VILNIUS UNIVERSITY
BUSINESS SCHOOL

DEEPTECH ENTREPRENEURSHIP

Modestas Gražys

MASTER THESIS

GENERATYVINIS DIRBTINIS INTELEKTAS - REVOLIUCINĖ JĖGA MARKETINGE: AUKŠTŲJŲ TECHNOLOGIJŲ VERSLO GALIMYBĖS IR IŠŠŪKIAI	**GENERATIVE AI AS A DISRUPTIVE FORCE IN MARKETING: OPPORTUNITIES AND CHALLENGES IN DEEPTECH ENTREPRENEURSHIP**

Supervisor ______________________

Doc. Dr. Eglė Terminė

Vilnius, 2025

SUMMARY

VILNIUS UNIVERSITY BUSINESS SCHOOL

DEEPTECH ENTREPRENEURSHIP

MODESTAS GRAŽYS

GENERATIVE AI AS A DISRUPTIVE FORCE IN MARKETING: OPPORTUNITIES AND CHALLENGES IN DEEPTECH ENTREPRENEURSHIP

Supervisor - Doc. Dr. Eglė Terminė

Master’s thesis was prepared in Vilnius, in 2025

Scope of Master’s thesis - 64 pages.
Number of tables used in the FMTP - 3 pcs.
Number of figures used in the FMTP - 8 pcs.
Number of bibliography and references - 57 pcs.

The FMTP described in brief:

Generative artificial intelligence systems have rapidly become integral to creative marketing infrastructure, yet their adoption is advancing more quickly than the understanding of their strategic and ethical implications. This work investigates the transformation of creative production process automation in marketing communication through generative AI, highlighting key implications for DeepTech entrepreneurs and marketing practitioners. The study utilizes a qualitative-dominant mixed-methods design, combining semi-structured expert interviews (n=3) with marketing practitioners and audience intercepts (n=30) in Stockholm, Sweden, which compared AI-generated and human-created advertising. Expert interviews were analyzed using thematic coding, while audience preferences were assessed through a pre-post disclosure design. The findings indicate that entrepreneurs primarily adopt generative AI in response to budget constraints and competitive pressures, implementing hybrid workflows that maintain human oversight during both ideation and approval stages. Experts report substantial efficiency gains and significant increases in variant generation capacity, while cautioning that unchecked automation may result in brand homogenization. Audience intercepts revealed 93.3% preference stability after disclosure, suggesting that execution quality is more influential than production method in low-involvement contexts. The research concludes that while AI offers scalable production capabilities, sustainable implementation requires pragmatic governance, context-sensitive disclosure, and continuous investment in human capability development to preserve brand distinctiveness.

Problem, objective and tasks of the FMTP:

The aim of this research is to analyse the impact of generative artificial intelligence on creativity automation within marketing communication, by exploring how advertising, social media, and brand content professionals integrate GenAI tools, and to identify strategic, ethical, and organizational implications of this transformation.

1. To analyse the theoretical foundations of generative artificial intelligence and its role in automating creativity within marketing communication.
2. To review and systematise existing research on the use of GenAI in advertising, social media marketing, and content branding, with a focus on opportunities and challenges.
3. To conduct expert-based empirical research to reveal current practices, perceived value, and challenges in the use of GenAI in creative marketing.
4. To assess the expert insights on how GenAI affects efficiency, emotional and creative quality, organizational structures, and ethical decision-making.
5. To conduct audience-based empirical research to examine consumer preferences and disclosure effects when viewing AI-generated versus human-created advertising.
6. To develop a conceptual model describing the relationship between generative AI adoption, creative performance, authenticity, and consumer trust.
7. To propose recommendations for entrepreneurs and marketing practitioners on how to balance human creativity and AI automation to ensure authenticity, effectiveness, and sustainable brand development.

Research methods used in the FMTP:

Qualitative-dominant mixed methods: semi-structured expert interviews (n=3) with thematic coding; audience intercept survey (n=30) comparing AI vs human Coca-Cola ads with pre/post disclosure; McNemar test for shifts; conceptual synthesis.

Research and results obtained:

Adoption is driven by cost and competitive pressure; workflows are hybrid with human oversight at start and end; efficiency gains coexist with homogenization risk; audience preferences stayed largely stable after disclosure (93.3% retained AI choice); governance is pragmatic and judgment-based.

Conclusions of the FMTP:

Generative AI is now a baseline capability, but teams must keep human quality gates, apply context-aware disclosure, build prompting/refinement skills, and formalize governance as volume scales.

Information about the publication of FMTP results or adaptation for publication

Results not published.

SANTRAUKA

VILNIAUS UNIVERSITETO VERSLO MOKYKLA

AUKŠTŲJŲ TECHNOLOGIJŲ VERSLAS

MODESTAS GRAŽYS

GENERATYVINIS DIRBTINIS INTELEKTAS - REVOLIUCINĖ JĖGA MARKETINGE: AUKŠTŲJŲ TECHNOLOGIJŲ VERSLO GALIMYBĖS IR IŠŠŪKIAI

Darbo vadovė - Doc. Dr. Eglė Terminė

Magistro baigiamasis darbas parengtas Vilnius, 2025 m.

Magistro baigiamojo darbo apimtis - 64 psl.
Lentelių skaičius magistro baigiamajame darbe - 3 vnt.
Paveikslėlių skaičius magistro baigiamajame darbe - 8 vnt.
Literatūros ir šaltinių sąrašo įrašų skaičius - 57 vnt.

Magistro baigiamasis darbas trumpai:

Generatyvinės dirbtinio intelekto sistemos greitai tapo neatsiejama kūrybinio marketingo infrastruktūros dalimi, tačiau jų diegimas vyksta greičiau nei jų strateginių ir etinių pasekmių supratimas. Šiame darbe nagrinėjama kūrybinės produkcijos proceso automatizavimo transformacija rinkodaros komunikacijoje naudojant generatyvųjį AI, išryškinant pagrindines pasekmes aukštųjų technologijų verslininkams ir marketingo praktikams. Tyrime naudojamas kokybinis mišrių metodų dizainas, derinant pusiau struktūrizuotus ekspertų interviu (n = 3) su rinkodaros specialistais ir atsitiktinių žmonių apklausas (n = 30) Stokholme, Švedijoje, kuriose buvo lyginama DI generuojama ir žmogaus sukurta reklama. Ekspertų interviu buvo analizuojami naudojant teminį kodavimą, o atsitiktinių žmonių pageidavimai buvo tiriami naudojant prieš ir po informacijos atskleidimo metodą. Išvados rodo, kad verslininkai generatyvines dirbtinio intelekto sistemas taiko daugiausia esant biudžeto apribojimams ir konkurenciniam spaudimui, įdiegdami hibridinę darbo eigą, kurioje žmogaus priežiūra išlieka idėjos generavimo ir patvirtinimo etapuose. Nors ekspertai praneša apie reikšmingą efektyvumo padidėjimą ir daugkartinį variacijų generavimo pajėgumų padidėjimą, jie įspėja, kad nekontroliuojama automatizacija kelia prekės ženklo homogenizacijos riziką. Atsitiktinių žmonių apklausos parodė 93,3 % nuomonių stabilumą po informacijos atskleidimo, o tai rodo, kad reklamos įvykdymo kokybė yra svarbesnė nei gamybos metodas. Tyrimo išvadose teigiama, kad nors dirbtinis intelektas siūlo keičiamo mastelio gamybos galimybes, tvariam įgyvendinimui reikalingas pragmatiškas valdymas, nuo konteksto priklausantis DI turinio žymėjimas ir dėmesys kūrybininkų gebėjimų stiprinimui, siekiant išsaugoti išskirtinumą.

Magistro baigiamojo darbo problema, tikslas ir uždaviniai:

Tyrimo tikslas - ištirti generatyvaus DI poveikį kūrybinei automatizacijai rinkodaros komunikacijoje, tiriant, kaip reklamos, socialinių tinklų ir turinio ženklodaros profesionalai integruoja GenAI įrankius, ir identifikuoti šios transfor-

macijos strategines, etines ir organizacines pasekmes.

1. Išanalizuoti generatyvaus DI teorinius pagrindus ir jo vaidmenį automatizuojant kūrybą rinkodaros komunikacijoje.
2. Peržvelgti ir susisteminti esamus tyrimus apie generatyvusis DI naudojimą reklamoje, socialinių tinklų rinkodaroje ir turinio ženklodaroje, sutelkiant dėmesį į galimybes ir iššūkius.
3. Atlikti ekspertinį empirinį tyrimą, kad būtų atskleista dabartinė praktika, suvokiama vertė ir iššūkiai generatyvusis DI naudojime kūrybiniame marketinge.
4. Įvertinti ekspertų įžvalgas apie tai, kaip generatyvusis DI veikia efektyvumą, emocinę ir kūrybinę kokybę, organizacines struktūras ir etinius sprendimus.
5. Atlikti empirinį tyrimą, siekiant ištirti vartotojų nuomones, kai lyginamos DI ir žmogaus sukurtos reklamos.
6. Sukurti konceptinį modelį, aprašantį ryšį tarp generatyvaus DI diegimo, kūrybinio rezultatyvumo, autentiškumo ir vartotojų pasitikėjimo.
7. Pateikti rekomendacijas verslininkams ir marketingo specialistams, kaip balansuoti tarp žmogaus kūrybiškumo ir DI automatizacijos, užtikrinant autentiškumą, efektyvumą ir tvarų prekės ženklo vystymą.

Magistro baigiamajame darbe naudoti tyrimo metodai:

Kokybinis-mišrus dizainas: pusiau struktūruoti ekspertų interviu (n=3) su teminiu kodavimu; gatvės apklausa (n=30) lyginant DI ir žmogaus kurtas Coca-Cola reklamas; McNemar testas ir konceptinė sintezė.

Tyrimas ir gauti rezultatai:

Diegimą lemia kaštų ir konkurencijos spaudimas; darbo eiga hibridinė su žmogaus kontrole; gaunami efektyvumo šuoliai, bet homogenizacijos rizika; atskleidus, kad vieną iš reklamų generavo DI 93,3 % dalyvių nepakeitė nuomonės; valdymas išlieka pragmatiškas.

Magistro baigiamojo darbo išvados:

DI tampa baziniu įrankiu, bet būtina žmogaus priežiūra ir kokybės įvertinimas; atskleisti DI naudojimą reikia priklausomai nuo aplinkybių konteksto; investuoti į įgūdžius ir formalizuoti valdymą augant mastui.

Informacija apie rezultatų publikavimą ar parengimą publikacijai:

Rezultatai nepublikuoti.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

1. **AI:** Artificial Intelligence
2. **B2B:** Business-to-Business
3. **B2C:** Business-to-Consumer
4. **BS:** Business School
5. **CG:** Computer Graphics
6. **CGI:** Computer-Generated Imagery
7. **CSV:** Comma-Separated Values
8. **DI:** Dirbtinis intelektas (Artificial Intelligence)
9. **DSA:** Digital Services Act
10. **EU:** European Union
11. **GDPR:** General Data Protection Regulation
12. **GenAI:** Generative Artificial Intelligence
13. **IEC:** International Electrotechnical Commission
14. **ID:** Identifier
15. **IP:** Intellectual Property
16. **ISO:** International Organization for Standardization
17. **LMS:** Learning Management System
18. **LLM:** Large Language Model
19. **QA:** Quality Assurance
20. **RAG:** Retrieval-Augmented Generation
21. **ROI:** Return on Investment
22. **SEO:** Search Engine Optimization

23. **SOC2:** Service Organization Control 2
24. **SSO:** Single Sign-On
25. **UI:** User Interface
26. **UK:** United Kingdom
27. **USA:** United States of America
28. **UX:** User Experience

INTRODUCTION

Background and Problem Statement

Generative artificial intelligence systems have rapidly evolved from experimental tools to configurable infrastructure embedded within marketing operations (Ford et al., 2023; Serra-Simón et al., 2025). Large language models that draft campaign briefs and refine copy, diffusion engines that materialise visual concepts from text prompts (Ramesh et al., 2021; Rombach et al., 2022), video generators that animate brand narratives, virtual avatar platforms that scale spokesperson content across languages, and integrated creative suites embedding AI across design workflows now form a reconfigured creative stack accessible to agencies, in-house teams, and solo practitioners alike (Canva, 2023; Midjourney, 2023; OpenAI, 2022; Runway, 2025; Synthesia, 2025a). Industry fieldwork documents that agencies report efficiency gains of 50–80% on production tasks and time savings of 40–50% when integrating these tools into routine content development, compressing cycles that previously required multi-person coordination (Serra-Simón et al., 2025). Practitioners simultaneously describe adoption as a survival response to flat client budgets and competitive pressure to match rivals deploying similar capabilities.

Yet adoption speed outpaces understanding of consequences. Marketing teams face an unresolved tension: generative systems accelerate variant generation and enable hyper-personalisation at scale, but agencies observe that unchecked automation weakens brand judgement, produces the visible "super-AI look" audiences reject, and threatens brand distinctiveness when competitors use the same AI tools to create similar content (Osadchaya et al., 2024; Serra-Simón et al., 2025). This gap limits evidence-based guidance for teams navigating the speed-versus-meaning trade-offs inherent in AI-assisted creative work. Understanding generative AI's impact on brand storytelling, productivity, and market perception is essential for DeepTech entrepreneurs, marketers, researchers, and consumers navigating this technological shift. This research is particularly timely as brands increasingly integrate AI tools into their creative workflows, making it critical to evaluate both opportunities and challenges this shift presents. Addressing this issue will provide strategic insights that help DeepTech ventures optimize AI adoption strategies while maintaining brand integrity and creative differentiation in an era of automated creativity.

Research Aim and Objectives

The aim of this research is to analyse the impact of generative artificial intelligence on creativity automation within marketing communication, by exploring how advertising, social media, and brand content professionals integrate GenAI tools, and to identify strategic, ethical, and organizational implications of this transformation.

To achieve this aim, the following objectives are set:

1. To analyse the theoretical foundations of generative artificial intelligence and its role in automating creativity within marketing communication.

2. To review and systematise existing research on the use of GenAI in advertising, social media marketing, and content branding, with a focus on opportunities and challenges.
3. To conduct expert-based empirical research to reveal current practices, perceived value, and challenges in the use of GenAI in creative marketing.
4. To assess the expert insights on how GenAI affects efficiency, emotional and creative quality, organizational structures, and ethical decision-making.
5. To conduct audience-based empirical research to examine consumer preferences and disclosure effects when viewing AI-generated versus human-created advertising.
6. To develop a conceptual model describing the relationship between generative AI adoption, creative performance, authenticity, and consumer trust.
7. To propose recommendations for entrepreneurs and marketing practitioners on how to balance human creativity and AI automation to ensure authenticity, effectiveness, and sustainable brand development.

Methods Overview

This exploratory study employs a qualitative-dominant mixed-methods design combining semi-structured expert interviews (n=3) with an audience intercept survey (n=30) to contextualize and extend the theoretical synthesis. Expert interviews were conducted via voice messages (November 25–29, 2025) with marketing practitioners who actively use generative AI in creative workflows, exploring adoption drivers, implementation pathways, observed outcomes, and organizational challenges. Interviews followed a 12-question protocol organized into five thematic sections and were analyzed through thematic coding with AI assistance, using structural coding based on interview sections alongside open coding to identify patterns.

The audience component employed a pre-post disclosure design in Stockholm public spaces (November 29–30, 2025), where participants viewed AI-generated and human-created Coca-Cola commercials, indicated preference before and after disclosure, and explained their reasoning. Analysis triangulated expert predictions against observed audience reactions using McNemar's test for paired preference data and categorical grouping for open explanations, yielding context-rich insights into how practitioners deploy generative AI and how audiences respond to disclosed AI-generated content.

Novelty, Benefits, and Scope

This research makes three contributions. Theoretically, it develops an integrated model linking adoption drivers, implementation mechanisms, organizational outcomes, and contextual moderators, synthesizing perspectives from Behavioral Reasoning Theory, paradox research, and field studies to map how generative AI reshapes creative marketing. This framework moves beyond simple efficiency narratives to surface inherent tensions between automation and authenticity, speed and distinctiveness, that practitioners must navigate rather than resolve.

Practically, findings provide DeepTech entrepreneurs and marketing practitioners with evidence-based guidance on where AI augments creative work versus where human oversight remains essential, informing decisions on tool selection, workflow integration, and disclosure strategies. DeepTech ventures and established organizations alike gain a structured approach to evaluating when automation strengthens a competitive position, and when it risks the brand homogenization.

The study focuses deliberately on advertising, social media, and content branding domains where generative AI adoption is most visible, examining creative workflows rather than technical AI capabilities. It excludes other marketing functions, such as analytics, customer relationship management, or pricing, that engage AI differently. Geographic constraints center audience data in Stockholm, Sweden, and temporal boundaries reflect contemporary tools available between 2022 and 2025. Small purposive samples prioritize depth over breadth, yielding context-rich insights whose transferability readers must judge against their own settings.

Thesis Structure and AI Usage Disclosure

This thesis proceeds through several sections: the theoretical part synthesizes literature on generative AI capabilities, automation debates, and domain-specific applications in advertising, social media, and content branding, concluding with an integrated conceptual model. The Methodology section details the qualitative-dominant mixed-methods design, sampling strategies, data collection instruments, and analytical procedures. Research Findings present expert interview themes, audience preference patterns, and triangulated insights linking practitioner predictions to consumer responses. Conclusions and Recommendations distill key findings and propose strategic guidance for DeepTech entrepreneurs and marketing practitioners.

Artificial intelligence tools, including Claude, ChatGPT, Gemini, Grammarly Pro, and TurboScribe, supported the literature review, expert interview transcription, thematic coding, and drafting phases. LLM-based tools were utilized to generate initial drafts of subsections and paragraphs, which were subsequently thoroughly reviewed and rewritten by the author to ensure contextual alignment and accuracy. Following this authorial review, LLMs and Grammarly Pro were employed to refine grammar and ensure academic writing style. The drafting process used a methodology inspired by Retrieval-Augmented Generation (RAG), in which the LLM was first given all necessary context and instructed to generate content solely based on the provided information. Additionally, LLMs assisted in structural planning and served as a consultative resource during data analysis, similarly adhering to this RAG-inspired methodology to ensure insights were grounded in the provided context. TurboScribe was used to transcribe the English-language expert interview, after which the transcript was manually verified and corrected. All research decisions, analyses, and conclusions represent the author’s original work. All empirical data collection adhered to the informed consent and anonymization protocols detailed in the Methodology section.

1 A THEORETICAL PART

1.1 Defining Generative AI in Marketing

1.1.1 Large Language Models for Campaign Ideation

Large language models function as conversational assistants, maintaining context across exchanges, answering follow-up questions, adjusting responses when users identify issues, refusing unsafe prompts, and, in multimodal variants, processing and generating images, audio, or code within a single dialogue while adhering to provider-defined safety protocols.

These documented behaviors address common user needs. Launch examples illustrate assistants explaining topics, self-correcting when challenged, guiding users through processes, and integrating text and visual content seamlessly. Collectively, these capabilities demonstrate that modern large language models can answer research questions, outline campaign plans, and draft or revise content linked to supporting assets within a single conversation.

In practice, these models are accessed via chat-style interfaces that answer questions, guide tasks, and revise drafts while supporting multimodal inputs subject to provider safety reviews. ChatGPT and Gemini exemplify this approach, offering comparable capabilities across text and other media formats (Google, 2023; OpenAI, 2022).

1.1.2 Diffusion and Transformer Image Generators

Diffusion and transformer image generators convert written prompts or reference images into visuals, offering controls for style, composition, and edits directly through conversational interfaces (Ramesh et al., 2021; Rombach et al., 2022). Public documentation highlights their ability to render detailed scenes, support guided styling, insert legible text, and iteratively adjust outputs (DeepMind, 2025; Fortin et al., 2025; B. F. Labs et al., 2025; Midjourney, 2023).

These controls enable users to prompt alternative visual styles, adjust lighting or layout details through iterative re-prompts, edit text overlays, and regenerate specific elements within the same interface. Marketing teams can leverage these functionalities to develop concept boards, localize campaign imagery, and rapidly test creative directions prior to allocating production budgets.

Midjourney, FLUX, Imagen 4, and Gemini 2.5 Flash Image exemplify this category. Midjourney emphasizes stylized rendering via prompt weighting, FLUX offers high prompt adherence and open-weight deployments for controlled workflows, Imagen 4 prioritizes photorealistic results with precise typography, and Gemini 2.5 Flash Image supports conversational edits and in-context revisions. These represent variations on a shared workflow that marketers can operationalize.

1.1.3 Generative Video Platforms

These platforms transform text prompts or reference frames into moving imagery, providing controls for camera movement, shot composition, and scene consistency while allowing targeted edits to regenerate segments that need refinement. (OpenAI, 2024; Runway, 2025; Wang et al., 2024)

Demonstrated behaviours include drafting entire shots from descriptions, regenerating segments to smooth transitions, and iterating on light, motion, or subject placement within the same workflow. Marketing teams can deploy these capabilities to craft short-form ads or reels, spin up quick-turn variations of campaign footage, and simulate experiential concepts before approving large production budgets.

Runway, Seedance, and Sora exemplify these traits: each supports text-to-video generation with controllable cameras, composition, and iterative regeneration, spanning outputs from short-form vertical reels to longer, cinematic sequences. Other prominent platforms in this category include Kling, Veo, Dream Machine, Hailuo, and Higgsfield (AI, 2024; Google, 2024; L. Labs, 2024; MiniMax, 2025; Technology, 2024).

1.1.4 Virtual Avatar Platforms

Virtual avatar platforms turn typed scripts and uploaded materials into presenter-led videos, combining AI avatars, synthetic voices, brand templates, and publishing controls so teams can scale spokesperson content without cameras or post-production. (Synthesia, 2025a)

Key capabilities include access to extensive avatar and voice catalogs, creation of custom or personalized avatars that maintain brand consistency, translation of content into over 140 languages with lip-synchronized speech, alignment of layouts with brand kits and media libraries, embedding of interactive elements and LMS-compatible exports, and coordination of reviews, analytics, and compliance within a unified workspace. Marketing and communications teams can apply these tools for sales enablement updates, product explainers, internal training, and executive announcements requiring rapid localization.

