. (2024). AI-generated journalism: Do the transparency provisions in the AI Act give news readers what they hope for? *Internet Policy Review*, 13(4). <https://doi.org/10.14763/2024.4.1810>
- Pleyers, G. (2025). *Neuromarketing and experimental approach [Diapositives]*. <https://uclouvain.be/en-cours-2025-llsms2005>
- Sansome, K., Conduit, J., & Wilkie, D. (2025). Consumers demand transparency... but do they actually engage? Exploring motives and interactions with brand transparency information. *Journal of Business Research*, 194, 115386. <https://doi.org/10.1016/j.jbusres.2025.115386>

- Schepman, A., & Rodway, P. (2023). The General Attitudes towards Artificial Intelligence Scale (GAAIS): Confirmatory Validation and Associations with Personality, Corporate Distrust, and General Trust. *International Journal of Human–Computer Interaction*, 39(13), 2724–2741. <https://doi.org/10.1080/10447318.2022.2085400>
- Schrift, R. Y., Kivetz, R., & Netzer, O. (2016). Complicating decisions: The work ethic heuristic and the construction of effortful decisions. *Journal of Experimental Psychology: General*, 145(7), 807–829. <https://doi.org/10.1037/xge0000171>
- Seymour, M., Riemer, K., & Kay, J. (2017). *Interactive Realistic Digital Avatars—Revisiting the Uncanny Valley*. Hawaii International Conference on System Sciences. <https://doi.org/10.24251/HICSS.2017.067>
- Shi, Y., & Jiang, Z. (2026). Consumer Responses to AI Disclosure Labels: The Role of Novelty and Authenticity. *Sage Open*, 16(1), 21582440261417793. <https://doi.org/10.1177/21582440261417793>
- Shimp, T. A., Hyatt, E. M., & Snyder, D. J. (1991). A Critical Appraisal of Demand Artifacts in Consumer Research. *Journal of Consumer Research*, 18(3), 273. <https://doi.org/10.1086/209259>
- Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355–374. <https://doi.org/10.2307/1882010>
- Sui, M. (2026). When the AI author is not disclosed: How cognitive dispositions affect audience perceptions of AI-generated news across topics. *Communication and Change*, 2(1), 3. <https://doi.org/10.1007/s44382-026-00023-6>
- Suresh, K. (2011). An overview of randomization techniques: An unbiased assessment of outcome in clinical research. *Journal of Human Reproductive Sciences*, 4(1), 8–11. <https://doi.org/10.4103/0974-1208.82352>
- Swaminathan, V., Gupta, S., Keller, K. L., & Lehmann, D. (2022). Brand actions and financial consequences: A review of key findings and directions for future research. *Journal of the Academy of Marketing Science*, 50(4), 639–664. <https://doi.org/10.1007/s11747-022-00866-7>
- Tavakol, M., & Dennick, R. (2011). Making sense of Cronbach’s alpha. *International Journal of Medical Education*, 2, 53–55. <https://doi.org/10.5116/ijme.4dfb.8dfd>

- Tzara, A. (2025). *Made with... AI? The Impact of AI Disclosure on Ad Attitudes* [University of Gothenburg]. <https://gupea.ub.gu.se/server/api/core/bitstreams/ff776d74-430f-4370-92e8-ac0cc9eb0ed6/content>
- Velásquez-Salamanca, D., Martín-Pascual, M. Á., & Andreu-Sánchez, C. (2025). Interpretation of AI-Generated vs. Human-Made Images. *Journal of Imaging*, 11(7), 227. <https://doi.org/10.3390/jimaging11070227>
- Voss, K. E., Spangenberg, E. R., & Grohmann, B. (2003). Measuring the Hedonic and Utilitarian Dimensions of Consumer Attitude. *Journal of Marketing Research*, 40(3), 310–320. <https://doi.org/10.1509/jmkr.40.3.310.19238>
- Wild, H., Kyröläinen, A.-J., & Kuperman, V. (2022). How representative are student convenience samples? A study of literacy and numeracy skills in 32 countries. *PLOS ONE*, 17(7), e0271191. <https://doi.org/10.1371/journal.pone.0271191>
- Wortel, C., Vanwesenbeeck, I., & Tomas, F. (2024). Made with Artificial Intelligence: The Effect of Artificial Intelligence Disclosures in Instagram Advertisements on Consumer Attitudes. *Emerging Media*, 2(3), 547–570. <https://doi.org/10.1177/27523543241292096>
- Yang, J., & Battocchio, A. F. (2021). Effects of transparent brand communication on perceived brand authenticity and consumer responses. *Journal of Product & Brand Management*, 30(8), 1176–1193. <https://doi.org/10.1108/JPBM-03-2020-2803>
- Zeng, R., Zhu, D., & Evans, R. (2026). When algorithms fail us: Perceived algorithmic ineffectiveness, psychological reactance, and implicit personality as drivers of algorithm aversion behavior on short-form video platforms. *Humanities and Social Sciences Communications*, 13(1), 266. <https://doi.org/10.1057/s41599-026-06573-w>
- Zhang, L., & Hur, C. (2025). The Impact of Generative AI Images on Consumer Attitudes in Advertising. *Administrative Sciences*, 15(10), 395. <https://doi.org/10.3390/admsci15100395>
- Zhang, Y., & Gosline, R. (2022). Understanding Algorithm Aversion: When Do People Abandon AI After Seeing It Err? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4299576>
- Zhou, S., Gordon, M. L., Krishna, R., Narcomey, A., Fei-Fei, L., & Bernstein, M. S. (2019). *HYPE: A Benchmark for Human eYe Perceptual Evaluation of Generative Models* (Version 4). arXiv. <https://doi.org/10.48550/ARXIV.1904.01121>

Abstract :

In the constant evolution of our modern world, few innovations have been as transformative and revolutionary as the emergence of artificial intelligence. In addition to changing how people live, work, and make decisions, artificial intelligence (AI) is now redefining the boundaries separating human and machine capabilities, an effect increasingly affecting marketing activities.

This study lies at the intersection of three dimensions: disclosure effects, product-category moderation, and deception dynamics. It employs an experimental design, and was conducted using a convenience sample that gathered 164 participants.

The results indicate that isolated AI-generated labels do not negatively affect consumer trust, emotions, perceived authenticity or purchase intentions. However, disclosed human-made visuals were found to reduce perceived authenticity by triggering consumer skepticism. In addition, direct comparisons reveal a preference for human-made advertisements, likely explained by perceived effort heuristics, while no significant interaction effects were observed across product categories. Finally, falsely claiming human authorship may have severe consequences: when consumers discover the true AI origin of content, trust and purchase intentions decline significantly.

Further research should continue exploring this rapidly evolving context, especially given its growing relevance for both companies and consumers. Firms must navigate the trade-off between the cost-efficiency of AI use and potential consumer backlash, while consumers themselves are becoming increasingly aware of this new era of hyper-realistic content.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management

Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique
www.uclouvain.be/lsm