Synthesia exemplifies these traits with its end-to-end AI video workflow, collaborative editing and version control, interactive modules, multilingual delivery, and enterprise safeguards (SOC2, GDPR, ISO/IEC 42001, SSO) that keep scaled spokesperson content governed (Synthesia, 2025b).

1.1.5 Audio and Sonic Branding Models

Audio generation platforms transform written prompts into branded voices, music, and soundscapes, offering control over tone, pacing, and arrangement while providing safety tooling for responsible deployment.

These tools enable users to design custom voices, adjust emotion and pronunciation, or compose complete songs by specifying genre, mood, and lyrical content within a single workflow. Marketers can leverage these capabilities to maintain consistent voiceovers across languages, provide accessible narration, develop sonic logos, and generate campaign-specific tracks without the need for traditional studio resources.

ElevenLabs and Suno exemplify this model family: ElevenLabs v3/Voice Design v3 combine expressive text-to-speech, cloning controls, multilingual output, and similarity safeguards, while Suno v5 turns prompts into complete vocal tracks with commercial licensing, making rapid sonic experimentation practical for marketing teams.(ElevenLabs, 2025a, 2025b; Suno, 2025)

1.1.6 Integrated Creative Suites

Integrated creative suites bundle generative copy, imagery, video, and design layouts with brand governance, collaboration, and publishing controls so teams can manage end-to-end visual production inside one workspace.(Adobe, 2023; Canva, 2023)

Across these platforms, users can begin with templates, apply brand kits, reformat assets for various channels, translate content, review edits collaboratively, and track usage, thereby maintaining workflow continuity without switching between applications. Marketing teams can utilize these toolkits to launch multi-format campaigns, update evergreen assets for new markets, and coordinate brand moodboards within a single suite.

Canva's Magic Studio and Adobe Firefly exemplify the integrated creative suites: both combine prompt-driven generation with asset libraries, brand controls, localisation presets, and admin tooling that covers permissions, audit trails, and safe commercial use—traits that make integrated suites a central command point for creative operations.

1.1.7 Implications for Creative Operations

Understanding these model families demonstrates how generative AI extends beyond individual tools to reshape marketing taxonomies. Large language models translate insights into briefs, vision models generate concepts, video engines animate narratives, audio systems develop sound identities, and integrated suites unify workflows into a manageable system. Collectively, these components structure the creative process into modular, interoperable services that marketing organizations can combine, govern, and scale across campaigns, channels, and regions.

1.2 Automation Versus Augmentation Debates

The automation versus augmentation debate frames generative AI as either automation that displaces human labor or augmentation that enhances creative capabilities. The displacement framing captures practitioner anxieties as AI systems execute tasks faster and on a larger scale, including fears that in-house brand automation will bypass agency roles (Osadchaya et al., 2024). Empirical studies report routine task substitution alongside the need for higher-level cognitive work, signaling role restructuring rather than uniform replacement (Amankwah-Amoah et al., 2024). Augmentation research, conversely, positions AI as a collaborative partner that accelerates content creation by analysing campaign data, market trends, and consumer reactions to propose resonant executions (Sands et al., 2024), while requiring practitioners to build competencies that leverage AI within existing creative processes (Serra-Simón et al., 2025).

The practical implications diverge. The automation lens targets standardization and efficiency, but agency interviews caution that overuse can dilute brand judgment, prompting safeguards that keep human review over client fit and originality. The augmentation lens preserves human creative agency yet demands reskilling and hybrid workflows that integrate AI-generated options with editorial oversight. For marketing organisations, these framings are not mutually exclusive. Agency evidence shows production-oriented tasks face greater automation exposure, while creative direction and client stewardship remain human-led (Serra-Simón et al., 2025), and audience studies find human–AI collaborative ads outperform fully synthetic versions on perceived creativity and brand morality, underscoring the value of mixed approaches (Madathil, 2025). Effective integration, therefore, depends on task-level differentiation and governance that match automation to repeatable work and reserve human oversight for narrative coherence and brand trust (Amankwah-Amoah et al., 2024).

1.3 Framing Creative Automation

Generative AI reconceptualizes creative automation as a dynamic interplay between machine-driven scalability and human judgment, moving beyond a simple substitution model. Interviews with agency leaders in Rio de Janeiro and Catalonia indicate that automation facilitates research synthesis, benchmarking, and initial ideation, while strategists maintain responsibility for evaluating outputs in relation to client context, regulatory requirements, and brand standards. In operational terms, automation reduces delays in information retrieval and enables rapid generation of variants; however, human augmentation remains essential for determining which machine-generated options advance to the storyboard stage (Osadchaya et al., 2024; Serra-Simón et al., 2025).

In production workflows, automation is implemented as modular components that teams reuse to ensure quality and scalability. Recent benchmarking demonstrates that large text-to-image models surpass human freelancers in perceived realism, aesthetic control, and prompt fidelity, as evidenced by 10,320 AI-generated variants derived from 2,400 human-authored briefs. This indicates that generative engines can achieve ‘superhuman’ quality when provided with precise instructions. Marketers leverage this capacity to produce cost-effective variants for testing and selection, while production leads assume a curator role, filtering machine outputs according to channel objectives and campaign structure before assets advance (Hartmann et al., 2025; Serra-Simón et al., 2025).

Creativity theory clarifies why this redistribution still hinges on human intervention. Evaluations of AI-authored products show that audiences penalise “useful but not novel” artefacts once they learn the creator was an AI system rather than a person, with the bias strongest among people who feel threatened by generative tools. The canonical novelty–usefulness framework therefore remains active: automation can flood marketers with permutations that maximise functional fitness, but maintaining distinctiveness demands human editors who can read cultural codes, recalibrate prompts, or remix assets so the final work earns originality credit. Teams that institutionalise human “gatekeepers of humanity” are effectively building safeguard rails around the creativity paradox—letting AI expand the feasible set while deploying humans to decide when mechanised solutions undermine brand meaning or audience esteem (Hattori et al., 2024; Madathil, 2025).

Research on brand storytelling highlights that the effectiveness of creative automation is contingent upon the distribution of narrative authority. Controlled experiments show that consumers who are informed that social posts are AI-authored rate brand authenticity, perceived image, and self-brand congruity lower than those exposed to human-authored content. However, the association between self-brand congruity and recommendation intent is stronger in the AI condition. These findings indicate that marketers should curate automated content with a focus on narrative belonging: while automation can expand scene-setting and persona diversity, human storytellers are necessary to craft emotionally coherent narratives that reinforce the brand’s stewardship of the experience. Managerial recommendations emphasize transparent disclosure and the integration of the brand’s established tone, values, and signature elements around AI-generated material to address authenticity concerns while preserving the scalability advantages of generative systems (Ali et al., 2025).

Collectively, creative automation now reallocates ideation, production, and narrative control across marketing functions in a multi-layered manner. Strategists employ automation for horizon scanning but retain authority over which opportunities should define the brand. Production leads manage generative services while instituting human review processes to safeguard novelty. Storytellers incorporate automated variations into campaigns but reassert authorship by contextualizing assets and explicitly communicating the brand’s voice, values, and signature experiences. The practical implication is that organizations should avoid framing automation and augmentation as mutually exclusive. Instead, they should develop operating models that position generative systems as collaborative infrastructure—enabling breadth, speed, and personalization—while empowering marketers to curate, test, and emotionally anchor outputs. This redistribution of creative labor presents a governance challenge: organizations must determine when automation enhances value creation and when human intervention is necessary to preserve meaning, trust, and strategic differentiation (Ali et al., 2025; Hartmann et al., 2025; Osadchaya et al., 2024; Serra-Simón et al., 2025).

1.4 Generative AI in Advertising

1.4.1 Application Landscape for Campaign Development

Generative AI in advertising is now applied across the core stages of campaign development. Recent reviews show that advertisers use AI for ad generation, copywriting, image production, media auctions, delivery, and optimisation, treating these tools as part of the standard planning toolkit (Ford et al., 2023). Agency surveys further report that more than one fifth of clients already view generative tools as a very positive force in advertising, while others remain cautious about pace and governance (Serra-Simón et al., 2025). Within ideation and copy development, practitioners report that large language models support early research synthesis, brainstorming, and tagline drafting, while major platforms already offer commercial services that create multiple text versions from a single prompt and adjust messaging to match landing page content (Duivenvoorde, 2024; Osadchaya et al., 2024). Interviews with digital marketing leaders add that the same systems are producing automated video posts, voiceovers, scripts, and SEO assets, with teams emphasising the need for human review to maintain quality and brand alignment (Joshi et al., 2025).

Visual production is another area of rapid deployment. Benchmark comparisons covering thousands of AI-generated images derived from human briefs show that text-to-image models can meet detailed specifications with high realism and control, so teams now rely on them for large volumes of creative variants (Hartmann et al., 2025). Interviews with agencies report cost savings where generative tools replace elements of traditional shoots, reducing the need for models, actors, and bespoke props, while emphasising that outputs remain under creative supervision before they are shared with clients (Duivenvoorde, 2024; Serra-Simón et al., 2025). In video advertising, mixed human–AI executions outperform fully synthetic characters on perceived creativity, attitudes, and brand morality while reducing audience anxiety, indicating that collaboration models remain important when teams assemble generative video spots (Madathil, 2025).

Personalisation and optimisation functions are increasingly automated. Computational advertising research documents the use of predictive modelling, programmatic creative, and recommender algorithms to tailor formats for specific behaviours and channels, supporting location-based targeting, mobile advertising, and chatbot-led interactions in the same workflow (Ford et al., 2023). Practitioner studies describe marketers using generative analytics to adjust tone and offers in response to observed reactions, framing the technology as a source of rapid customer insight that can feed back into copy and asset selection (Joshi et al., 2025). Experimental work also indicates that message configuration needs to respect the production source: rational appeals are processed more fluently in AI-attributed campaigns, suggesting that teams should match claim framing to the declared creator when they test variants (Song et al., 2024).

Platform providers are accelerating adoption by embedding generative services directly into advertising environments. Meta's creative suite now auto-generates background scenes and copy variants, Google Ads offers image and headline suggestions calibrated to the advertiser's landing page, and brands reported to the World Federation of Advertisers that these tools help reduce low-value production tasks so teams can focus on higher-order strategy (Duivenvoorde, 2024; Serra-Simón et al., 2025). In parallel, agency interviews describe planners leaning on GenAI to scan trend data, assemble inputs for account planning, and produce first-draft mood boards that specialists then refine during formal creative development (Joshi et al., 2025; Serra-Simón et al., 2025). Practitioners caution that these capabilities act as accelerants rather than replacements: they still route deliverables through human editors who validate brand tone, legal compliance, and cultural fit before anything reaches external audiences (Joshi et al., 2025; Serra-Simón et al., 2025).

Together these applications show how campaign teams are treating generative systems as configurable infrastructure across ideation, asset production, and optimisation. Creative, media, and analytics specialists draw on the same toolset for scale while maintaining manual checkpoints, ensuring that workflow speed gains do not displace brand fit or regulatory readiness (Ford et al., 2023; Serra-Simón et al., 2025).

1.4.2 Creative Effectiveness and Efficiency

Empirical studies document measurable productivity gains from generative AI deployment in advertising workflows. Practitioner surveys reported by the World Federation of Advertisers cite up to 50–80% efficiency gains and 40–50% time savings on low-value production tasks when agencies integrate text-to-text and image-generation tools into routine content development (Serra-Simón et al., 2025). These gains concentrate in asset iteration, brief refinement, and variant production, where automated generation reduces manual cycles and frees creative personnel for strategic work

(Osadchaya et al., 2024; Serra-Simón et al., 2025). Interviews with agency executives in Rio de Janeiro and Catalonia confirm that GenAI accelerates brainstorming, prototype creation, and report assembly, compressing timelines for deliverables that previously required multi-person coordination (Serra-Simón et al., 2025).

Testing workflows benefit from automation in scale and speed, though quality assurance remains human-gated. Experimental and field A/B tests of AI-generated marketing imagery demonstrate that automated variants can be evaluated rapidly against human-made assets (Hartmann et al., 2025). The efficiency paradox surfaces when outputs require fact-checking, tone adjustment, or brand-fit validation; practitioners note that supervision overhead can offset initial speed gains when hallucinations or generic phrasing demand rework (Osadchaya et al., 2024; Serra-Simón et al., 2025). Advertising professionals interviewed across both regions emphasize that GenAI-assisted testing still routes through human editors who verify legal compliance, cultural nuance, and alignment with brand guidelines before deployment (Serra-Simón et al., 2025).

Creative effectiveness metrics show mixed results when comparing AI-generated assets to human-authored content. Hartmann's multi-study evidence finds AI-generated visuals achieving comparable quality and realism ratings to professional photography in controlled evaluations and, in some A/B field tests, outperforming a stock image on click-through rates (Hartmann et al., 2025). Yet the same production pipelines mirror training data patterns rather than delivering conceptual breakthroughs, prompting concerns that reliance on generative systems can homogenize campaign outputs (Osadchaya et al., 2024). Practitioners adopt GenAI as an "antimodel" to test the originality of human ideas—if a copywriter's tagline matches ChatGPT output, agencies discard it to ensure differentiation (Osadchaya et al., 2024). This benchmarking practice reflects awareness that efficiency alone can erode brand distinctiveness without deliberate differentiation safeguards.

Cost implications remain under negotiation within the industry. Agency leaders acknowledge that GenAI reduces billable hours for tasks like mock-up creation, report writing, and email drafting, but fieldwork in 2024 found most sampled agencies had not adjusted fee structures while human oversight and strategic judgment remained indispensable (Serra-Simón et al., 2025). Legal and consumer-protection analyses also highlight that automated copy generation can lower labor requirements for writers and on-set talent, though downstream review costs persist (Duivenvoorde, 2024). The tension between operational efficiency and creative effectiveness underscores that productivity gains do not automatically translate to improved campaign outcomes unless agencies maintain manual checkpoints for quality, originality, and brand alignment (Osadchaya et al., 2024; Serra-Simón et al., 2025).

1.4.3 Audience Reception and Trust in AI-Produced Ads

Audience reception of AI-generated ads hinges on perceived creativity, trust, and humanness. Experiments in luxury advertising show that when viewers judge AI outputs as genuinely creative, they grant higher trust and humanness ratings, which lift informativeness, entertainment, credibility, and purchase intent relative to baselines lacking those qualities. Yet the same luxury studies find that authenticity remains fragile: trust depends on clear signals of human oversight and quality, underscoring the need for transparent disclosure and brand-fit safeguards. Studies of perceived eeriness also note that human-like cues increase trust only when they feel natural; if they seem off or "uncanny," they can erode trust instead (Jung et al., 2025). These dynamics map to concerns about sameness in advertising: overly standardized AI outputs can push campaigns to look alike, echoing the institutional isomorphism risk noted in industry interviews (Osadchaya et al., 2024).

Consumer responses vary sharply by service type and involvement. In hospitality, real images are preferred for hedonic or high-involvement choices, while AI- generated visuals are tolerated in utilitarian or low-involvement decisions, reflecting authenticity concerns and risk sensitivity. Qualitative follow-ups attribute rejection of AI visuals in hedonic settings to doubts about accuracy and the perceived absence of human craft, whereas low- involvement or functional tasks reduce scrutiny and make AI imagery acceptable. Boundary cases such as privacy-sensitive appeals show conditional acceptance: donors tolerate synthetic child imagery to avoid exploitation, but still demand clarity about the manipulation (Belanche et al., 2025). These findings are specific to hospitality contexts and show that disclosure and context alignment shape trust more than technical image quality alone.

Performance tests of AI-generated visuals indicate partial parity with human work but reveal dependence on format and task. Controlled studies report comparable quality and realism ratings between AI and professional photography, and field A/B campaigns in business-education messaging found a DALL-E 3 banner outperforming a curated stock image on click-through rates within the same budget. However, these results are confined to visual assets and a traffic objective; they do not eliminate concerns about conceptual sameness or authenticity raised in high-involvement categories. Brands applying these gains still route outputs through human review to maintain tone and ensure claims match the underlying offer Belanche et al. (2025).

Trust gaps persist even when emotional reactions converge. A study comparing AI and human-authored ad text found no significant differences in pleasure, arousal, or dominance, but participants expressed lower trust in AI content absent clear transparency and ethical cues (Sallaku et al., 2025). Interviews in advertising and consumer-protection analyses warn that AI-generated copy can over-promise or misrepresent features, reinforcing calls for human gate-keeping and factual verification before publication (Duivenvoorde, 2024). Emerging regulation echoes these concerns: the EU AI Act and platform deepfake policies require that synthetic or manipulated ads be flagged to avoid deception and preserve brand trust, and flagging AI use can itself influence how audiences judge the content (Duivenvoorde, 2024). Together these findings position disclosure, human oversight, and context-sensitive use as prerequisites for audience trust in AI- produced advertising.

1.4.4 Section Summary

Advertising teams use GenAI as configurable infrastructure for ideation, asset creation, testing, and optimisation, gaining speed and scale while keeping human checkpoints for tone, legality, and brand fit. Effectiveness results are mixed: AI visuals can perform well in small, controlled split tests, yet novelty and differentiation still rely on human judgment to avoid sameness. Audience trust depends on perceived creativity, humanness, and clear disclosure, with higher barriers in hedonic or high-involvement settings and under new labeling rules. Open questions remain about how quality assurance costs net against savings, whether fee models will change, and how long-term brand distinctiveness holds up under scaled automation.

1.5 Generative AI in Social Media

1.5.1 Content Creation and Engagement Automation Across Social Platforms

Generative AI deployment in social media spans content production, scheduling automation, and community interaction, with marketing teams drawing on the same model families to scale output across platforms. Practitioner interviews document that digital marketing specialists use AI tools for automated video posts, voiceovers, script preparation, SEO asset generation, and conditional posting workflows, where scheduling systems publish content based on predefined triggers or audience activity patterns. These automation layers free teams from repetitive execution tasks, potentially allowing strategists to concentrate on campaign design and performance interpretation while generative systems handle format adaptation and timed distribution.

Visual and video content generation tools are increasingly integrated into social media production pipelines. The same text-to-image and video platforms covered in model-family taxonomies—Midjourney, DALL-E, Adobe Firefly, Canva, and generative video engines—now supply assets tailored for platform specifications, from square Instagram posts to vertical Stories and Reels. Marketing experts report that generative AI supports rapid visualisation of campaign concepts, enabling creative teams to produce mood boards, test creative directions, and generate platform-optimised imagery without commissioning external shoots for every iteration. Industry analyses note that traditional marketing techniques are losing traction as businesses shift toward technologically advanced strategies that include social media marketing and search engine optimisation, with generative content creation positioned as a competitive requirement rather than an optional enhancement.

Textual content creation and copywriting assistance extend across captions, post drafts, and response templates. Tools such as ChatGPT enable marketers to generate multiple caption variants, adjust tone to match brand voice guidelines, and draft replies that community managers then review before publication. Research participants emphasise that AI-generated text serves as a starting point requiring human editing to ensure factual accuracy, cultural relevance, and alignment with brand messaging, a pattern reported across the interviewed marketing specialists, where automation accelerates first-draft production while human oversight remains the quality gate.

Customer engagement automation manifests through AI-powered chatbots embedded in social platform messaging systems. Interview data reveal that brands deploy chatbots to handle routine inquiries around the clock, with the systems learning from prior conversations to refine responses over time. These conversational agents pre-feed brand data and frequently asked question libraries, then surface answers without requiring live staff intervention for standard queries. The same generative models support personalised content recommendations and dynamic messaging: marketing platforms analyse user behaviour—such as product pages viewed or past engagement patterns—and generate tailored messages or offers that appear when the user revisits a social channel or branded environment. Practitioners note that this capability strengthens the bond between consumers, brands, and content by delivering customised experiences at scale, though it depends on robust data integration.

Efficiency and speed gains emerge as recurring justifications for adopting generative AI in social media workflows. Experts report significant time savings when automating tasks such as email marketing analytics, customer service transcription, and post scheduling, with some processes compressed from hours to minutes. However, adoption remains uneven: while larger organisations integrate AI into end-to-end social strategies, resource constraints and learning curves slow uptake among smaller teams. The literature underscores that successful deployment requires balancing automation's scale benefits with human curation to maintain authenticity and ensure that generated content aligns with evolving audience expectations (Amankwah-Amoah et al., 2024; Joshi et al., 2025).

1.5.2 Engagement Dynamics and Community Perception

Audience responses to AI-generated social media content reveal systematic differences in brand perception compared to human-authored posts. Experimental studies using disclosed AI-generated restaurant social media posts show that perceived brand authenticity, brand image, and self-brand congruity register significantly lower scores than human-created equivalents, suggesting that disclosure of AI authorship activates skepticism even when content quality remains comparable (Ali et al., 2025). These findings extend earlier advertising research by demonstrating that social media contexts—where audiences expect personal connection and community interaction—intensify authenticity concerns relative to traditional ad formats. Interview data from marketing practitioners confirm that AI-generated text requires human editing to ensure cultural relevance and brand alignment, implicitly acknowledging that automated outputs risk disconnection from community norms and audience expectations (Joshi et al., 2025).

The disclosure question looms large across regulatory and consumer-expectation domains. The U.S. Executive Order on Safe, Secure, and Trustworthy AI mandates clear disclosure of AI-generated material, while the EU AI Act and platform deepfake policies require synthetic or manipulated content to carry explicit flags to prevent deception (Duivenvoorde, 2024). These regulatory frameworks reflect broader consumer sentiment: studies report lower trust in AI-authored ad text absent transparency cues, even when emotional responses such as pleasure, arousal, and dominance show no significant difference between AI and human conditions (Sallaku et al., 2025). The implication for social media marketers is that undisclosed AI content may invite backlash if audiences discover automation post-publication, yet disclosed AI use triggers preemptive credibility penalties. Brands navigating this paradox increasingly adopt hybrid approaches, using AI for first drafts while routing final approval through human community managers who can interpret platform-specific discourse norms and adjust tone to maintain perceived authenticity (Joshi et al., 2025).

Consumer behavior shifts when AI involvement becomes known. The same restaurant study finds that while brand authenticity and brand image exert weaker effects on electronic word-of-mouth and behavioral intention under AI-disclosed conditions, self-brand congruity—the alignment between consumer identity and brand perception—gains significantly stronger predictive power (Ali et al., 2025). This suggests that when audiences recognize content as automated, they compensate by weighting identity fit more heavily, relying on whether the brand “feels like me” rather than whether it “seems genuine” in an absolute sense. Social media managers can interpret this as a strategic signal: AI-generated posts may succeed if they reflect audience values and lifestyle aspirations, even if they sacrifice the warmth of human authorship.

Emerging evidence also indicates that engagement outcomes depend on content type and platform context. Controlled evaluations show AI-generated visuals achieving comparable quality ratings to professional photography, and field tests report that AI image banners outperform stock photos on click-through rates in specific business-education campaigns (Hartmann et al., 2025). Yet these performance gains do not eliminate concerns about conceptual sameness or long-term community trust, particularly in high-involvement or hedonic categories where audiences tolerate AI imagery less readily (Belanche et al., 2025). Practitioners must therefore assess engagement not only through quantitative metrics—likes, shares, click-through rates—but also through qualitative signals such as comment sentiment, follower churn, and community feedback loops that surface when automation erodes perceived relationship quality.

1.5.3 Governance of Social AI Content

Regulatory frameworks targeting AI-generated social media content impose disclosure and marking obligations that reach beyond advertising-specific rules. The EU AI Act establishes transparency requirements under Article 50.2, mandating that providers of AI systems generating synthetic content mark such content as artificially generated or manipulated to help distinguish it from authentic material and protect consumers against deception. Article 50.4 extends disclosure duties to deployers of systems producing deepfakes—AI-generated or manipulated image, audio, or video resembling real persons or objects—requiring clear, distinguishable disclosure at first interaction or exposure. These provisions apply to professional users, including marketing teams publishing on social platforms. Parallel obligations appear in the Digital Services Act, which requires very large platforms to implement measures marking deepfakes through prominent markings when presented on their interfaces, covering synthetic content in organic posts alongside paid placements (Duivenvoorde, 2024).

Misinformation and deepfake risks surface as particular governance challenges in social media contexts where content spreads rapidly, and platform architectures amplify emotionally resonant material. Consumer protection scholarship highlights generative AI's capability to produce synthetic content at scale, while deepfake techniques have been weaponised to impersonate celebrities and public figures, damaging reputations and fueling consumer scams. Digital marketing experts interviewed in adoption studies cite high-profile deepfake incidents targeting celebrities as external concerns that deter AI adoption, fearing similar tactics could undermine consumer trust in brand communications (Joshi et al., 2025). The risk extends beyond outright fraud: synthetic content that appears real yet distorts product capabilities or misrepresents endorsements can mislead consumers into transactional decisions they would not have taken with accurate information. These dynamics position social media as a high-stakes environment where governance failures cascade into wider harm because sharing mechanisms and algorithmic distribution outpace manual review capacity (Duivenvoorde, 2024).

Content moderation challenges compound regulatory intent with practical enforcement gaps. While transparency rules specify disclosure obligations, they offer limited guidance on what constitutes “clear and distinguishable” marking in practice, leaving platforms and marketers to interpret compliance thresholds. Detection complexity arises when synthetic elements blend seamlessly with human-created posts, making automated filtering unreliable without continuous model retraining as generation techniques evolve. Scale issues further strain moderation: platforms process billions of daily posts, and manual review teams cannot inspect every item flagged by detection algorithms before publication. Enforcement asymmetries emerge when rules target deployers—the entities publishing content—yet some generative systems operate in automated pipelines where human oversight occurs only after content reaches audiences. The same transparency mandates apply across jurisdictions, but national enforcement agencies vary in resources, technical capacity, and willingness to sanction non-compliant actors, creating uneven pressure on platforms and publishers to meet disclosure standards consistently (Duivenvoorde, 2024).

Transparency mandates center on disclosure mechanisms, yet effectiveness depends on implementation details that current regulations leave underspecified. EU law and platform deepfake policies do not prescribe specific disclosure formats, introducing uncertainty about what wording, prominence, and timing satisfy legal obligations. Practitioners face a tension between minimal compliance—adding small-print disclaimers that technically meet rules—and substantive transparency that genuinely informs audiences. Scholarly frameworks for responsible AI advertising propose that human oversight and transparent communication build consumer trust, framing disclosure as an ethical obligation rather than a regulatory checkbox, though critical questions remain about whether disclosure alone can address risks when consumers may not attend to or comprehend technical labels (Sands et al., 2024). The practical implication is that governance of social AI content hinges on bridging the gap between legal mandates and operational execution: clear rules matter less if platforms lack tools to detect synthetic posts reliably, and disclosure obligations lose force if audiences cannot parse the labels marketers attach (Duivenvoorde, 2024).

1.5.4 Section Summary

Generative AI enables social media teams to scale content production across platforms while automating scheduling and customer interactions, yet disclosed AI authorship systematically reduces perceived brand authenticity, image, and congruity compared to human-created posts. This disclosure paradox intensifies under EU transparency mandates—the AI Act and DSA require clear synthetic-content marking to prevent deception, yet audiences penalize revealed automation by shifting reliance from authenticity to self-brand identity alignment. Deepfake risks and content moderation at scale create governance challenges: while laws require AI content labeling, they leave specific formats and implementation methods undefined, and enforcement struggles to keep pace with AI generation speed. Critical questions remain: which transparency mechanisms satisfy both legal obligations and consumer trust, and how content type and platform context influence whether AI content successfully engages audiences.

1.6 Generative AI in Content Branding

1.6.1 Brand Identity and Story Coherence

Brand identity comprises the distinctive voice, visual vocabulary, and narrative frameworks that firms deploy to signal their values and foster recognition across touchpoints (Ali et al., 2025). Traditional branding theory posits that coherent identity systems—encompassing tone, messaging consistency, and symbolic continuity—anchor consumer perceptions and facilitate self-brand congruity, the alignment between a consumer's self-concept and the brand's projected image (Ali et al., 2025). Creative content that enables brands to stand out and become memorable is central to differentiation in competitive markets (Osadchaya et al., 2024). Generative AI's capacity to produce text, imagery, and video at scale introduces both efficiency gains and identity risks for content branding strategies.

Empirical evidence from service-intensive contexts suggests that AI-generated branded content can erode perceived authenticity. In a restaurant-marketing context, AI-authored communications significantly lowered participants' assessments of brand authenticity and brand image relative to human-created equivalents; moreover, brand self-alignment became a stronger predictor of consumer response when AI generation was disclosed, implying that audiences scrutinize fit more intensely when machine authorship is known (Ali et al., 2025). Similarly, perceived humanness and trust mediate audience reactions to AI-infused luxury advertising, suggesting that brand narratives relying on emotional resonance risk dilution if AI origins undermine perceived sincerity (Jung et al., 2025). These findings underscore a tension: while generative tools expedite content production, disclosure or detection of AI use may weaken the affective bonds that brand storytelling traditionally cultivates.

Beyond authenticity concerns, AI adoption may homogenize brand voices across competitors. A "creativity versus sameness" paradox emerges in which generative models, trained on existing corpora, tend to reproduce prevailing tropes rather than yield genuinely novel narratives; practitioners interviewed report using AI outputs as "antimodels"—benchmarks that reveal generic ideas to be discarded—rather than as final creative material (Osadchaya et al., 2024). This echoes the concept of institutional isomorphism, wherein brands converge on similar messaging because their content creators draw from overlapping training data and industry norms (Osadchaya et al., 2024). For content branding, the risk is that AI-enabled efficiency amplifies uniformity, compromising the narrative distinctiveness that underpins brand differentiation, though brand-specific fine-tuning and prompt engineering may mitigate such convergence (Sands et al., 2024).

Maintaining coherent brand identity in AI workflows, therefore, demands deliberate human oversight. Humans serve as "gatekeepers of humanity," with advertisements blending real human characters with AI-generated elements eliciting higher perceived brand morality and lower audience anxiety than fully automated executions; this effect stems from audiences' valuation of human effort and emotional intelligence in creative processes (Madathil, 2025; Osadchaya et al., 2024). Transparency and human checkpoints serve as foundational principles for responsible AI advertising, since unchecked automation risks reputational damage when outputs misalign with brand values or contain factual errors (Sands et al., 2024). Taken together, these studies—conducted primarily in consumer-facing service and luxury sectors—suggest that generative AI can support brand storytelling when configured as a collaborative tool under continuous editorial stewardship, but risks diluting narrative coherence and authenticity when deployed as a wholesale replacement for human-authored content.

1.6.2 Personalisation and Customer Experience

Generative AI enables marketers to create tailored brand narratives at scale by generating content aligned with individual preferences. In expert interviews (n=11, indicating sample size), digital marketing professionals identified enhanced customization and personalization as a primary benefit of generative AI adoption, with 36% emphasizing its capacity to deliver deeper customer insights that inform product development and messaging strategies (Joshi et al., 2025). These systems process user behavior, demographic profiles, and contextual signals to produce individualized communications intended to match individual consumers better than generic messaging, though research confirming how consumers perceive this personalization remains limited (Ali et al., 2025).

Customer journey orchestration represents a second application domain where generative AI personalizes brand touchpoints across multichannel interactions. Algorithms can dynamically optimize messaging, adjusting tone, imagery, and promotional offers in response to real-time consumer behavior. Practitioners mentioned AI-driven chatbots as one application for maintaining engagement and delivering contextually relevant responses. Location-based personalization was cited as an illustrative example—displaying event promotions to users within specific geographic regions (Joshi et al., 2025). This allows brands to quickly update content at different points in the customer journey as consumer needs change.

Survey research examining perceived relevance outcomes reveals mixed patterns in how personalization influences brand perception and engagement. Restaurant industry experiments demonstrated through scenarios that self-brand congruity—the alignment between a consumer's self-concept and brand identity—emerged as the strongest predictor of behavioral intention and electronic word-of-mouth under AI-generated content conditions. When restaurants deployed AI to create social media content, consumers relied more heavily on psychological alignment with brand values than on traditional authenticity cues, suggesting that AI-generated personalization may preserve identity alignment even when traditional cues weaken (Ali et al., 2025). Expert interviews reported that enhanced customization and personalized recommendations improved customer engagement through chatbot interactions and dynamically optimized communication strategies. However, these gains must be weighed against data quality requirements, privacy trade-offs inherent in data-intensive personalization, and the risk that AI-generated content may lack the emotional depth consumers associate with human-crafted brand interactions (Ali et al., 2025; Joshi et al., 2025).

Boundary conditions emerge when personalization strategies misalign with consumer expectations for authenticity and transparency in brand communication. Restaurant industry research found that brand authenticity and brand image scores declined significantly when participants were exposed to AI-generated content compared to human-generated alternatives, indicating consumer skepticism toward automated personalization in high-touch service contexts where warmth and genuine human connection remain valued (Ali et al., 2025). Advertising agency practitioner interviews (n=25) revealed that agencies use generative AI primarily for brainstorming and prototyping rather than delivering final client-facing assets, citing concerns over intellectual property rights and copyright uncertainties surrounding AI-generated content (Serra-Simón et al., 2025). These findings suggest that personalization effectiveness varies by industry context and hinges on whether AI-enhanced individualization reinforces or undermines the brand's human-centric positioning, with consumers in the restaurant industry more accepting of automated personalization for routine orders than in situations where human warmth and service quality matter.

1.6.3 Ethical and Reputational Boundaries

Deploying generative AI in content branding raises ethical challenges across disclosure obligations, intellectual property rights, algorithmic bias, and cultural representation—domains where lapses can inflict lasting reputational damage. Transparency stands as a foundational principle: marketers must explicitly disclose AI involvement when it materially influences consumer perceptions or purchasing decisions (Sands et al., 2024). The EU Artificial Intelligence Act (AI Act) mandates that deployers of systems generating deepfake content—synthetic images, audio, or video resembling real persons or events—clearly flag such material as artificially generated at first exposure, while the Digital Services Act requires very large platforms to prominently mark deepfakes. These regulations aim to prevent consumer deception, particularly concerning fake celebrity endorsements or AI-generated influencers presented as authentic humans; failure to disclose may constitute misleading commercial practices under EU consumer protection law. However, regulatory clarity remains incomplete: the AI Act leaves marking formats underspecified, creating uncertainty about whether disclosure must be visible or embedded, and enforcement mechanisms lag the pace of content generation (Duivenvoorde, 2024).

Intellectual property concerns compound disclosure challenges. Using third-party creative works—advertisements, imagery, or branded content— to train generative models without authorization infringes creators' rights and exposes brands to legal liability. When organizations deploy employee or model likenesses in AI-generated brand narratives, ethical obligation demands active consent and fair compensation, such as royalties recognizing individuals' contributions to the final output (Sands et al., 2024). Practitioners interviewed by (Osadchaya et al., 2024) identified intellectual property protection as a primary concern, uncertain whether proprietary information fed into AI systems might be exposed or misappropriated, underscoring the need for contractual safeguards governing data usage and model training.

AI systems can show bias when their training data does not represent all groups, leading them to repeat and strengthen existing stereotypes. Image generators may create similar-looking people unless specifically told to show diversity; text generators may use stereotyped language or scenarios. Companies can instruct AI to create diverse content and use inclusive language, but these instructions only work if the company genuinely cares about diversity rather than just appearing to care. Adapting content for different regions adds more challenges: AI must respect local cultural differences while keeping the brand's core values, and humans must check that automated translations avoid offensive content. These ethical issues require governance systems that include transparency rules, rights checks, bias audits, and cultural sensitivity reviews before AI content reaches the public, making compliance both a legal requirement and a way to protect brand reputation (Duivenvoorde, 2024; Sands et al., 2024).

1.6.4 Section Summary

Generative AI offers content branding efficiency through scaled personalization and rapid iteration, yet introduces authenticity risks: consumers perceive AI-generated brand content as less authentic than human-created alternatives, particularly in service contexts valuing emotional connection. A creativity paradox emerges where AI-trained models reproduce common patterns rather than novel narratives, threatening brand differentiation. Ethical boundaries demand disclosure protocols, intellectual property safeguards, bias mitigation, and cultural sensitivity reviews to prevent reputational damage. Effectiveness varies by context—routine transactional touchpoints accept automation better than relationship-driven interactions. These findings suggest AI functions best as a collaborative tool under continuous human oversight rather than an autonomous content generator, raising questions about sustainable integration models that preserve brand coherence while capturing efficiency gains.

1.7 Organisational Readiness and Human–AI Collaboration

1.7.1 Capability Building and Workflow Integration

Interviews with digital marketing leaders emphasise that adoption hinges on closing a learning curve: organisations need to invest in training so staff can issue precise prompts, evaluate outputs, and blend AI assistance with their own expertise, especially where teams are not uniformly tech-savvy (Joshi et al., 2025). Broader creative industry analyses recommend structured upskilling programs and continuous learning schemes so marketers can adapt to shifting job boundaries and repurpose AI as a creative partner rather than a substitute for human craft (Amankwah-Amoah et al., 2024). These findings position capability building as an ongoing organisational investment that equips teams to steer generative tools confidently instead of treating them as drop-in automation.

Fieldwork with agencies reports that many generative tools are intuitive and slot into existing workflows without heavy training budgets, accelerating implementation in day-to-day production (Serra-Simón et al., 2025). Those practitioners still route outputs through human review to prevent homogenised results and maintain brand fit, underscoring that workflow integration pairs easy onboarding with deliberate oversight. In these contexts, production tasks such as mock-up creation shift toward AI support while creative staff retain the authority to approve variants and align them to client needs—reducing rework but preserving human judgement where sameness risk is high.

Interviewed marketers describe training and change programs as continuous, with learning-curve barriers addressed through short courses and on-the-job pairing rather than one-off tool rollouts (Amankwah-Amoah et al., 2024; Joshi et al., 2025). These programs focus on prompt clarity, output evaluation, and brand-safe reuse of assets, framing AI as an augmentation that frees time for higher-order tasks instead of replacement. By keeping skill development ongoing, teams reduce supervision overhead as familiarity grows, yet retain checkpoints for tone, accuracy, and context before external release.

Collaboration models that foreground human presence outperform automation-only setups in audience-facing video work, suggesting workflows should keep humans visible when brand morality or perceived creativity are at stake. In video ad experiments, blended human–AI executions scored higher on creativity and brand morality than fully synthetic ads, highlighting the value of pairing AI generation with human roles rather than removing people from the frame (Madathil, 2025). Agency interviews likewise noted that even when final assets are produced entirely with AI, a creative lead reviews them before client delivery to ensure channel fit and authenticity (Serra-Simón et al., 2025).

Empirical comparisons of AI and human asset production show that, in controlled image benchmarks, generative systems can match or surpass human-made stock photography on realism and quality while lowering production cost, shifting human labour toward orchestration and quality control rather than first-pass generation (Hartmann et al., 2025). Interviewed creatives also use ChatGPT as an “antimodel” to discard taglines that mirror its outputs, treating the model as a novelty check rather than a production source and elevating prompt design, curation, and variant selection to core daily tasks (Osadchaya et al., 2024). As these tasks expand, teams rebalance effort toward localisation, QA, and editorial reviews so human judgement governs how generative capacity is translated into brand-safe, channel-ready deliverables.

1.7.2 Decision Rights, Quality Assurance, and Governance

Decision rights in AI-assisted creative work remain anchored in human approval chains. Field interviews note that AI-generated ad copy and visuals are routed through creative leads for review, with teams treating automated drafts as inputs that must be checked before client or public release (Serra-Simón et al., 2025). Consumer protection guidance emphasises accountability for AI-generated content, placing responsibility on marketers to validate outputs against brand, legal, and audience standards rather than delegating final say to models (Duivenvoorde, 2024). Positioning humans as gatekeepers keeps automated assistance from displacing judgment on fit-for-purpose assets.

Quality assurance workflows focus on detecting hallucinations, bias, and brand misalignment before publication. Responsible-use guidance highlights that generative systems can surface plausible but incorrect or discriminatory content, requiring manual checks for factual accuracy, inclusivity, and tone (Sands et al., 2024). In practice, reviewers calibrate prompts, compare variants, and reject outputs that feel generic or misaligned with context, maintaining diversity and authenticity in advertising, social, and branding outputs (Serra-Simón et al., 2025). These checkpoints formalise AI use as draft generation rather than autopilot, with human reviewers accountable for any content that reaches audiences.

Approval processes extend to intellectual property, data, and disclosure considerations. Legal analyses warn that AI-generated copy and imagery can create uncertainty over copyright ownership and may infringe third-party rights when training data is mirrored in outputs, prompting clearance steps before use (Duivenvoorde, 2024). Governance policies, therefore, require teams to verify licenses for embedded references, document prompt inputs, and record modifications so provenance is traceable. These controls give brand and legal functions the authority to halt or rework assets that risk IP breaches or misleading claims.

Risk controls also address organisational exposure from supplier and platform dependencies. Industry fieldwork reports limited visibility into how agencies or platforms apply generative tools on behalf of brands (Serra-Simón et al., 2025). Marketers, therefore, seek greater transparency on model use, safety filters, and data handling, and may add audit rights to contracts to close visibility gaps. Internally, teams can log prompt histories, model versions, and human approval decisions to support remediation if errors surface after deployment (Sands et al., 2024). These practices are recommended safeguards rather than evidence of uniform adoption.

Governance ties quality controls directly to compliance and reputation protection. By setting clear human review thresholds, IP checks, and disclosure checks plus supplier transparency expectations, organisations lower the likelihood that automated content will trigger consumer deception, bias incidents, or rights disputes, while recognising that residual risk remains (Duivenvoorde, 2024; Sands et al., 2024). These safeguards apply across advertising, social, and branding outputs, aligning rapid AI-assisted production with accountable decision rights and auditable quality assurance.

1.7.3 Section Summary

Marketing teams treat generative AI as an augmentation that demands continuous training, prompt clarity, and manual review, pairing fast onboarding with human oversight to avoid homogenisation and safeguard brand fit. Quality assurance remains human-led to screen hallucinations, bias, and tone, while approval chains address IP, disclosure, and data concerns before release. Visibility gaps with agencies and platforms prompt safeguards such as transparency commitments, optional audit rights, and traceable edits, with residual risk acknowledged. Taken together, these capability investments and governance controls let teams speed up AI-assisted production while keeping people responsible for what is approved and published.

1.8 Conceptual Synthesis

1.8.1 Integrated Conceptual Model

This model synthesizes the preceding theoretical analysis into a unified framework that explains how and why organizations adopt generative AI for creative marketing workflows, and under what conditions these adoptions lead to desired outcomes. It maps the adoption decision, implementation mechanisms, contextual factors, and resulting outcomes, directly addressing the research aim of analyzing generative AI's impact on creative workflows, brand storytelling, and consumer perception across advertising, social media, and content branding domains. Figure 1 visualises the driver–mechanism–outcome linkages that structure this model.

Organizational decisions to integrate generative AI stem from five interconnected drivers: - Strategic innovation pressures push firms to seek novel creative methods (Amankwah-Amoah et al., 2024; Corvello, 2025). - Speed and efficiency imperatives arise from content-hungry digital environments where deadlines shorten (Osadchaya et al., 2024; Serra-Simón et al., 2025). - Competitive parity concerns drive defensive adoption as agencies fear losing clients to faster rivals (Joshi et al., 2025). - Cost and resource constraints motivate leaner teams to substitute AI-generated drafts for expensive human labor (Serra-Simón et al., 2025). - Data and technological readiness determines integration feasibility for firms with mature infrastructures (Joshi et al., 2025).

Adoption translates into practice through four core mechanisms shaping daily workflow: - End-to-end integration embeds AI across ideation, asset production, testing, and refinement, speeding cycles and widening exploration (Joshi et al., 2025; Serra-Simón et al., 2025). - Human–AI collaboration patterns define task allocation, with organizations assigning AI to data synthesis and rapid prototyping while reserving human judgment for strategic framing and emotional nuance (Osadchaya et al., 2024; Serra-Simón et al., 2025). - Brand-safety and guardrail processes institutionalize oversight through prompt-engineering protocols, hallucination checks, IP scans, and bias audits to mitigate reputational risks (Sands et al., 2024). - Feedback and iteration loops enable continuous refinement as agencies test AI-generated variants and feed performance signals back into prompt adjustments (Osadchaya et al., 2024).

The strength and direction of these relationships depend on several factors. Industry and brand maturity moderate success: established brands with documented voice guidelines find AI easier to tune (Amankwah-Amoah et al., 2024). Campaign type matters: data-driven personalization fits utilitarian, high-involvement products, while hedonic, low-involvement contexts need emotional tone that AI may miss (Ali et al., 2025). Regulatory sensitivity varies by geography, shaping disclosure transparency (Sands et al., 2024), while team structure and organizational size influence whether

AI augments workflows or causes disruption (Joshi et al., 2025; Serra-Simón et al., 2025).

These outcomes span six dimensions organizations must balance: - Creative quality and novelty capture whether campaigns produce distinctive ideas or replicate training patterns (Osadchaya et al., 2024). - Brand consistency and authenticity reflect alignment with established identity (Ali et al., 2025). - Time-to-market measures speed gains (Serra-Simón et al., 2025). - Cost efficiency quantifies budget savings (Joshi et al., 2025). - Consumer response tracks behavioral metrics. - Perceived authenticity reflects whether audiences think the content is genuine, with disclosure effects varying by context (Ali et al., 2025; Jung et al., 2025).

Firms embedding AI across full lifecycles while maintaining rigorous oversight gain speed without losing distinctiveness, while loose, unchecked automation makes outputs uneven and undermines trust (Osadchaya et al., 2024; Sands et al., 2024). Feedback loops complete the cycle as outcome data inform iterative mechanism adjustments.

The model surfaces inherent tensions that cannot be resolved through simple optimization. Research identifies operational paradoxes where generative AI can both support and hinder research rigor, expand and limit creative exploration, and speed production while adding oversight work (Osadchaya et al., 2024). Other work stresses the psychological side: teams need to treat AI as a collaborator while keeping humans as gatekeepers of brand authenticity (Madathil, 2025). These are not mere trade-offs but paradoxical relationships requiring ongoing negotiation rather than one-time allocation decisions.

This driver-mechanism-outcome architecture draws on Behavioral Reasoning Theory's distinction between facilitating factors and barriers (Joshi et al., 2025), a paradox lens highlighting operational and psychological tensions in AI-mediated advertising work (Osadchaya et al., 2024), and field evidence showing how implementation choices vary with organizational readiness and governance maturity (Serra-Simón et al., 2025). By integrating these perspectives, the model shows both why organizations choose generative AI and the tensions that surface after implementation, giving a clear map for studying how generative AI affects creative marketing.

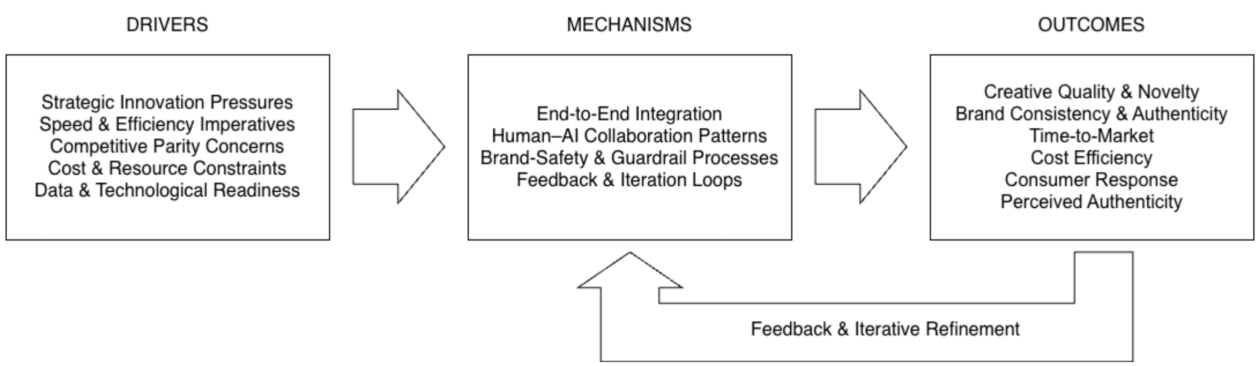

Figure 1. Integrated Conceptual Model of Generative AI in Creative Marketing

1.8.2 Theoretical Propositions or Guiding Questions

These questions operationalise the integrated model into empirical inquiries that inform data collection and thematic analysis. What organizational pressures, infrastructure capabilities, and competitive dynamics shape initial decisions to explore generative AI for creative workflows, and how do governance concerns around intellectual property, brand safety, and regulatory compliance either accelerate or constrain that exploration? When do cost efficiencies and speed imperatives outweigh concerns about creative dilution or workforce displacement, and conversely, under what conditions do organizations retreat from automation toward selective human-AI task allocation?

How do organizations embedding AI across ideation, asset production, and campaign testing perceive shifts in creative output quality, brand voice consistency, and operational tempo, and what feedback mechanisms allow teams to detect when AI-generated content drifts from brand identity or audience expectations? When AI-generated materials are labeled for consumers—by choice or by rule—how do brand-safety checks, editorial review, and transparency notices shape audience views of authenticity, and do those effects differ by product type and disclosure format?

To what extent do product characteristics—particularly whether offerings emphasize functional performance or emotional resonance—condition consumer acceptance of AI-assisted creative work, and do practitioners observe systematic differences in how audiences respond to disclosed AI content in hedonic versus utilitarian contexts? How do broader contextual factors, such as industry creative maturity, existing brand voice codification, geographic regulatory environments, and team structures, influence whether AI adoption enhances creative differentiation and operational agility or instead homogenizes output and undermines authenticity? What tensions do organizations experience between efficiency gains enabled by scaled AI content generation and the imperative to maintain creative distinctiveness in crowded markets, and how do practitioners conceptualize the boundary between AI as a productivity tool and AI as a creative collaborator?

These questions prioritize practitioner-level adoption dynamics and consumer-facing outcomes over technical AI capabilities because the thesis examines generative AI's impact on creative marketing practice, not algorithmic performance. The emphasis on contextual factors—campaign type, industry maturity, regulatory environment, team structure—reflects evidence that generative AI effects are highly context-dependent rather than universal. Answers to these questions will advance understanding of generative AI's practical integration challenges, success factors, and boundary conditions in advertising, social media, and content branding domains.

The methodological section that follows specifies how these questions guide participant selection, interview protocol design, case-study boundaries, and analytical procedures.

2 A METHODOLOGICAL PART

2.1 Research Design and Approach

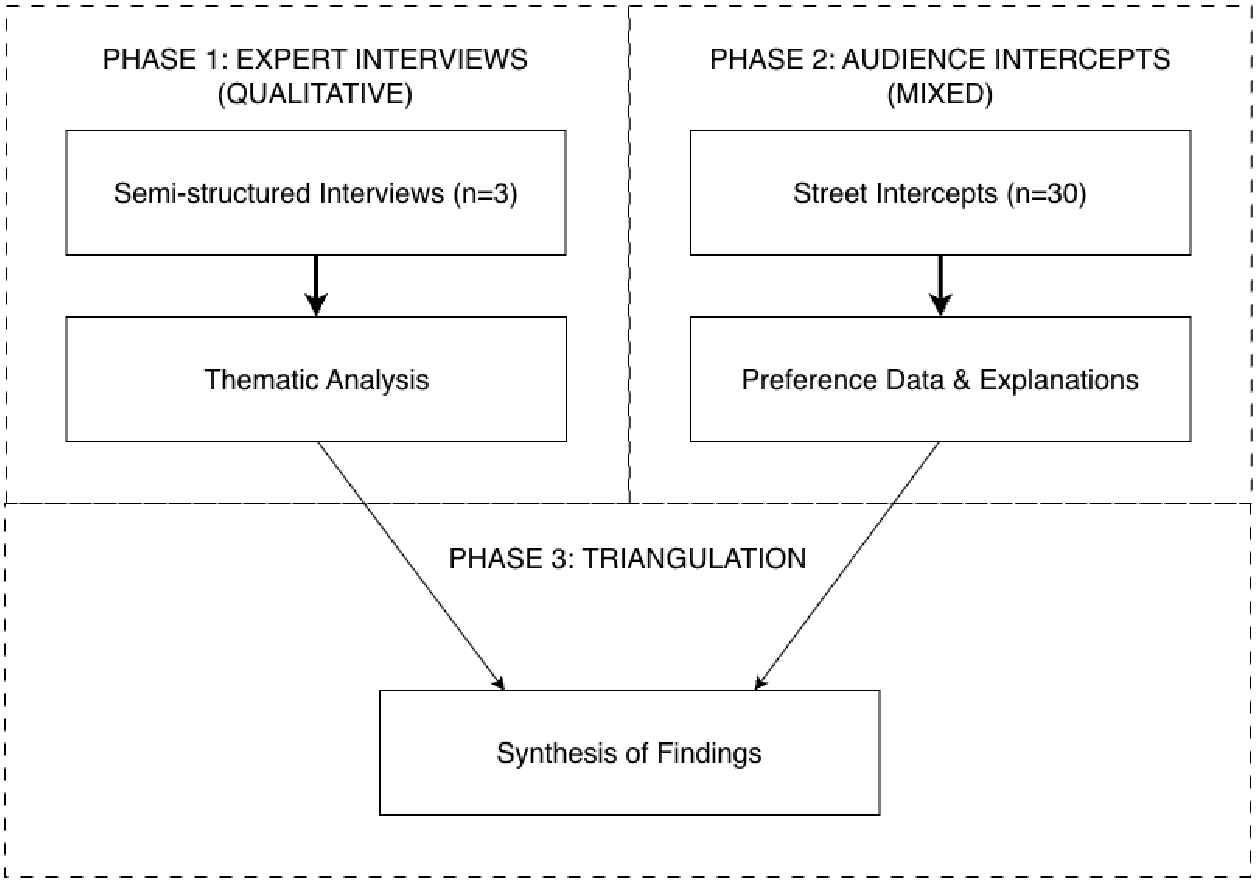

Figure 2. Methodology Flowchart: Mixed-Methods Design

This exploratory study adopts a qualitative-dominant mixed-methods design (Creswell & Plano Clark, 2014) combining semi-structured expert interviews with a quasi-experimental audience intercept survey to contextualize and extend the theoretical synthesis presented in preceding sections. The research does not aim to produce standalone empirical generalisations but rather to combine practitioner perspectives with consumer responses, grounding the conceptual model in lived experience and observable preference patterns. Different methods address different aspects of the research aim: interviews capture workflow transformation from the practitioner viewpoint, while audience intercepts reveal reception patterns that complement expert testimony.

The qualitative core comprises semi-structured interviews (Brinkmann & Kvale, 2013) with marketing practitioners who actively use generative AI in creative workflows. These interviews explore adoption drivers, implementation pathways, observed outcomes, and organizational challenges, addressing the guiding questions posed in the conceptual synthesis. Expert testimony provides insight into how teams balance automation with human oversight, what efficiency and cost gains they achieve, and which safeguards maintain quality and brand alignment. This method aligns with qualitative research principles that prioritize depth and information-rich cases over numerical breadth when research questions are focused, and participants offer specialized expertise (Guest et al., 2006). The approach suits the thesis aim to analyse how generative AI reshapes creative workflows and brand storytelling from the perspective of those deploying the technology.

The audience survey component employs a quasi-experimental pre-post disclosure design administered via street intercepts (n=30). Participants view one AI-generated and one human-created commercial in a fixed sequence (AI-generated first, then human-created), indicate their preference and reasoning, then receive disclosure of the production method and state whether this knowledge changes their preference. The fixed sequence examines one disclosure timing scenario where AI provenance is revealed after initial evaluation, a design choice that captures preference shifts while acknowledging that alternative orderings may yield different patterns. This approach examines transparency and disclosure effects relevant to findings documented in the literature (Ali et al., 2025; Jung et al., 2025; Sallaku et al., 2025), combining quantitative preference measures with qualitative reasoning through open-ended explanations.

Together, the expert and audience methods address the dual focus of the research aim: experts reveal workflow transformation that practitioners observe directly, while audience intercepts capture reception patterns that practitioners cannot observe without external validation. The design treats qualitative depth as primary while using structured preference measures to explore practitioner claims about disclosure and authenticity effects, following purposive sampling principles (Patton, 2015) appropriate for exploratory studies that contextualize theoretical predictions rather than produce standalone empirical contributions.

2.2 Sampling and Participant Criteria

The study employs purposive sampling for both expert interviews and audience intercepts. This sampling strategy prioritizes theoretical insight and depth over statistical generalizability, accepting the constraint that findings will be context-bound rather than population-representative.

Expert Interview Sampling. The expert sample (n=3) targets creative professionals or marketing managers who have directly implemented generative AI solutions in brand processes. Inclusion criteria require participants to be directly responsible for applying tools such as ChatGPT, Midjourney, DALL-E, Runway, or Sora to ideation, copywriting, and visual production workflows. Eligible experts must evaluate AI platforms by quality, controllability, speed, and cost while coordinating human-AI collaboration; assess return on investment, team competency development, and ethics and transparency standards; and demonstrate the ability to identify near-term opportunities and risks in AI-driven creative work. Recruitment proceeded through convenience and snowball methods (professional networks, industry conferences, social media outreach on LinkedIn, Facebook, and Instagram), targeting practitioners with active deployment experience rather than conceptual familiarity alone. The sample size aligns with qualitative research precedents for exploratory interviews with specialized practitioners, where depth and relevance outweigh numerical breadth. Research shows that 3–6 interviews can achieve saturation when questions are focused and participants are information-rich (Guest et al., 2006); the present sample (n=3) sits at the lower bound of this threshold, justified by the study's exploratory validation scope and intensive semi-structured format within master's thesis constraints.

Table 1. Expert Participant Profiles

ID	Role	Experience	Sector
E1	Agency Head	Senior	Digital Marketing (B2B/B2C)
E2	Production Studio Lead	Senior	Visual Production
E3	Independent Creator	Senior	Video Production & Training

Audience Survey Sampling. The audience sample (n=30) draws from public spaces in Stockholm, Sweden — a shopping mall in the city center and commuter trains running from the outskirts to the city center. These locations were selected to access a range of ages and origins within Stockholm's metropolitan population, recognizing geographic and temporal constraints, with the commuter route connecting suburban residential areas to central commercial districts. Inclusion criteria required participants to be adults (18+) or minors accompanied by a parent or guardian who provided consent, willing to participate, and able to view video content on a laptop. Exclusion criteria eliminated participants under 18 unaccompanied by a legal guardian to avoid potential ethical concerns when involving minors in research. The sample size (n=30) balances exploratory scope with practical feasibility of street intercepts conducted in public spaces, aligning with mixed-methods sampling conventions that accommodate modest quantitative components alongside qualitative depth (Onwuegbuzie & Collins, 2007). The semi-structured protocol captures qualitative reasoning over four questions with a range of participant ages and national origins, enabling preliminary observation of boundary conditions across consumer segments within a single metropolitan context.

2.3 Data Collection Instruments and Procedures

Data collection employed two instruments aligned with the mixed-methods design: a semi-structured expert interview protocol (Brinkmann & Kvale, 2013) and a pre-post disclosure audience questionnaire (Madathil, 2025).

Expert Interview Instrument and Procedures. The expert protocol comprises 12 open-ended questions organized into five thematic sections: (1) background and context, establishing the participant's role and AI experience; (2) application and tools, identifying usage areas and platform preferences; (3) perceived value and outcomes, probing efficiency, cost, and creative quality impacts; (4) ethical, strategic, and organizational aspects, exploring human-AI balance, transparency challenges, and team adaptation; and (5) strategic outlook, capturing forward-looking opportunities and risks. Each question includes an explanation of intent and an illustrative example answer to guide participant responses while maintaining flexibility for emergent themes. The full protocol appears in Appendix A (Expert_Questionnaire_GenerativeAI_Guided.pdf).

Interviews were conducted between November 25 and 29 through voice messages exchanged via social media messaging platforms (Instagram, Facebook Messenger), with one real-time interview recorded in its entirety. This approach allowed participants to respond when and where they preferred, fitting the interview around their work schedules while providing time for thoughtful answers. Participants selected their preferred language (English or Lithuanian), and interviews ranged generally from 15 to 25 minutes. Voice messages were transcribed verbatim into text for analysis by the researcher, with Lithuanian responses translated to English for coding, preserving participant phrasing and tone while enabling systematic coding.

Audience Survey Instrument and Procedures. The audience protocol employs a four-question pre-post disclosure design. Participants first view two Coca-Cola commercials in fixed sequence—the 2025 “Holidays Are Coming” advertisement created using generative AI, followed by the 2020 “Holidays Are Coming” commercial produced through conventional filming methods—both presented without sound on a laptop. After viewing, Question 1 asks which commercial participants prefer (first or second) and invites open explanation. The researcher then discloses production methods, and Question 2 asks which commercial participants now prefer (AI-generated or human-made), whether this knowledge changed their view, and how. Questions 3 and 4 collect demographic data (age and origin). The complete protocol appears in Appendix B (Audience_Questionnaire.md).

Street intercepts were conducted on November 29–30 in the mornings and afternoons at the designated Stockholm locations. Recruitment targeted individuals already seated and visibly scrolling on mobile devices rather than those walking with apparent destination intent, minimizing disruption and selecting for available attention. For younger-appearing participants, the researcher verbally confirmed age eligibility; when minors were present with a parent or guardian, guardian consent was obtained before proceeding. Participants viewed commercials on the researcher’s laptop, and the researcher recorded responses using the iPhone Voice Memos application after obtaining verbal consent for audio capture. Intercepts ranged from 1 to 4 minutes, averaging approximately 2 minutes. Interviews were conducted in English.

2.4 Data Analysis Methods

Analysis followed a two-stream approach aligned with an exploratory mixed-methods design (Creswell & Plano Clark, 2014), with expert interviews and audience surveys analyzed separately before triangulation.

Expert Interview Analysis. Transcribed interviews were analyzed using thematic coding with AI assistance. The approach combined structural coding based on the interview protocol’s five sections (Background, Tools, Value, Ethics, Outlook) with open coding to identify patterns within each section (Saldaña, 2016). Rather than multiple coding rounds, analysis proceeded in a single pass appropriate for small, manageable datasets (Guest et al., 2006): transcripts were read, and quotes were tagged with theme labels as patterns emerged. AI-assisted coding tools (Claude 4.5 Sonnet, ChatGPT-5.1, Gemini 2.5, and 3 Pro) were used to generate initial theme suggestions, which were then manually reviewed, refined, and validated against transcript data to ensure grounding in participants’ actual language. A simple codebook was developed iteratively to maintain consistency, recording theme definitions and example quotes as coding progressed. All data were organized in CSV format and analyzed using Python with the pandas library, enabling organized theme tracking across the three expert transcripts.

Audience Preference Analysis. Pre-post disclosure preference shifts were analyzed through both descriptive and inferential methods. Descriptive analysis counted how many participants changed their preference after learning production methods, categorizing shifts as “no change,” “shifted toward AI,” or “shifted away from AI.” McNemar’s test (McNemar, 1947) for paired categorical data was applied to determine whether disclosure significantly changed preference distributions, appropriate for the within-subjects design where each participant provided both pre- and post-disclosure choices. Demographic variables (age, origin) were cross-tabulated with preference patterns to explore whether responses varied by participant characteristics. Analysis was conducted in Python using pandas for data organization and statsmodels for statistical testing.

Audience Explanation Analysis. Open-ended explanations accompanying preference choices were categorized into broad response types rather than exhaustively coded. Common reasoning patterns—such as quality-focused responses, authenticity concerns, indifference to production method, or surprise at disclosure—were identified by reading through all explanations and grouping similar rationales. This lighter categorization approach suited the brief, conversational nature of street-intercept responses while still capturing the range of participant reasoning.

Triangulation Strategy. After separate analysis, expert and audience findings were triangulated by extracting expert predictions from interview transcripts and systematically comparing these predictions with observed audience preference shifts and explanation patterns to identify areas of alignment or contradiction. Expert predictions were documented with supporting quotes and categorized by topic (e.g., disclosure effects, quality concerns, authenticity sensitivity). Each prediction was then matched against audience data and classified as supported, partially supported, unsupported, or unable to assess based on available evidence. For example, if experts predicted that disclosure would reduce trust in AI-generated content, audience preference shifts and post-disclosure explanations were reviewed to confirm or challenge that prediction. Conversely, when experts emphasized specific concerns such as authenticity or visual quality, audience explanations were examined for evidence of those same concerns. This bidirectional comparison allowed expert testimony to contextualize audience reactions while audience data grounded expert claims in observable consumer behavior.

2.5 Validity, Reliability, and Ethics

This study adopts Lincoln and Guba's (1985) trustworthiness framework, establishing credibility, transferability, dependability, and confirmability as quality criteria for naturalistic inquiry (Lincoln & Guba, 1985). Credibility was strengthened through integrating multiple participant perspectives—expert interviews and audience surveys—comparing practitioner predictions about consumer reactions with actual audience responses (Patton, 2015).

Dependability relies on a codebook documenting each theme with a definition, inclusion criteria, and exemplars, providing a clear record of how insights were categorized. Because the expert sample is small (n=3) and interviews were conducted by a single researcher, inter-coder reliability checks were not performed. This is an acknowledged limitation, balanced by transparent codebook documentation, preservation of full transcripts, and systematic categorization protocols that allow readers to trace how conclusions were reached. Audience explanations follow a similar categorization process, grouping responses into 4–6 common reasoning types rather than exhaustive thematic coding.

Transferability depends on a thick description of the research context. The audience survey occurred in a shopping mall in the city center, and commuter trains running from the outskirts to the city center of Stockholm on November 29–30, 2025, targeting diverse foot traffic in public spaces. Participants viewed Coca-Cola holiday commercials from 2020 (human) and 2025 (AI-generated) on a laptop screen. Demographic information (age, origin country) and open-ended explanations provide context for interpreting preference patterns. Expert interviews occurred November 25–29, 2025, via asynchronous voice messages with marketing practitioners selected for hands-on experience implementing generative AI tools in campaign work. Sampling criteria and participant characteristics are detailed in the preceding subsection. This contextual specificity allows readers to judge whether findings transfer to their own settings, a core principle of naturalistic inquiry's transferability criterion (Lincoln & Guba, 1985).

Verbal informed consent was obtained from all participants before data collection, and all were informed of their right to withdraw at any time without penalty. For minors in the audience sample, guardian consent was obtained on-site; one participant's Swedish-language responses were translated into English by their parent. Anonymity is preserved by omitting names, employer identities, and specific job titles from expert accounts. Given the small expert sample (n=3) in a specialized practitioner community, contextual details in findings have been edited to reduce re-identification risk. Audience data reports origin countries and age ranges rather than exact ages. Expert interview transcripts are stored in password-protected private repositories with local encrypted backups. Original voice messages remain in private social media conversations (Facebook, Instagram), representing a data storage limitation mitigated by maintaining encrypted transcript copies. Expert interview data will be retained permanently as research artifacts, while audience intercept recordings will be deleted following the thesis defense.

2.6 Limitations and Delimitations

Several limitations constrain the interpretation of findings. The audience survey used a fixed presentation order (AI-generated commercial first, then human-created), introducing confounded order and disclosure effects that prevent isolating preference drivers and limit causal interpretation. Small sample sizes restrict generalizability: three expert interviews provide depth but cannot represent the full range of marketing practitioner perspectives, while thirty audience intercepts offer preliminary patterns rather than statistically robust population estimates (Onwuegbuzie & Collins, 2007). Geographic constraints further narrow the scope—audience data come exclusively from central Stockholm locations during late November 2025, and findings may not transfer to other Scandinavian cities, broader European contexts, or cultural settings outside northern Europe. The stimulus materials represent one brand (Coca-Cola), one advertising genre (holiday emotional appeal), one disclosure timing (post-viewing), one ad format (video), and one involvement level (low-involvement product category), restricting claims about AI-generated content in other product categories, creative styles, disclosure approaches, or high-involvement purchases.

Expert interviews used a voice-message exchange format that constrained real-time probing and clarification compared to live conversations. Audience intercepts relied on convenience sampling in high-traffic public locations, introducing selection bias—participants who stop differ systematically from those who decline in time availability, social orientation, and language comfort (intercepts conducted primarily in English in multilingual Stockholm). The paired-sample analysis (n=30) has limited statistical power, meaning null findings may reflect Type II error rather than true absence of effects (McNemar, 1947). Self-reported preferences may diverge from actual behavior due to social desirability bias or attitude-behavior gaps. Short engagement time (approximately two minutes per participant) limits the depth of responses. Qualitative coding was conducted by a single researcher without inter-coder reliability checks, a standard limitation in solo exploratory studies but one that constrains claims of interpretive consistency (Lincoln & Guba, 1985).

Delimitations reflect deliberate scope choices constrained by access and resources. This study examines advertising, social media, and content branding domains where generative AI adoption is most visible in professional discourse, while excluding other marketing functions such as analytics, customer relationship management, or pricing that engage AI differently. Expert recruitment prioritized hands-on practitioners accessible within study timelines rather than organizational case studies requiring extended access negotiation. The exploratory design trades breadth for depth, yielding rich contextual insights whose transferability readers must judge against their own settings (Lincoln & Guba, 1985).

2.7 Section Summary

This exploratory study employs a qualitative-dominant mixed-methods design combining semi-structured expert interviews (n=3) with quasi-experimental audience intercepts (n=30) to ground the conceptual model in practitioner experience and consumer response patterns. Expert interviews were conducted via asynchronous voice messages (November 25–29, 2025) and analyzed through thematic coding with AI assistance, while audience intercepts in Stockholm public spaces (November 29–30, 2025) employed a pre-post disclosure design examining preference shifts after learning production methods. Analysis triangulated expert predictions against observed audience reactions using McNemar's test for preference data and categorical grouping for open explanations. Credibility was strengthened through multiple participant perspectives and systematic codebook protocols, though single-researcher coding and absence of inter-coder checks constrain interpretive consistency claims. Key limitations include small sample sizes precluding generalization, geographic clustering in one metropolitan area restricting transferability, fixed presentation order confounding causal interpretation, and single-brand stimulus limiting cross-category claims. The design prioritizes depth over breadth, yielding context-rich insights into how practitioners deploy generative AI in creative marketing and how audiences respond to disclosed AI-generated content, directly addressing the guiding questions on adoption drivers, implementation mechanisms, outcomes, and boundary conditions articulated in the conceptual synthesis.

3 A RESEARCH PART

3.1 Data Corpus and Preparation

The study generated two complementary datasets: expert interview transcripts and audience intercept responses, each prepared through distinct workflows aligned with the exploratory mixed-methods design.

Expert Interview Corpus. Three expert interviews (n=3) were conducted via asynchronous voice messages exchanged through social media platforms (Instagram, Facebook Messenger) between November 25 and 29, 2025. Audio recordings were transcribed verbatim by the researcher, preserving participant phrasing, pauses, and conversational tone to maintain semantic fidelity (Brinkmann & Kvale, 2013). Two of the three interviews were conducted in Lithuanian and subsequently translated into English by the researcher, while one interview was conducted entirely in English. Lithuanian transcripts were retained separately to preserve the original linguistic context. English versions served as the primary corpus for thematic coding, ensuring consistent analytical treatment across all three participants.

Anonymization procedures removed all personally identifying information from expert transcripts before analysis. Participant names were replaced with numeric identifiers, employer names and specific client references were redacted, and contextual details that could enable re-identification (such as precise geographic locations or project names) were generalized while preserving substantive content. Anonymized transcripts were separated from original recordings, with access to identifiable materials restricted to the researcher following informed-consent commitments made during recruitment.

Audience Intercept Corpus. Thirty street intercept interviews (n=30) were conducted on November 29–30, 2025, in central Stockholm public spaces (shopping mall, commuter trains). Each intercept followed a structured four-question protocol: participants viewed two Coca-Cola commercials (AI-generated 2025 and human-created 2020) in a fixed sequence, stated their preference with explanation, received production-method disclosure, and indicated whether disclosure changed their view. Audio responses were captured using a digital voice-recording application after obtaining verbal consent, then transcribed verbatim. Transcripts preserved participant reasoning verbatim, including hesitations and clarifications, to support qualitative explanation analysis. One participant's responses were provided in Swedish and translated into English by an accompanying parent; all other intercepts were conducted in English. Demographic data (age, origin country) were recorded during the intercept and linked to each transcript.

Participant names were never collected, and transcripts identify individuals only by sequential identifier, age, origin, and gender. Audio recordings remain in the researcher's secure archive and will be deleted following the thesis defense.

Data Staging for Analysis. Both expert and audience corpora were transformed into structured tabular formats to enable systematic coding and statistical testing. Expert transcripts were segmented into thematic units and coded using AI-assisted tools (Claude, ChatGPT, Gemini) to generate initial theme suggestions, which were then manually reviewed, refined, and organized into a codebook documenting theme names, definitions, and exemplar quotes. Coded expert data were organized in linked tables capturing quote-to-theme mappings and theme frequency summaries across participants. Audience responses were tabulated to capture pre-disclosure choice, post-disclosure choice, preference shift category, and demographics for each participant. Qualitative explanations were extracted and categorized by reasoning type. McNemar test inputs, demographic breakdowns, and preference shift summaries were organized in separate analytical tables. All structured data were processed systematically, enabling transparent linkage between raw transcripts and reported findings.

3.2 Expert Interviews: Participant Profile

The three expert participants represent complementary professional domains within creative marketing and production, offering perspectives across strategy, execution, and education. Expert 1 operates a full-service marketing agency focused on LinkedIn communication and broader brand strategy for B2B and B2C clients spanning technology, logistics, and consumer sectors. This participant integrates generative AI across the entire agency workflow—from brainstorming and strategy development to offer preparation, campaign personalization, and content creation—employing ChatGPT, Gemini, Claude, MidJourney, DALL-E, Higgsfield, Sora, and automation platforms (n8n, Zapier AI agents) (AI, 2024; *n8n*, 2025; Zapier Inc., 2025). The agency has institutionalized AI adoption by hiring a dedicated AI lead responsible for process automation and staff training, reflecting organizational commitment to scaled implementation.

Expert 2 leads a visual production studio specializing in computer-generated content and AI-assisted video creation, serving clients across commercial advertising sectors. This participant entered generative AI experimentation in 2021 using early diffusion models (Stable Diffusion, Deforum) and now deploys a comprehensive tool suite including MidJourney, DALL-E, Runway, Sora, Kling, Comfy UI, Freepik, and ChatGPT across ideation, mood boarding, character creation, and final asset production (comfyanonymous, 2023; Freepik, 2023; Rombach et al., 2022; Technology, 2024). The studio's workflow treats AI as a cost-effective alternative to traditional CGI pipelines while maintaining manual post-processing to achieve client-specific precision.

Expert 3 brings a technical production and education background, combining cinematography, editing, and animation expertise with four years of teaching experience in Adobe Creative Suite (Photoshop, Premiere, After Effects) (Adobe Inc., 1990, 1991, 1993) and AI tools. Currently active in both commercial video production (clips, commercials) and corporate training (conferences, workshops), this participant emphasizes cutting-edge video generation platforms—Kling, Veo 3.1, Hailuo, Luma Dream Machine—selected for their market-leading capabilities (Google, 2024; L. Labs, 2024; MiniMax, 2025; Technology, 2024). The dual role as practitioner and educator offers insight into both hands-on production constraints and the skill development challenges facing marketing teams adopting generative tools.

Together, these profiles span strategic, operational, and technical adoption contexts, covering B2B and B2C applications, agency and studio environments, and diverse creative outputs (text, static imagery, video, automation). This variation supports pattern identification across organizational structures, client expectations, and tool ecosystems, directly addressing the study's focus on workflow transformation, human–AI collaboration, and capability building outlined in the conceptual model.

3.3 Expert Interviews: Adoption Drivers and Objectives

Expert interviews show adoption reasons that match the theoretical drivers from the conceptual model while revealing limits and competitive pressures missing from explanations focused only on efficiency. Cost and resource pressures are the main reasons across all three experts. Expert 2 described adoption as a survival response to flat or declining client budgets: "their budgets have not changed and to be fair in some cases they've just dropped for the visual production… the AI here is very good for marketing agencies for production agencies who are trying to survive in this game." The studio used AI as a replacement for computer-generated imagery workflows, presenting the technology as delivering similar visual quality at lower cost. Expert 1 agreed with this reasoning, noting that AI tools let the agency "accept more clients under the same amount of people and simply improve business efficiency" while reducing dependence on external production houses. Expert 3 confirmed cost-reduction benefits, noting AI "speeds up, facilitates and reduces costs for sure" when used for appropriate content types. These accounts position AI adoption as a protective strategy when revenue is limited and traditional production costs no longer support profit margins.

Speed and efficiency goals are linked to cost concerns rather than separate reasons. Expert 1 reported that AI "gives much more speed to deliver things" while enabling hyper-personalization through rapid variant creation: "you can have not ten manually created visuals anymore, but twenty, thirty various variations… and you can test again and advertise very directly to that, who is attracted by that specific one visual out of the hundred." This ability to create more content does more than make work easier—it changes campaign strategy, allowing large-scale testing that was previously too expensive. Expert 2 emphasized content volume gains, stating "we are able to do more content… for the same client" within unchanged budget limits. Expert 3 described speed benefits in practical terms: "everything is done much faster, easier, simpler," though noting that output quality still needs basic creative knowledge. Across all experts, speed gains lead to expanded service capacity and testing options rather than just compressing existing workflows.

Competitive pressure shows up as deep worry about market position. Expert 1 stated this directly: "if you don't upload, don't use AI, you will lag behind as a company, as an agency and alone you won't prevent anything." Expert 3 framed adoption as unavoidable: "those who completely don't want to use it will lag behind… Soon we won't be able to distinguish what is real, what is fake." These statements reflect protective adoption patterns where firms use AI not mainly to improve capabilities but to avoid falling behind as other agencies adopt similar tools. Expert 2 noted that costs across generative platforms are becoming similar, predicting future differences will depend on "convenience of using the actual tool with UI and features" rather than performance or pricing gaps. This suggests early adoption driven by competitive pressure may shift to improving operations as the technology matures and more companies adopt AI.

Strategic innovation reasons are less important than practical efficiency and survival concerns. Expert 1 described

AI as opening up creative work to more people: “Previously you had to rely on the designer’s time, creativity, and now much more people can get involved in the creative process, not just the designer,” presenting adoption as enabling team brainstorming rather than just automating tasks. Expert 2’s early testing in 2021 with Stable Diffusion and Deforum positioned the studio as an innovation leader, noting “we are trying to use as much AI as possible to automate everything,” though this innovative approach exists alongside cost-driven reasons. Expert 3’s dual role as practitioner and educator shows innovation through teaching: training corporate clients on AI tools positions the participant’s services as cutting-edge while creating training revenue. These innovation stories appear secondary to economic needs or help justify cost-focused adoption by calling it strategic.

Data and technological readiness limits appear indirectly through tool selection criteria and skill development challenges. Expert 1 focused on price, quality, and user experience when evaluating platforms, noting that non-technical staff abandon tools “if they don’t succeed in the first few prompts,” requiring “many such trainings in the company team” to build adoption skills. The agency hired a dedicated AI lead to make process automation part of the organization, showing that readiness requires active organizational investment rather than just installing technology. Expert 2 emphasized that effective prompting needs existing creative knowledge: “it’s very important that the designer still would have that skill and craft of knowing all the… colour schemes the composition because you have to have that knowledge in order to prompt well.” Expert 3 agreed, stating “in order to do something high quality, you need to understand: cinematography, how shots connect… Ideas, thoughts, execution, technical execution—very many things connect into one pile.” These accounts suggest technological readiness is not all-or-nothing but exists on a range where basic creative expertise limits effective AI use, and organizations must invest in training to bridge adoption gaps.

Limits emerge around content type and client context. Expert 1 separated functional from luxury product categories, noting that low-cost functional products accept “ugly” AI-generated ads focused on price and delivery speed, while luxury products need lifestyle storytelling where AI reduces production costs for emotional stories. Expert 2 noted consumer acceptance thresholds: when AI outputs are “just terribly hallucinating… it can bring bad feedback,” but execution “with a good taste… can bring good feedback,” citing Coca-Cola’s 2024 holiday campaign as a positive example. Expert 3 described AI adoption as context-dependent: sending casual content “as a joke” differs from producing corporate-acceptable clips, with quality standards changing by use case. These details show adoption goals are not uniform but adjusted to audience expectations, brand positioning, and acceptance of visible AI features, challenging simple adoption models that treat AI integration as a universal efficiency upgrade.

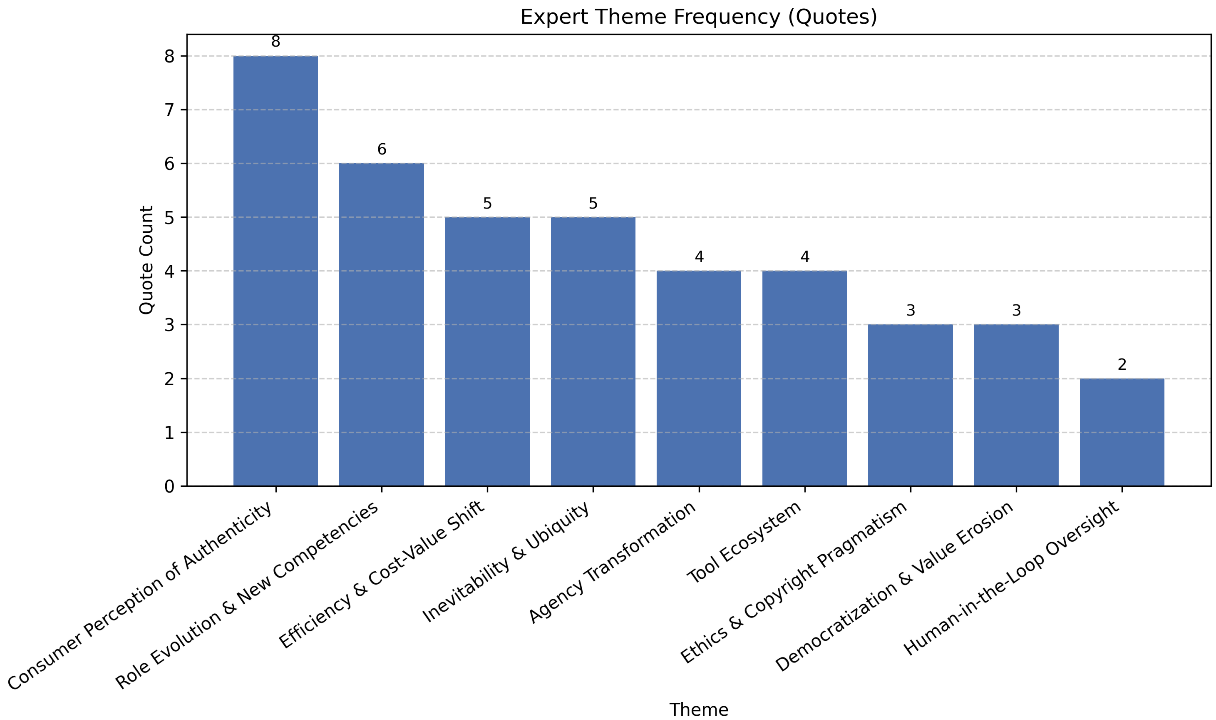

Figure 3. Expert Interview Themes: Adoption Drivers and Workflow Impacts

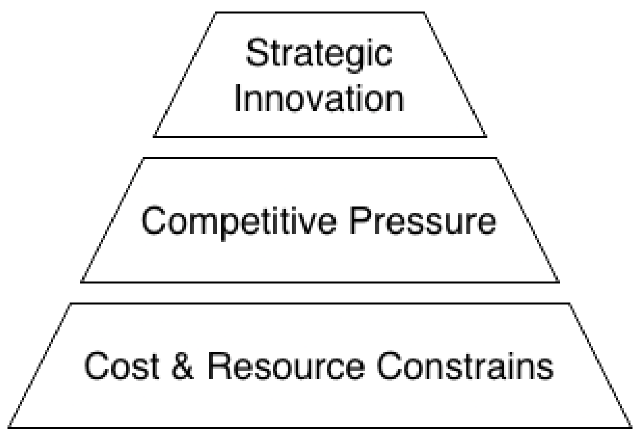

Figure 4. Adoption Drivers Hierarchy: Cost, Competition, Innovation

3.4 Expert Interviews: Workflow Reconfiguration and Outcomes

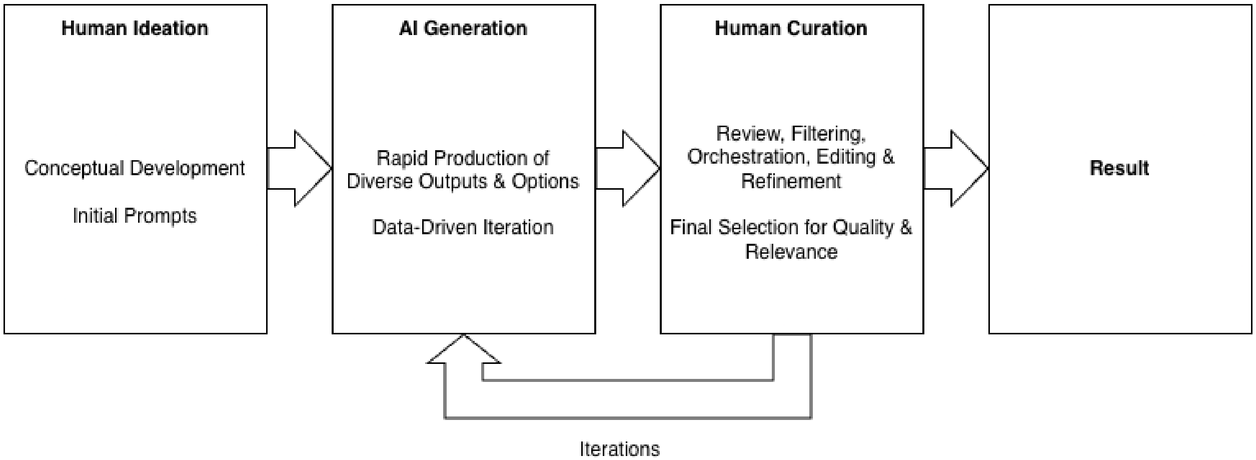

Figure 5. Hybrid Human-AI Workflow

AI integration reshaped workflows across ideation, content production, and quality control stages, with experts reporting consistent patterns in how tasks are distributed between human and AI entities. Changes varied by creative domain but shared common themes around volume expansion, hybrid human-AI execution, and mandatory human oversight checkpoints.

Ideation and Brainstorming. All three experts positioned AI as a front-end ideation accelerator rather than autonomous creative agent. Expert 2 described using ChatGPT for "brainstorming, scripts, and making variations of them definitely for briefings," emphasizing the tool's role in rapid option generation during planning phases. Expert 1 reported that AI "gives much more creativity" by enabling broader team participation: non-designers can now "project a question and get an answer faster," democratizing early-stage brainstorming beyond traditional creative roles. Expert 1 extended this to strategic work, noting AI assists in "preparing offers for clients, when we do marketing strategy offer preparation or the same audit, or when a client comes and asks to put together a half-year plan we personalize it a lot with ChatGPT." This shift suggests AI moved upstream into strategic planning rather than remaining confined to asset execution. Expert 3 noted that ideation value depends on the human's baseline knowledge: "Ideas, thoughts, execution, technical execution—very many things connect into one pile," indicating AI amplifies existing creative capacity rather than replacing it.

Visual Production Workflows. Image and video generation workflows followed a consistent hybrid pattern across all experts: AI generates initial outputs, humans refine to client specifications. Expert 2 described the studio's process: "we generate it with AI and then we tweak it with Photoshop and you know by hand if we're talking about just like images" (Adobe Inc., 1990), noting that pure AI output is "not deliberate" enough for client-specific requirements. The workflow treats AI as a rapid prototyping tool that produces visually appealing starting points requiring manual adjustment for precision. Expert 1 emphasized that the agency "unequivocally do not give out content written only with AI immediately, we write, or make and then a person reviews checks, tweaks and only then gives to the client." Expert 2's studio uses AI image generators for "mood boarding, for character creation inspiration and for actual production of the visuals," spanning early conceptualization through final asset delivery. Expert 3 reinforced that quality output demands understanding "cinematography, how shots connect," suggesting effective AI use requires pre-existing production knowledge to judge output quality and guide refinement.

Testing and Iteration. AI's volume capacity enabled workflow expansion into large-scale variant testing previously limited by production costs. Expert 1 reported creating "twenty, thirty various variations of visuals changed with AI tools, modified slightly" where previously only ten manual variants were feasible, then testing variants through targeted advertising: "you can test again and advertise very directly to that, who is attracted by that specific one visual out of the hundred." This transformed testing from limited A/B comparisons to multivariate optimization at scale. Expert 2 noted the studio can now "do more content… for the same client" within unchanged budgets, implying testing breadth increased without proportional cost increases. Expert 3 described iteration quality as context-dependent, distinguishing casual content shared "as a joke" from "real good clip, which a company/corporation will accept," suggesting testing workflows adapt thresholds based on output destination.

Quality Control and Brand Voice. All experts described mandatory human oversight as non-negotiable despite automation gains. Expert 2 framed this as start-and-end human control: "at the beginning and at the end a human should have its own judgment and input and be responsible for what AI is doing," comparing this to factory automation where humans design parameters and validate outputs even when production is automated. Expert 1 described this as preventing the "super-AI look, when everyone sees that this is AI," which "does not evoke good emotions," requiring human refinement to achieve output that audiences accept as professional rather than visibly artificial. Expert 2 reported consumer reactions depend on execution quality: poorly executed AI content where the tool is "just terribly hallucinating" generates "bad feedback," while "good taste" execution that masks AI origins generates positive responses, citing Coca-Cola's 2024 holiday campaign as successful concealment example. Expert 3 noted that "people don't care" whether content is AI-generated if the result is aesthetically successful, comparing it to plastic surgery: "if done badly—we say cha cha, if done well, we say—what a beautiful person." This suggests brand voice control focuses on output quality rather than process transparency.

Perceived Outcomes. Experts reported divergent speed-quality-cost trade-offs depending on content type and client expectations. Expert 1 described efficiency gains as enabling client capacity expansion: hiring an AI lead for process automation allowed the agency to handle more clients without staff increases. Expert 2 quantified the studio's cost advantage: AI enables production "that looks like it's been made with CG, which takes time and money… with just AI," delivering "bigger value for the client" within compressed budgets. However, quality outcomes appeared more conditional. Expert 3 noted that while AI "speeds up, facilitates and reduces costs for sure" for appropriate content, quality depends on the creator's knowledge: "the majority do nonsense" when lacking foundational skills, and true quality requires understanding how creative elements combine. Expert 1 distinguished functional product categories where "ugly" AI ads focused on price and speed are acceptable from luxury categories requiring lifestyle storytelling where production values remain high even when AI-assisted. Expert 2 observed that consumer tolerance for AI inconsistencies varies by viewing context: "on a phone on a small screen you actually barely see the AI inconsistencies," suggesting mobile consumption lowers quality bars compared to desktop or print contexts. These nuanced accounts indicate workflow outcomes are not uniformly positive but depend on content destination, audience expectations, and creator expertise levels.

3.5 Expert Interviews: Governance, Risk, and Human Oversight

Expert accounts reveal pragmatic, context-dependent governance approaches rather than formal risk management frameworks, with oversight practices shaped by client relationships and brand positioning requirements more than regulatory compliance concerns.

Transparency and Disclosure Practices. Experts distinguished between client-facing transparency and end-consumer disclosure, with different expectations for each. Expert 2 emphasized mandatory honesty with clients about AI use: “we definitely never say to our clients that oh we can make this with CG and then we do it with AI. We definitely transparent and we have to be.” This client transparency serves contractual clarity and manages expectations about production methods. However, end-consumer disclosure practices varied. Expert 2 noted that brand-specific positioning determines disclosure approach: brands like Dove with “natural only products and… models” face pressure to avoid AI given brand identity commitments, citing Dove’s public statement rejecting AI use in marketing. In contrast, Expert 1 reported minimal internal concern about disclosure ethics: “no question marks arose for us, because we are quite an innovative and youthful team,” with management setting a tone that adoption outweighs theoretical ethical concerns. Expert 2 observed that even companies making anti-AI proclamations may reverse positions, noting Warner Bros initially sued music generator Suno but “now apparently they are collaborating,” suggesting disclosure norms remain unstable as “companies definitely change because they see that AI is inevitable in every step of business.”

IP and Copyright Management. Experts acknowledged copyright challenges but lacked clear resolution strategies. Expert 2 emphasized human responsibility for dataset selection: “the IP like the dataset that you’re using, you know, a human being has to choose that,” positioning compliance as upstream decision-making rather than output monitoring. The participant distinguished consent-based uses from prohibited ones: “deep fakes without consent of the person of course shouldn’t in generally happen… you always should have an agreement,” though acknowledging caricature exceptions where intent is obvious. Expert 3 described copyright as unresolved: “About copyright I don’t even want to talk, here is a long topic. If briefly — we see the problem, but we don’t know, how to solve it.” This admission suggests experts recognize IP risks but operate without clear mitigation protocols. Expert 1 reported selective data protection: “we don’t put all contracts into tools, but as far as photos are concerned, it is uploaded,” indicating pragmatic risk-benefit calculations where competitive advantage from AI use outweighs potential IP exposure for non-sensitive materials.

Approval Gates and Quality Control. All experts described mandatory human checkpoints at workflow start and end points. Expert 2 framed this as non-negotiable oversight: “at the beginning and at the end a human should have its own judgment and input and be responsible for what AI is doing,” comparing this to factory automation where “a human being sets up the design… and only the production is being given to the robots and AI” with final validation by humans. Expert 1 reinforced this pattern: “A human is needed and their decisions remain necessary before going on air or before going to the client,” positioning AI as middle-process automation rather than autonomous decision-maker. Expert 2 noted organizational accountability drives oversight: “if a client hires us we are responsible for what we bring to them,” making human validation essential regardless of AI’s production role. No experts mentioned automated bias checks, hallucination monitoring, or systematic quality audits beyond human review, suggesting governance relies on expert judgment rather than structured protocols.

Human–AI Division of Labor. Experts described emerging role specialization around AI tool management. Expert 2 identified “prompt engineer” as a new profession requiring model knowledge and prompting skills, though emphasizing this role still demands “creative direction mindset” and understanding of “professional outcome” standards rather than purely technical skill. Expert 1 institutionalized this through dedicated hiring: the agency added “a person who works purely with AI process automation… like a kind of AI lead,” signaling organizational recognition that AI adoption requires specialized coordination roles. Expert 3 offered a contrasting view: “generally no team is needed, if you know things,” noting that skilled individuals can now accomplish tasks requiring entire teams previously, with clients moving to solo practitioners “because normally cheaper.” This suggests labor division varies by organizational scale: larger agencies adding AI coordination roles while individual practitioners consolidate multiple functions. Expert 2 emphasized that effective AI use still requires baseline creative expertise: the designer must know “colour schemes the composition because you have to have that knowledge in order to prompt well,” indicating AI amplifies rather than replaces domain knowledge in the division of labor.

3.6 Expert Interviews: Interim Summary

Experts adopted AI mainly because of budget limits and competitive pressure. All three described flat or declining client budgets as the main reason, with worry about falling behind other agencies pushing adoption. Workflows followed similar patterns: AI generates first versions, humans adjust them to client needs, with people making decisions at the start and end of each project. No expert mentioned formal risk rules; instead, they made case-by-case choices based on client relationships and brand needs. AI lets teams create more content variants for testing than was possible before, changing how campaigns work rather than just making existing work faster. Quality results depended on context: experts said functional products accept rougher AI output while luxury products need higher production values, and mobile screens hide AI problems that show up on desktops. Larger agencies hired dedicated AI staff while solo practitioners took on tasks that used to need whole teams. In all cases, AI worked better when people already had creative skills—it made existing knowledge stronger rather than replacing it, with quality depending on how well creators understood composition, camera work, and professional standards.

3.7 Audience Intercepts: Sample and Stimuli Check

The audience sample (n=30) spanned ages 10 to 71 across six age groups, with participants from twelve countries, including Sweden (n=17), Somalia (n=2), and the USA (n=2). Gender distribution showed 16 females and 14 males. All viewed the AI-generated commercial first, then the human-created commercial, stating preferences before and after disclosure of production methods. Participants readily distinguished between commercials and commented on production qualities, visual style, and tone, confirming they processed both stimuli as intended.

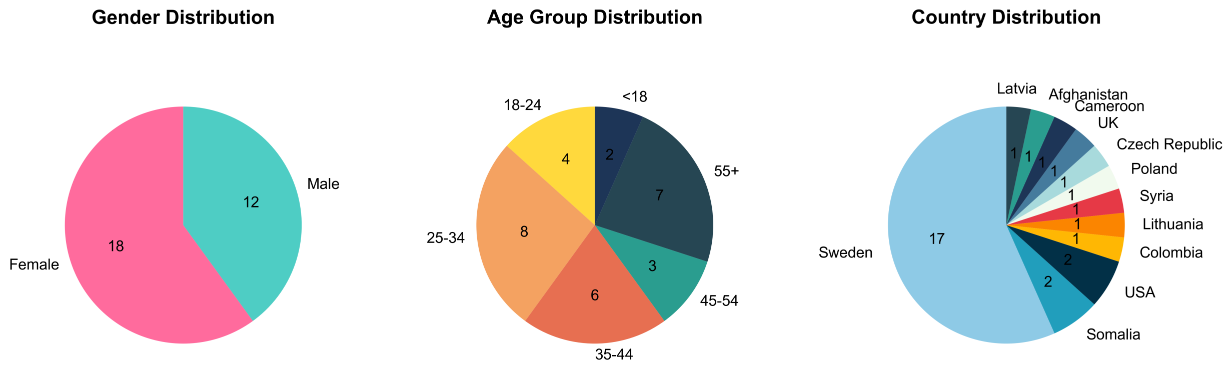

Figure 6. Audience Demographics by Gender, Age, and Origin

3.8 Audience Intercepts: Preference Shifts (Quantitative)

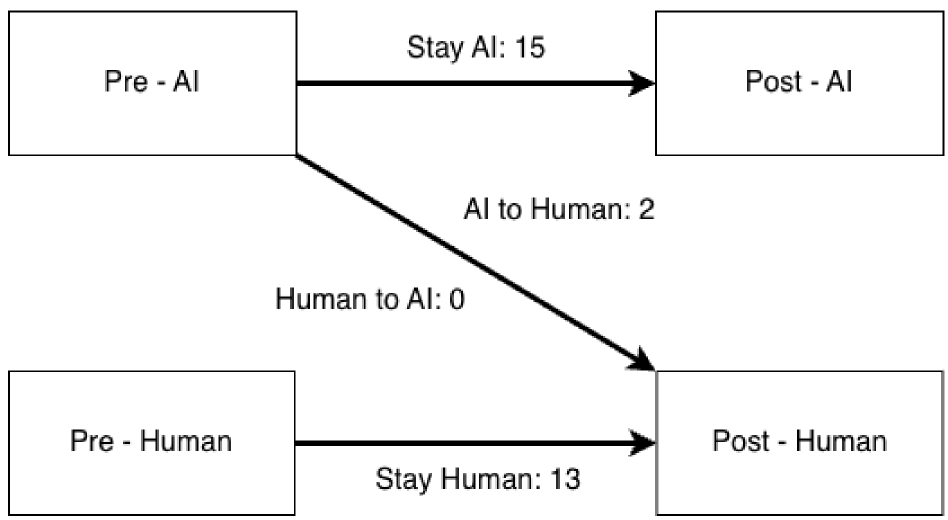

Figure 7. Audience Preference Shifts: Pre- and Post-Disclosure Flow

Before disclosure of production methods, participants split nearly evenly between the AI-generated commercial (n=17, 56.7%) and the human-created commercial (n=13, 43.3%). This initial distribution shows the AI-generated commercial competed effectively with the human-created one when evaluated on perceived quality alone, without viewers knowing which production method created which ad. The slight AI advantage in blind evaluation contradicts assumptions that AI-generated advertising is immediately recognizable as inferior or off-putting. After learning which commercial was AI-generated and which was conventionally filmed, the distribution shifted slightly: 15 participants (50%) preferred the AI commercial and 15 (50%) preferred the human commercial. Overall, 28 participants (93.3%) maintained their initial preference regardless of disclosure, while 2 participants (6.7%) shifted away from the AI commercial after learning it was AI-generated. No participants shifted toward the AI commercial following disclosure.

Preference Stability and Shift Patterns. The dominant pattern was preference stability. Among the 17 participants who initially preferred the AI commercial, 15 (88.2%) maintained that preference after disclosure, while 2 (11.8%)

switched to the human commercial. All 13 participants who initially preferred the human commercial maintained their preference (100% stability). The two participants who shifted away from AI represented the only directional change observed, reducing AI preference from 56.7% to 50% of the sample.

McNemar Test Results. The paired-sample McNemar test examined whether disclosure caused significant preference shifts. The test compares discordant pairs—participants who changed preferences—against the null hypothesis of no systematic change. With 2 participants shifting from AI to human and 0 shifting from human to AI, the test yielded $\chi^2(1) = 0.5$ with continuity correction, $p = 0.48$. The exact binomial test gave $p = 0.50$. Both p-values exceed conventional significance thresholds ($p < 0.05$), indicating disclosure did not produce statistically detectable preference changes in this sample.

Table 2. Audience Preference Patterns

Preference Pattern	Count	Percentage
Stayed with AI	15	50.0%
Stayed with Human	13	43.3%
Shifted from AI to Human	2	6.7%
Shifted from Human to AI	0	0.0%
Total	**30**	**100%**

Interpretation of Observed Patterns. The 11.8% shift rate among initial AI supporters (2 out of 17 switching to the human commercial) suggests a directional negative disclosure effect. If this pattern held at larger sample sizes, it would support rejecting the null hypothesis that disclosure has no impact, indicating that revealing AI production methods causes modest but measurable preference shifts away from AI-generated advertising. The directional asymmetry reinforces this interpretation: all movement went away from AI after disclosure, with zero shifts toward AI, creating a 2-to-0 split among those who changed preferences. If scaled proportionally to larger samples, such consistent one-way movement would reach statistical significance (Cohen, 1988).

However, the more likely interpretation given the dominant stability pattern (93.3% no change) is that most viewers base preferences on perceived quality, style, and emotional resonance rather than production method. The near-even post-disclosure split (50-50) and high retention rates (88.2% for AI supporters, 100% for human supporters) suggest disclosure functions as a secondary factor that occasionally tips marginal preferences but rarely overrides initial quality-based reactions. The absence of any shifts toward AI after disclosure is notable: knowing an ad was AI-generated never improved its standing, only maintained or reduced it. This asymmetry aligns with expert predictions that disclosure can hurt AI ads but never helps them—some viewers reacted negatively when they learned an ad was AI-generated, but no viewers reacted more positively.

Statistical Power Limitations. These interpretations remain tentative due to sample size constraints. The paired-sample McNemar test could not reject the null hypothesis ($p = 0.48$), but this does not prove disclosure has no effect. With only 2 discordant pairs in $n=30$, the test lacks statistical power to detect modest effects even if they exist (Cohen, 1988). The observed 11.8% shift rate among initial AI supporters (2 of 17 switching away) falls below the threshold detectable at this sample size, meaning we cannot rule out that a similar negative disclosure effect exists in the broader population but remains invisible in this exploratory sample. The results do suggest disclosure does not cause large-scale preference reversals—if half the sample or more changed their minds, the test would detect it—but subtle effects remain undetectable. These limitations frame this finding as a preliminary pattern requiring validation in a larger sample rather than a definitive conclusion about disclosure effects.

Practical Implications. The dominant stability pattern shows that viewers base their preferences on execution quality and perceived appeal rather than production method. Brands should focus on output refinement rather than production-method messaging, since 93.3% of viewers maintained their initial judgments regardless of disclosure. The 88.2% retention rate among initial AI supporters demonstrates that well-executed AI advertising sustains approval even when its origins are revealed, contradicting fears of wholesale audience rejection upon disclosure. The findings indicate authenticity concerns may be overstated—production method does not dominate viewer judgment in the way disclosure debates often assume.

3.9 Audience Intercepts: Explanation Patterns (Qualitative)

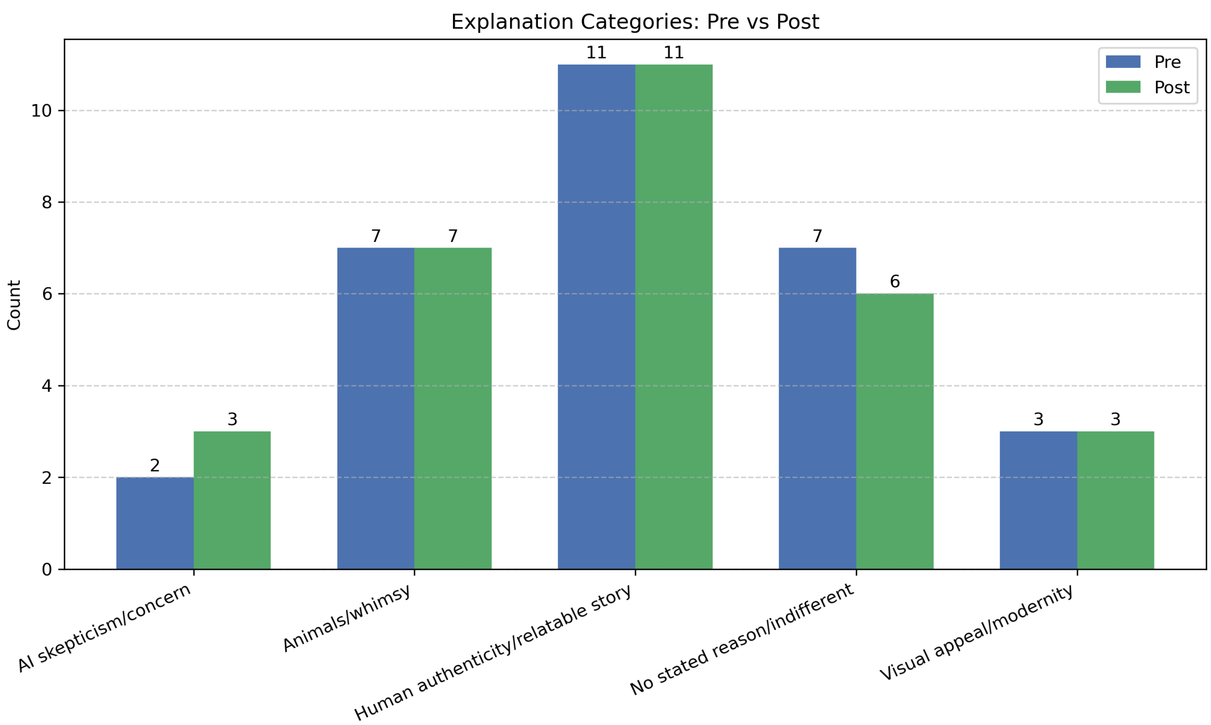

Figure 8. Audience Explanation Categories: Reasoning Patterns

When asked to explain their preferences, participants invoked distinct reasons that clustered into five categories: authenticity and relatability (the dominant theme, accounting for 11 pre-disclosure and 10 post-disclosure responses), visual appeal and whimsy (10 participants across both conditions), indifference or unexplained preference (7 participants), AI skepticism (2 pre-disclosure, 3 post-disclosure), and visual modernity (3 participants). These patterns reveal what cues viewers attend to when evaluating advertising and whether production method disclosure reshapes those priorities.

Table 3. Audience Explanation Categories

Category	Description	Count (Pre)	Count (Post)
Authenticity	Relatability, ”real feel”	11	10
Visual Appeal	Aesthetics, ”cute animals”	10	10
Indifference	No strong reason	7	7
AI Skepticism	Principled objection	2	3
Total		**30**	**30**

Authenticity and Relatability. The most common explanation pattern centered on perceived authenticity, familiar scenarios, and emotional resonance. Eleven participants initially cited these reasons, with 10 maintaining this framing after disclosure. Participants who preferred the human-created commercial described it as invoking “memories” and “familiar family gathering,” offering “Christmas vibe and more people,” and presenting a “family and Christmas spirit [that] felt more relatable.” One participant explained that “people drinking cola…made more sense” than animals doing so, while another found the human commercial “more real/relatable [and] felt like a normal commercial.” Nostalgia appeared frequently, with multiple participants selecting the human ad because it reminded them of past experiences or felt “old school” in a positive sense. Two participants who initially preferred the AI commercial for its nostalgic or timeless qualities switched to the human ad after disclosure, explicitly stating they “wanted human-made and more genuine” or “preferred the filmed, more real feel.” This shift suggests authenticity concerns can override initial quality judgments for some viewers, but the small number (2 of 11) indicates this reaction is not universal.

Visual Appeal and Whimsy. Ten participants based their preferences on visual qualities unrelated to authenticity. Seven chose the AI-generated commercial because they “liked animals,” found it “cute,” appreciated the “lighter landscape and animals,” or described it as “more fantasy/fairytale.” None of these participants changed their preference after disclosure, suggesting production method had no bearing on their enjoyment of the visual content. Three additional participants preferred the AI commercial for its “newer/modern” feel versus the “old fashioned” human ad or cited “brighter colors” and superior visual execution. These responses reflect judgments about aesthetic style and production polish rather than concerns about origin or authenticity.

Indifference. Seven participants offered no substantive explanation for their preference, either stating a simple choice such as “number 1 maybe” or declining to elaborate when prompted. None of these participants switched after disclosure, and their post-disclosure responses remained equally terse (“For me - no”; “Not for me”). This pattern suggests a subset of viewers make rapid, intuitive judgments without consciously processing specific features or caring about production backstory.

AI Skepticism and Disclosure Impact. Only two participants expressed concerns about AI before disclosure. One stated they “dislike AI taking creative jobs” and found the human “family story more relatable,” while another described the AI commercial as “not appealing” compared to the “old school” human version. After disclosure, one additional participant joined this group, switching from the AI commercial because they “wanted human-made and more genuine.” The increase from 2 to 3 participants citing AI skepticism aligns with the quantitative finding that disclosure caused minimal preference shifts—most viewers either did not care about production method or had already incorporated such concerns into their initial judgment.

Patterns Across Disclosure. The stability of explanation categories before and after disclosure reinforces the quantitative finding that viewers base preferences on perceived quality and appeal. Only one category changed meaningfully: AI skepticism grew by one participant. The dominant authenticity theme lost one adherent (11→10) as a participant moved to the skepticism category, but this represented a reframing of the same underlying concern rather than a wholesale shift in priorities. No participants who initially cited visual appeal, whimsy, or indifference changed their reasoning after learning production methods, indicating disclosure does not alter how these viewers process advertising stimuli.

3.10 Audience Intercepts: Demographic Cross-Tabs

Demographic breakdowns reveal one notable age-related pattern but offer no reliable insights about origin due to severe sample-size constraints. These findings should be treated as exploratory observations requiring validation in larger samples rather than stable demographic effects.

Age Patterns. Both participants who shifted away from the AI commercial after disclosure fell within the 18-34 age range: one from the 18-24 group (n=4, yielding a 25% shift rate within that bracket) and one from the 25-34 group (n=8, yielding a 12.5% shift rate). In contrast, participants aged 35 and older showed zero preference changes regardless of disclosure. The under-18 group (n=2) and the 35-44 (n=6), 45-54 (n=3), and 55+ (n=7) groups all maintained their initial preferences. This pattern suggests younger viewers may be slightly more responsive to production-method information, potentially reflecting greater exposure to AI-related discourse or less entrenched viewing habits. However, the absolute numbers remain small—two individuals shifting does not establish a reliable age effect, and the 18-24 bracket in particular contains too few participants to support strong claims. The pattern aligns with the possibility that younger audiences weigh production authenticity differently than older cohorts, but this remains speculative without replication in a larger sample stratified by age.

Origin Patterns and Small-Cell Constraints. The sample's origin distribution prevents meaningful analysis. Swedish participants dominated the sample (n=17, 56.7% of total) and showed zero shifts, maintaining perfect preference stability after disclosure. The remaining 13 participants came from 10 different countries, with most countries represented by a single individual: Afghanistan, Cameroon, Colombia, Czech Republic, Latvia, Lithuania, Poland, and Syria each contributed one participant who did not shift, while the USA contributed two stable participants. The two participants who shifted came from Somalia (n=2, one shifted, one stable) and the UK (n=1, shifted). Drawing any conclusions about origin-based disclosure effects from these distributions would be methodologically unsound. The single-participant cells offer no statistical power, and even the Swedish subsample (n=17) cannot be compared to any comparably sized group. Whether Swedish participants' stability reflects cultural factors, selection effects, or random variation cannot be determined from this data.

Interpretation Limits. The demographic breakdowns confirm the overall preference stability pattern across all groups but add limited explanatory value. The age-related observation that both shifts occurred among younger participants warrants further investigation but does not alter the core finding that disclosure caused minimal overall change. The origin data's severe imbalance forecloses demographic comparison and underscores the need for purposive sampling designs if future research aims to test cultural or nationality-based hypotheses about AI disclosure effects.

3.11 Audience Intercepts: Interim Summary

The audience intercepts establish three core findings relevant to the research questions. First, AI-generated advertising competed effectively with human-created content when evaluated on perceived quality alone: 56.7% of participants preferred the AI commercial before learning production methods, demonstrating that well-executed generative content can match or exceed traditional advertising in initial audience appeal. Second, disclosure of production methods caused minimal preference shifts—only 2 of 30 participants (6.7%) changed their choice after learning which ad was AI-generated, and both shifts moved away from AI rather than toward it. The dominant pattern (93.3% stability) indicates viewers base preferences on execution quality, emotional resonance, and visual appeal rather than production method. Third, qualitative explanations reinforce this finding: participants cited authenticity, relatability, visual whimsy, or offered no substantive reasoning, with explanation categories remaining stable across disclosure. Only 3 of 30 participants expressed AI-related concerns, and disclosure added just one person to that group. Demographic patterns remain tentative due to sample size, though both shifts occurred among viewers aged 18-34. Together, these findings suggest that the production method operates as a secondary consideration that rarely overrides quality-based judgments, challenging assumptions that audiences uniformly reject AI-generated advertising upon disclosure.

3.12 Triangulation: Expert Predictions vs Audience Responses

Comparing expert predictions with observed audience responses reveals both accurate anticipations and notable misalignments, offering insight into how practitioners' operational assumptions align with viewer behavior.

Supported: Quality Execution Determines Acceptance. Experts predicted that execution quality would drive audience reactions more than production method. Expert 2 stated that poorly executed AI content "just terribly hallucinating" generates "bad feedback" while "good taste" execution generates positive responses, citing Coca-Cola's 2024 holiday campaign as a successful example. Expert 3 reinforced this view, noting "people don't care" whether content is AI-generated if the result succeeds aesthetically. The audience data strongly supports this prediction: 93.3% of participants maintained their preferences after disclosure, with explanation patterns dominated by quality-based reasoning (authenticity, visual appeal, emotional resonance) rather than production-method concerns. Only 3 of 30 participants expressed AI-related concerns, confirming experts' operational assumption that well-executed AI advertising performs competitively when quality matches human-created alternatives.

Supported: Disclosure Can Hurt, But Never Helps. Experts anticipated that revealing AI production methods carries downside risk without upside potential. This prediction appeared implicitly in disclosure strategies that prioritize client transparency while remaining cautious about consumer disclosure, particularly for brands with authenticity positioning. The audience data confirm this asymmetry: both preference shifts moved away from AI after disclosure (2 of 17 initial AI supporters switched to human), while zero participants shifted toward AI upon learning production methods. The directional pattern validates experts' risk assessment that disclosure may tip marginal preferences against AI advertising but offers no competitive advantage, supporting their operational preference for producing content indistinguishable from human-created work rather than emphasizing AI involvement.

Partially Supported: Product Category and Context Dependencies. Expert 1 distinguished functional product categories that accept "ugly" AI ads focused on price from luxury categories requiring lifestyle storytelling with high production values. Expert 2 noted that viewing context matters: "on a phone on a small screen you actually barely see the AI inconsistencies." The current audience study cannot fully validate these predictions because the Coca-Cola stimulus represents a single brand (low-involvement hedonic product (Belanche et al., 2025)) viewed on laptop screens in controlled settings. However, the finding that visual appeal and whimsy drove 10 participants' preferences regardless of disclosure suggests that viewing context and product category may indeed moderate reactions, as experts predicted. The lack of differentiation by product type in this study reflects methodological constraints rather than evidence against the experts' claims, leaving this prediction tentatively supported but requiring validation across varied categories and viewing conditions.

Partially Supported: Audience Skepticism and Disclosure Risk. Experts accurately predicted that AI skepticism exists among some viewers and that disclosure carries downside risk. The audience data confirm both predictions: 10% of participants (3 of 30) expressed AI-related concerns, including one with a marketing background who explicitly spoke about his job displacement worries, and 6.7% (2 of 30) shifted away from AI after disclosure while none shifted toward it. This validates experts' core insight that quality execution matters most but some viewers do care about production methods. However, the level of caution experts show—avoiding consumer disclosure, emphasizing brands like Dove proclaiming they will not use AI, framing disclosure primarily as risk management—suggests they may treat this concern as more common than the observed 10% rate. Whether their caution is appropriate risk management for a 10% concern rate or excessive relative to it remains unclear. Their exposure to high-profile negative reactions to specific AI campaigns may make the risk feel larger than the general audience data indicate, though the 93.3% preference stability shows most viewers base their judgments on quality rather than production method.

The triangulation reveals strong alignment between expert predictions and audience responses on core quality-driven reception patterns, with partial alignment on skepticism prevalence and caution calibration. Experts accurately identified that quality execution determines acceptance, that AI skepticism exists among some viewers, and that disclosure carries asymmetric risk—all confirmed by audience data showing 93.3% preference stability, 10% expressing AI concerns, and 6.7% shifting away from AI with none shifting toward it. Their operational strategies—investing in production refinement, maintaining human oversight, avoiding obvious AI artifacts—directly match what drives viewer preferences. The product-category distinctions experts make remain plausible but require empirical validation across contexts beyond the current study's scope, which examined a single brand type in controlled viewing conditions. Expert predictions prove most accurate when describing creative execution standards and quality drivers.

3.13 Integration with Conceptual Model

The empirical findings confirm the conceptual model's core architecture while refining the relative importance of drivers, adding nuance to mechanism descriptions, and revealing outcome measurement challenges that the theory underspecified.

Drivers: Confirmed Hierarchy with Cost Dominance. The model proposed five adoption drivers without ranking them. Expert interviews establish a clear hierarchy: cost and resource constraints dominate all three cases, with competitive parity concerns appearing as defensive reinforcement rather than independent motivation. Expert 2's statement that client budgets "have not changed and to be fair in some cases they've just dropped" frames AI adoption as cost-driven strategic response. Strategic innovation pressures appear secondary—mentioned by experts but framed as justification for cost-driven adoption rather than primary trigger. This finding refines the model by indicating that drivers operate in nested layers: economic pressures create urgency, competitive anxiety accelerates timeline, and innovation framing legitimizes decisions already made for financial reasons. Data and technological readiness emerged as barrier rather than driver—experts described skill-building challenges and tool abandonment when non-technical staff "don't succeed in the first few prompts," suggesting readiness moderates implementation success rather than enabling initial adoption decisions.

Mechanisms: Human Oversight More Intensive Than Anticipated. The model proposed four mechanisms with human-AI collaboration and brand-safety guardrails as distinct categories. Empirical evidence shows these mechanisms often merge in practice: every expert emphasized mandatory human checkpoints "at the beginning and at the end" that simultaneously handle task allocation and quality control, with Expert 2 explicitly rejecting middle-process automation without validation. This integration pattern confirms intensive human involvement throughout production cycles, refining the theoretical model by showing how distinct conceptual categories operate together in implementation. End-to-end integration manifested as experts described but with persistent quality control requirements—Expert 2 noted AI outputs require refinement to avoid the "super-AI look, when everyone sees that this is AI," indicating integration extends workflow duration rather than compressing it when quality standards remain high. Feedback and iteration loops appeared informal rather than systematic: no expert mentioned structured performance monitoring or algorithmic bias audits beyond subjective human judgment, suggesting governance relies on practitioner expertise rather than institutionalized protocols the theoretical literature emphasizes.

Moderating Factors: Product Category Confirmed, Team Structure Shows Scale-Dependent Patterns. The model proposed that campaign type (hedonic versus utilitarian) and team structure moderate outcomes. Expert 1's distinction between functional products accepting "ugly" AI ads and luxury products requiring lifestyle storytelling confirms the campaign-type moderator with concrete examples. Team structure effects expand the theoretical model by revealing scale-dependent patterns: Expert 1 described larger agencies institutionalizing AI through dedicated coordination roles (hiring an "AI lead"), while Expert 3 showed solo practitioners consolidating functions previously requiring teams, with clients preferring individuals "because normally cheaper." This finding refines the model by demonstrating that organizational scale creates divergent pathways—larger agencies adding specialized roles while smaller operations enable individual practitioners to replace team structures—rather than team structure operating uniformly across all organization types.

Outcomes: Consumer Response Measured, Disclosure Effects Specified. The model specified six outcome dimensions including consumer response and perceived authenticity. Audience intercepts provide direct measurement previously absent from theoretical speculation: 93.3% preference stability demonstrates that consumer response to well-executed AI advertising matches human-created equivalents when quality remains high, validating the model's prediction that execution quality moderates reception. Disclosure effects reveal patterns the model left unspecified: they proved asymmetric (can hurt, never helps) but minimal in magnitude (only 2 of 30 shifted), suggesting disclosure operates as threshold concern—relevant when flagged but not salient unprompted—rather than continuous dimension requiring ongoing management. This finding refines outcome measurement by indicating that some dimensions (time-to-market, cost efficiency) demand continuous monitoring while others (disclosure effects) require attention only at decision points where production methods are revealed.

Tensions and Paradoxes: Validated with Operational Examples. The model identified operational paradoxes where AI simultaneously speeds production and adds oversight work. Expert accounts validate this tension through concrete workflow descriptions: AI enables rapid variant generation (Expert 1's "twenty, thirty various variations") while requiring human refinement at every stage, with Expert 3 noting that quality output still demands "cinematography, how shots connect… Ideas, thoughts, execution" expertise. Expert 2's factory automation analogy illustrates the division of labor: "a human being sets up the design… and only the production is being given to the robots and AI" with final validation by humans. These operational examples confirm the paradox exists in practice, with experts describing workflows that incorporate both rapid AI-enabled production and intensive human oversight requirements.

The integration reveals the conceptual model's predictive validity while refining driver hierarchies, mechanism intensiveness, moderator pathways, and outcome measurement approaches based on practitioner experience and observed audience behavior.

3.14 Practical Implications for Creative Operations

Workflow Design. Expert testimony (n=3) confirms AI accelerates content production; Expert 1 noted that "the coming of AI into our industry gives much more creativity, gives much more speed to deliver things… You can have not ten manually created visuals anymore, but twenty, thirty various variations of visuals changed with AI tools," while creating new quality control demands. All three experts emphasized mandatory human oversight "at the beginning and at the end" to refine outputs and avoid the "super-AI look" audiences reject. This pattern suggests organizations should plan for two workflow effects simultaneously: increased production capacity (variant generation, content volume) alongside persistent quality control requirements. Context determines emphasis: functional product campaigns may prioritize speed gains and accept lower refinement thresholds, while luxury and lifestyle categories require intensive human oversight to meet brand standards regardless of AI's production advantages. Organizations should allocate human expertise at workflow checkpoints rather than expecting automation to reduce oversight needs.

Quality Control Priorities. Experts demonstrate quality management through output assessment rather than process monitoring—Expert 2 noted poorly executed AI content "just terribly hallucinating" generates "bad feedback" while "good taste" execution generates positive responses. The finding that 93.3% of audience participants (n=30) maintained preferences when execution quality was high validates this output-focused approach. Organizations should prioritize visible quality standards (avoiding AI artifacts, maintaining brand voice consistency, ensuring output precision) over procedural documentation. Human expertise at validation checkpoints drives quality outcomes: Expert 3 noted quality depends on the creator's understanding of "cinematography, how shots connect," indicating effective oversight requires domain knowledge. Expert 2 observed viewing context affects quality thresholds—mobile screens mask AI inconsistencies better than desktop viewing—suggesting quality standards should adapt to content destination.

Disclosure Considerations. Audience intercept data (n=30) show 93.3% preference stability after production method disclosure, indicating well-executed AI content sustains approval even when origins are revealed. However, disclosure effects proved asymmetric: 6.7% of participants shifted away from AI-generated content after disclosure while none shifted toward it, and one participant expressed AI-related concerns including job displacement worries. These patterns suggest disclosure carries modest downside risk without competitive advantage. Brand positioning affects disclosure sensitivity: brands emphasizing authenticity or natural attributes face higher risk, as evidenced by Expert 2's observation that brands like Dove proclaimed they would not use AI given their positioning. The 10% concern rate indicates most viewers prioritize execution quality over production method, though organizations should recognize that a minority holds principled objections to AI-generated content regardless of output quality.

Capability Development. Expert accounts (n=3) indicate AI adoption creates skill-building demands rather than reducing expertise requirements. Expert 1 reported non-technical staff abandon tools when they "don't succeed in the first few prompts," requiring "many such trainings" and dedicated AI coordination roles to sustain adoption. Quality output demands baseline creative knowledge: Expert 2 emphasized designers must understand "colour schemes the composition because you have to have that knowledge in order to prompt well." This pattern means AI amplifies existing expertise rather than compensating for skill gaps. Organizations should invest in training that combines creative fundamentals with AI tool proficiency, recognizing that untrained users produce what Expert 3 described as "nonsense" despite tool access. Organizational scale creates different capability pathways: larger agencies may add specialized AI coordination roles, while smaller operations enable skilled solo practitioners to consolidate functions previously requiring teams.

3.15 Section Summary

Findings show practitioners adopt generative AI under budget pressure and competitive urgency, using it to expand variant generation and testing at speed while keeping humans at the start and end of workflows to prevent the visible "super-AI look" and align output to brand standards. Governance remains pragmatic: disclosure decisions depend on brand positioning, IP safeguards are case-by-case, and quality gates rely on expert judgment rather than formal protocols. Audience intercepts (n=30) found the AI Coca-Cola ad performed on par with the human version pre-disclosure and retained 93.3% preference stability after disclosure, with only 6.7% shifting away from AI and a small minority voicing principled concerns. Integration with the conceptual model validates cost/speed drivers, hybrid mechanism pathways, and moderators such as category positioning, screen context, and disclosure framing. Key limitations—small expert (n=3) and audience (n=30) samples in one city, fixed presentation order, single brand and format. The evidence base supports practice guidance on workflow checkpoints, quality standards, cautious disclosure, and capability investment, setting up the conclusions and recommendations.

CONCLUSIONS AND RECOMMENDATIONS

Key Conclusions

This research examined how generative artificial intelligence reshapes creativity automation within marketing communication, with implications for DeepTech entrepreneurs and established marketing practitioners. The following conclusions answer the research aim and objectives:

1. **Adoption drivers are cost-driven and competitive.** Entrepreneurs and marketing practitioners adopt generative AI primarily to address budget constraints and competitive pressure. Expert interviews (n=3) confirm that flat client budgets and rivals deploying similar tools force adoption as a survival response. These external pressures transform AI adoption from an experimental option into an operational baseline, making the question not whether to adopt but how quickly and strategically to integrate AI into existing workflows.
2. **Implementation follows hybrid human-AI workflows.** Practitioners position AI as augmentation rather than replacement, keeping human oversight at ideation start and output approval end. This hybrid model prevents the visible "super-AI look" that audiences reject while capturing efficiency gains. The conceptual model's mechanism pathway—where task allocation and quality gates moderate outcomes— accurately describes observed practice.
3. **Efficiency gains coexist with homogenization risks.** Practitioners report significant production-task efficiency improvements and time savings, enabling multi-fold increases in variant generation, but caution that

unchecked automation produces generic outputs that weaken brand distinctiveness when competitors use identical tools. The automation-authenticity tension identified in theoretical synthesis manifests directly in workflow decisions.

4. **Audience disclosure effects are minimal in low-involvement contexts.** Audience intercepts (n=30) comparing AI-generated and human-created Coca-Cola holiday commercials found 93.3% preference stability after disclosure, with only 6.7% shifting away from AI content. Performance parity pre-disclosure and limited post-disclosure penalties suggest that execution quality and brand fit matter more than production method for low-involvement emotional advertising.
5. **Governance remains pragmatic and judgment-based.** Practitioners implement case-by-case IP checks, selective disclosure tied to brand positioning, and expert judgment quality gates rather than formal protocols. This pragmatic approach enables speed but lacks scalability as AI use intensifies.

Practical Recommendations for DeepTech Entrepreneurs and Marketing Practitioners

Based on expert and audience evidence, the following recommendations provide actionable guidance for integrating generative AI into creative marketing workflows:

Maintain human oversight at workflow boundaries. Position AI as augmentation rather than replacement by keeping practitioners at ideation start (defining brief, tone, constraints) and output approval end (brand alignment, quality verification). This hybrid approach captures the significant efficiency gains experts report while preventing the “super-AI look” that audiences reject. Automated middle stages (variant generation, asset iteration) deliver speed without sacrificing distinctiveness.

Implement quality gates before publication. Establish clear checkpoints where expert judgment evaluates whether AI outputs align with brand voice, comply with IP constraints, and meet category expectations. Practitioners emphasize that unchecked automation produces generic content that weakens competitive positioning. Formal review protocols prevent homogenization and maintain creative differentiation even when competitors use identical tools.

Adopt context-sensitive disclosure strategies. Audience evidence (93.3% preference stability, n=30) suggests that disclosed AI authorship imposes minimal penalties in low-involvement emotional advertising when execution quality is high. Focus disclosure efforts where regulatory mandates require it or where brand transparency positioning demands it, rather than applying blanket policies. Test disclosure formats and timing within your specific product category and audience context.

Prioritize capability building over tool acquisition. Expert interviews reveal that prompt clarity, iterative refinement skill, and AI-human workflow coordination determine outcomes more than tool selection. Invest in training practitioners to write effective prompts, recognize when AI outputs drift from brand standards, and maintain judgment authority over generated content. Technical access alone does not guarantee a strategic advantage.

Formalize governance as adoption scales. Case-by-case IP checks and informal quality reviews work early but lack scalability. As AI use intensifies across campaigns, develop documented protocols for rights clearance, bias audits, disclosure requirements, and approval workflows. Pragmatic early governance must evolve into systematic oversight to manage legal, ethical, and reputational risks at scale.

Research and Implementation Roadmap

This thesis establishes a baseline for understanding generative AI's impact on marketing, but several areas require further investigation to validate findings across broader contexts.

1. DeepTech Pilot Implementation. Future work should move beyond audience intercepts to live pilot implementations within DeepTech ventures. A controlled experiment comparing human-only versus AI-augmented workflows in a real campaign would quantify efficiency gains (validating the reported 3x increase in production capacity) and measure conversion rates rather than just preference. This would directly test the "cost-driven" adoption driver in a high-stakes environment.

2. Cross-Category Validation. The current study focused on low-involvement emotional advertising (Coca-Cola). Future research must test the "functional vs. luxury" distinction identified by experts. A comparative study analyzing consumer reactions to AI-generated content for high-stakes products (e.g., real estate, financial services) versus functional goods would clarify where the "authenticity penalty" is most severe.

3. Longitudinal Trust Studies. As AI content saturates the market, consumer skepticism may evolve. A longitudinal study tracking sentiment over 12-24 months would reveal whether the "minimal disclosure effect" holds as audiences become more AI-literate, or if a "trust deficit" accumulates over time.

4. Scalable Governance Frameworks. Current governance relies heavily on manual expert judgment and case-by-case decisions, which may become a bottleneck as production volumes increase. Research should explore how to formalize these ad-hoc practices into scalable frameworks—such as standardized disclosure protocols or automated quality checks—to maintain safety and consistency without slowing down the expanded production capacity AI enables.

BIBLIOGRAPHY AND A LIST OF REFERENCES

Adobe. (2023, March 21). *Adobe firefly: The new family of creative generative AI models*. https://www.adobe.com/sensei/generative-ai/firefly.html

Adobe Inc. (1990). *Adobe photoshop* [Computer software]. Adobe Inc.

Adobe Inc. (1991). *Adobe premiere* [Computer software]. Adobe Inc.

Adobe Inc. (1993). *Adobe after effects* [Computer software]. Adobe Inc.

AI, H. (2024). *Higgsfield AI*. https://www.higgsfield.ai/

Ali, L., Ali, F., Abdalla, M. J., & Alotaibi, S. (2025). Beyond the hype: Evaluating the impact of generative AI on brand authenticity, image, and consumer behavior in the restaurant industry. *International Journal of Hospitality Management*, *131*, 104318. https://doi.org/10.1016/j.ijhm.2025.104318

Amankwah-Amoah, J., Abdalla, S., Mogaji, E., Elbanna, A., & Dwivedi, Y. K. (2024). The impending disruption of creative industries by generative AI: Opportunities, challenges, and research agenda. *International Journal of Information Management*, *79*, 102759. https://doi.org/10.1016/j.ijinfomgt.2024.102759

Belanche, D., Ibáñez-Sánchez, S., Jordán, P., & Matas, S. (2025). Customer reactions to generative AI vs. Real images in high-involvement and hedonic services. *International Journal of Information Management*, *85*, 102954. https://doi.org/10.1016/j.ijinfomgt.2025.102954

Brinkmann, S., & Kvale, S. (2013). *InterViews: Learning the craft of qualitative research interviewing* (3rd ed.). SAGE Publications.

Canva. (2023, October 4). *Canva unveils magic studio: The AI-powered design platform for the 99*. https://

www.canva.com/newsroom/news/magic-studio/

Coca-Cola. (2020). *Coca cola | holidays are coming 2020*. https://www.youtube.com/watch?v=6wtxogfPieA

Coca-Cola. (2025). *Coca cola | holidays are coming 2025*. https://www.youtube.com/watch?v=Yy6fByUmPuE

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.

comfyanonymous. (2023). *ComfyUI* [Computer software]. https://github.com/comfyanonymous/ComfyUI

Corvello, V. (2025). Generative AI and the future of innovation management: A human centered perspective and an agenda for future research. *Journal of Open Innovation: Technology, Market, and Complexity*, *11*(1), 100456. https://doi.org/10.1016/j.joitmc.2024.100456

Creswell, J. W., & Plano Clark, V. L. (2014). *Designing and conducting mixed methods research* (3rd ed.). SAGE Publications.

DeepMind, G. (2025). *Imagen 4 – capabilities and updates*. https://deepmind.google/models/imagen/

Duivenvoorde, B. (2024). *Generative AI and the future of marketing: A consumer protection perspective*. https://doi.org/10.2139/ssrn.4967461

ElevenLabs. (2025a). *v3 – text-to-speech model features*. https://elevenlabs.io/blog/eleven-v3

ElevenLabs. (2025b, November 7). *Introducing voice design v3*. https://elevenlabs.io/blog/voice-design-v3

Ford, J., Jain, V., Wadhwani, K., & Gupta, D. G. (2023). AI advertising: An overview and guidelines. *Journal of Business Research*, *166*, 114124. https://doi.org/10.1016/j.jbusres.2023.114124

Fortin, S., Ginosar, S., Cres, T., Yuan, P., Agarwal, A., Balamurugan, A., Miao, C., Patel, M., Pritts, L., Zhou, Y., & Ruan, T. (2025, August 26). *Introducing gemini 2.5 flash image (aka nano-banana): Generate and edit images with gemini*. https://developers.googleblog.com/en/introducing-gemini-2-5-flash-image/

Freepik. (2023). *Freepik AI image generator*. https://www.freepik.com/ai/image-generator

Google. (2023, December 6). *Introducing gemini: Our largest and most capable AI model*. https://blog.google/technology/ai/google-gemini-ai/

Google. (2024, December 3). *Introducing veo and imagen 3 on vertex AI*. https://cloud.google.com/blog/products/ai-machine-learning/introducing-veo-and-imagen-3-on-vertex-ai

Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, *18*(1), 59–82. https://doi.org/10.1177/1525822X05279903

Hartmann, J., Exner, Y., & Domdey, S. (2025). The power of generative marketing: Can generative AI create superhuman visual marketing content? *International Journal of Research in Marketing*, *42*(1), 13–31. https://doi.org/10.1016/j.ijresmar.2024.09.002

Hattori, E. A., Yamakawa, M., & Miwa, K. (2024). Human bias in evaluating AI product creativity. *Journal of Creativity*, *34*(2), 100087. https://doi.org/10.1016/j.yjoc.2024.100087

Joshi, S., Bhattacharya, S., Pathak, P., Natraj, N. A., Saini, J., & Goswami, S. (2025). Harnessing the potential of generative AI in digital marketing using the behavioral reasoning theory approach. *International Journal of Information Management Data Insights*, *5*(1), 100317. https://doi.org/10.1016/j.jjimei.2024.100317

Jung, T., Koghut, M., Lee, E., & Kwon, O. (2025). Artificial creativity in luxury advertising: How trust and perceived humanness drive consumer response to AI-generated content. *Journal of Retailing and Consumer Services*, *87*, 104403. https://doi.org/10.1016/j.jretconser.2025.104403

Labs, B. F., Dockhorn, T., Rombach, R., Blattmann, A., Dreuw, P., Ommer, B., Sauer, A., & Schuhmann, C. (2025). *FLUX.1 kontext: Flow matching for in-context image generation and editing in latent space*. https://arxiv.org/abs/2506.15742

Labs, L. (2024, June 12). *Introducing dream machine: A new chapter in video generation*. https://lumalabs.ai/blog/introducing-dream-machine

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. SAGE Publications.

Madathil, J. C. (2025). Generative AI advertisements and human–AI collaboration: The role of humans as gatekeepers of humanity. *Journal of Retailing and Consumer Services*, *87*, 104381. https://doi.org/10.1016/j.jretconser.2025.104381

McNemar, Q. (1947). Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, *12*(2), 153–157. https://doi.org/10.1007/BF02295996

Midjourney. (2023). *Midjourney documentation*. https://docs.midjourney.com/docs/quick-start

MiniMax. (2025). *Hailuo AI Video Generation Platform*. https://www.getimg.ai/blog/minimax-hailuo-ai-vid

n8n: Workflow automation. (2025). n8n.io GmbH; https://n8n.io/.

Onwuegbuzie, A. J., & Collins, K. M. T. (2007). A typology of mixed methods sampling designs in social science research. *The Qualitative Report*, *12*(2), 281–316. https://doi.org/10.46743/2160-3715/2007.1638

OpenAI. (2022, November 30). *Introducing ChatGPT*. https://openai.com/blog/chatgpt

OpenAI. (2024, February 15). *Video generation models as world simulators*. https://openai.com/research/video-generation-models-as-world-simulators

Osadchaya, E., Marder, B., Yule, J. A., Yau, A., Lavertu, L., Stylos, N., Oliver, S., Angell, R., Regt, A. de, Gao, L., Qi, K., Zhang, W. Z., Zhang, Y., Li, J., & AlRabiah, S. (2024). To ChatGPT, or not to ChatGPT: Navigating the paradoxes of generative AI in the advertising industry. *Business Horizons*, *67*(5), 571–581. https://doi.org/10.1016/j.bushor.2024.05.002

Patton, M. Q. (2015). *Qualitative research & evaluation methods: Integrating theory and practice* (4th ed.). SAGE Publications.

Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., & Sutskever, I. (2021). Zero-shot text-to-image generation. In M. Meila & T. Zhang (Eds.), *Proceedings of the 38th international conference on machine learning* (Vol. 139, pp. 8821–8831). PMLR. https://proceedings.mlr.press/v139/ramesh21a.html

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.

Runway. (2025, March). *Introducing runway gen-4*. https://runwayml.com/research/introducing-runway-g
Saldaña, J. (2016). *The coding manual for qualitative researchers* (3rd ed.). SAGE Publications.
Sallaku, K., Macca, L. S., Santoro, G., & Thu Phan, T. (2025). Advertising in the digital era: Can AI evoke emotions like humans? *Business Process Management Journal*, *31*(5), 1785–1811. https://doi.org/10.1108/bpmj-07-2024-0637
Sands, S., Campbell, C., Ferraro, C., Demsar, V., Rosengren, S., & Farrell, J. (2024). Principles for advertising responsibly using generative AI. *Organizational Dynamics*, *53*(2), 101042. https://doi.org/10.1016/j.orgdyn.2024.101042
Serra-Simón, J., Puntí-Brun, M., Espinosa-Mirabet, S., Faria Nogueira, M. A. de, Martín-Guart, R., & Tôrres de Azevedo, S. (2025). Generative artificial intelligence in advertising. Field applications in rio de janeiro and catalonia. *Telecommunications Policy*, *49*(8), 103033. https://doi.org/10.1016/j.telpol.2025.103033
Song, M., Chen, H., Wang, Y., & Duan, Y. (2024). Can AI fully replace human designers? Matching effects between declared creator types and advertising appeals on tourists' visit intentions. *Journal of Destination Marketing & Management*, *32*, 100892. https://doi.org/10.1016/j.jdmm.2024.100892
Suno. (2025, September 25). *Introducing v5*. https://help.suno.com/en/articles/8105153
Synthesia. (2025a). *Synthesia — AI video for business*. https://www.synthesia.io/
Synthesia. (2025b, October 1). *With synthesia 3.0, the next era of video is here*. https://www.synthesia.io/post/synthesia-3-0-the-next-era-of-video
Technology, K. (2024, June 10). *Kuaishou unveils KLING large model to rival sora, generating 2-minute videos up to 1080p*. https://www.kuaishou.com/en/news/kuaishou-unveils-kling-large-model-to-riv
Wang, Z., Qin, Y., Wang, H., Liu, J., Wang, J., Hu, Y., Dou, Y., Liu, Y., Wei, Y., & Hu, Y. (2024). *Seedance 1.0: Exploring the boundaries of video generation models*. https://arxiv.org/abs/2506.09113
Zapier Inc. (2025). *Zapier Workflow Automation*. https://zapier.com/.

ANNEXES

Annex A: Expert Interview Questionnaire

Semi-Structured Interview

Section 1: Background and Context

1. Could you briefly describe your current role and your experience with advertising, branding, or creative marketing involving AI tools?
2. How extensively is generative AI (e.g., ChatGPT, Midjourney, DALL-E, Runway, Sora) used in your organization's creative or marketing workflows?

Section 2: Application and Tools

3. In which areas do you or your team most frequently use generative AI (e.g. ideation, copywriting, image/video creation, personalization, campaign testing, etc.)?
4. Which tools or platforms have proven most effective, and what criteria (e.g., cost, speed, quality, controllability) guide your choice?

Section 3: Perceived Value and Outcomes

5. How has the introduction of generative AI affected the efficiency, speed, and cost of your creative production compared to traditional human-only workflows?
6. In your experience, how does AI-generated or AI-assisted content compare to fully human-created content in terms of creativity, emotional impact, and authenticity?
7. What have you observed regarding consumer reactions to AI-generated or AI-assisted marketing materials? (e.g., engagement rates, trust, brand perception, purchase intention)

Section 4: Ethical, Strategic, and Organizational Aspects

8. What is the ideal balance between human input and AI automation in creative work? Where should human judgment remain essential?
9. What ethical, transparency, or brand authenticity challenges have emerged when using AI-generated content?
10. How has generative AI adoption influenced team structures, creative processes, or required skill sets in your organization?

Section 5: Strategic Outlook

11. Overall, how would you rate generative AI's contribution to marketing effectiveness and ROI — has it enhanced storytelling, consumer connection, or sales results?
12. Looking ahead, what opportunities and risks do you foresee in the next 3–5 years regarding the role of generative AI in advertising, social media, and content branding?

Annex B: Audience Intercept Protocol

Audience Questionnaire: Consumer Perception of AI-Generated vs. Human-Created Advertising

Purpose: This questionnaire assesses consumer preferences for advertising content created using generative AI versus conventional production methods, and measures whether disclosure of AI authorship influences those preferences.

Stimuli:

- **Commercial A:** AI-generated advertisement (Coca-Cola | Holidays Are Coming 2025) (Coca-Cola, 2025)
- **Commercial B:** Human-created advertisement (Coca-Cola | Holidays Are Coming 2020) (Coca-Cola, 2020)

Procedure:

1. **Viewing:** Participant views Commercial A (AI) then Commercial B (Human) without disclosure.
2. **Question 1 (Pre-Disclosure):** "Which commercial do you prefer more—the first one or the second one?" (Record preference and explanation).
3. **Disclosure:** "The first commercial was created using artificial intelligence. The second commercial was made using conventional filming methods."
4. **Question 2 (Post-Disclosure):** "Now that you know the first commercial was AI-generated and the second was human-made, does it change your choice of preference?" (Record preference and explanation).
5. **Demographics:**
    - **Question 3:** "How old are you?"
    - **Question 4:** "And where are you from originally?"

Annex C: Expert Themes Codebook

Table C1. Thematic Analysis Codebook

Theme Name	Definition	Example Quote
Agency Transformation	A background in traditional creative fields (CG, editing, marketing) shifting to AI-first workflows, often driven by early adoption and a need to balance commercial orders with training.	”We are trying to use as much AI as possible to automate everything to make the CG sort of like work... save money instead of doing something with CG.”
Tool Ecosystem	A flexible, comprehensive suite of generative tools (Midjourney, ChatGPT, Kling, etc.) selected for price-to-value ratio, usability, and specific output needs.	”Mainly it is ChatGPT, Gemini, Sora, Higgsfield, Midjourney... Price is the first thing... UX is one of those nuances.”
Efficiency & Cost-Value Shift	The prioritization of speed and cost-efficiency, delivering high-quality (often CG-level) results at a fraction of the price, changing the client value equation.	”connection with the consumer - really really improved, from x fifty to x hundred times... sales results are increasing at least x five times.”
Role Evolution & New Competencies	The shift from pure creation to ”engineering,” direction, and hybrid roles, enabling individuals to replace larger teams and necessitating new skills like prompting.	”people simply move... to people like me - one can do that, what previously a whole team could do.”
Human-in-the-Loop Oversight	The mandatory requirement for human judgment, review, and refinement at the beginning and end of the process to ensure quality, safety, and brand alignment.	”at the beginning and at the end a human Should have its own Judgement and input and be responsible for what AI is doing.”
Inevitability & Ubiquity	The consensus that AI adoption is unavoidable, will become an invisible part of all workflows, and resisting it leads to lagging behind.	”Soon we won’t be able to distinguish what is real, what is fake and those who completely don’t want to use it will lag behind.”
Consumer Perception of Authenticity	The observation that audiences prioritize aesthetic appeal (”good taste”) over ”realness,” though low-quality ”slop” causes backlash.	”If you see a beautiful woman on screen, it doesn’t matter to you, whether this is a real photo or generated.”
Democratization & Value Erosion	The dual effect of making creativity accessible to non-designers while potentially devaluing the scarcity that made traditional art valuable.	”Previously you had to rely on the designer’s time, creativity, and now much more people can get involved in the creative process.”
Ethics & Copyright Pragmatism	A pragmatic, innovation-first approach where internal efficiency and ”don’t overthink” culture outweigh theoretical copyright concerns.	”I think that if someone needed to, they would get everything they need... if you don’t upload, don’t use AI, you will lag behind.”

Annex D: Sample Expert Interview Transcript

Transcript ID: Expert 1 **Role:** Agency Head **Sector:** Digital Marketing

Note: Transcript translated from Lithuanian/English mix. Identifying details removed.

[Response to Role/Experience]: I am the head of a marketing agency, we work mainly with LinkedIn communication, but we work with absolutely anything, from making slides, web to advertising campaigns, really whatever the client needs, we do or outsource, and clients: from tech companies to logistics, really anyone, even B2C - Coffee-in is our client.

[Response to Tools/Usage]: How widely used - basic ChatGPT, Midjourney, DALL-E, e.g., Higgsfield are used in almost all processes, from text writing (ChatGPT, Gemini, Claude) to visual generation also through the same ChatGPT, Midjourney, mainly we use them, sometimes Gemini and now we have a Higgsfield subscription, automations are: not necessarily only with marketing processes, but we use n8n and connect whatever we can with whatever. We use Zapier and its AI agents. Mainly it is ChatGPT, Gemini, Sora, Higgsfield, Midjourney.

[Response to Application Areas]: Second question about application, actually for absolutely everything, from brainstorming ideas themselves to visuals, to boards, to strategies, to video creation, also for example preparing offers for clients, when we do marketing strategy offer preparation or the same audit, or when a client comes and asks to put together a half-year plan we personalize it a lot with ChatGPT.

[Response to Selection Criteria]: Regarding the criteria of those same platforms, since I am the final decision maker, I look at several things - first of all price, because there are hundreds of those tools, if you choose one, but expensive, then you look very carefully that it really has, e.g., Higgsfield it is about 90 dollars a month, has integrated very many models from Sora to others and in one place has many subscriptions, or you look at others, which are ten, nine, six dollars a month and then you combine them, price is the first thing, quality is second, that you get what you are looking for and UX is one of those nuances, because if colleagues use it, who are not very tech savvy, so that they understand how to use it quickly, because simply the attention span is so short, because even with a new tool, when people try it, if they don't succeed in the first few prompts there, they simply drop it, because we have done many such trainings in the company team and I see how people simply try, they don't succeed, it doesn't work, it doesn't suit them etc.

[Response to Value/Efficiency]: Good question about value and results, if we talk about value, the coming of AI into our industry gives much more creativity, gives much more speed to deliver things, interesting things, Previously you had to rely on the designer's time, creativity, and now much more people can get involved in the creative process, not just the designer. Can give even potential guidelines, how they imagine much faster, much clearer and it really brings much more value, because not only one person is involved, also the same person can having a vision clearly project a question and get an answer faster, also countless AI functions, which allow designers or copywriters to work faster.

[Response to Human vs. AI Content]: Regarding content, actually I think that comparing with human created content, we have such a practice: we unequivocally do not give out content written only with AI immediately, we write, or make and then a person reviews checks, tweaks and only then gives to the client, or only then we release to the world, So from that side that fact is really very important, that it is never left only to AI, And reactions to content are good, higher engagement, higher interest, since we can make more unique content without having large budgets. Depends how you do it, style is very subjective, but basically if you do with AI do super-AI look, when everyone sees that this is

AI, it does not evoke good emotions, but if you do it so that no one understands that it is AI generated, then you can achieve many good results.

[Response to Human-AI Balance]: A human is needed and their decisions remain necessary before going on air or before going to the client, so this is the ideal balance, when you use AI to visualize what you want to visualize - your thoughts, your ideas, and then you still tweak what is needed.

[Response to Ethics/Authenticity]: Regarding ethical and brand authenticity, no question marks arose for us, because we are quite an innovative and youthful team, and there are no people with us who would overthink about it too much and I, as a manager, set the tone that I am not afraid, because I think that if someone needed to, they would get everything they need. So from that side yes, we don't put all contracts into tools, but as far as photos are concerned, it is uploaded, because if you don't upload, don't use AI, you will lag behind as a company, as an agency and alone you won't prevent anything, so that's my attitude.

[Response to Team Structure]: Regarding team structure, I think that it changed from that side, that it was necessary to start teaching and implementing new skills, about how to use [new tools] and for me until today it is so not understandable, when I see my colleague, who is 30 years old and she is capable of many things, trends, everything else, And I see how she doesn't know how to prompt simply ChatGPT, simply doesn't know basic things of the basics, even if it seems that some people can be tech savvy as if quickly catch things in the end I notice that they lack that skill, so it really changed that fact, that it is necessary to fill and improve team skills. Another thing is structure, so it helped me to accept more clients under the same amount of people and simply improve business efficiency, and I hired a person who works purely with AI process automation in my company, like a kind of AI lead and that function, I think, is extremely necessary and will be a must in every company in the future.

[Response to ROI/Effectiveness]: Regarding AI contribution to efficiency and return - it really improved in every way, I can't say immediately, what 10x or x how much, but storytelling, improved a hundred times, because you don't need large production houses to be able to fulfill something. Another thing - connection with the consumer - really really improved, from x fifty to x hundred times, depending what brand and what trademark, but the fact that you can, even talking about some ads, that you can have not ten manually created visuals anymore, but twenty, thirty various variations of visuals changed with AI tools, modified slightly And you can test again and advertise very directly to that, who is attracted by that specific one visual out of the hundred and this gives the consumer now hyper-personalization towards which everything is going now, and also gives the brand an extremely good result, and sales results for sure, because well both content created with AI avatars, because how much it raised e-commerce brands by hundreds of millions so all other businesses too, because you reduce time spent and at the same time increase client amount, for sure from business side sales results are increasing at least x five times

[Response to Future Outlook]: And how it will change in advertising in social media, I think, that really for a fact such hyper-personalization as everyone has been talking for a few years, that will continue everyone continues to look for that direct attention, because that content is increasing and will increase. Another thing, it can be that, I believe, that - people, who work as copy-writers, people will work as copy-writer engineers, who not that fully-on will write, but they will write with AI help, and that's okay, but it will take time for this to change, because it depends on generations and on everything, So yeah it will change a lot and positions, such large teams won't be needed anymore, such large budgets won't be needed, at the same time, consumers will be overfilled with attention of all those posts, videos, everything, and it will depend a lot how platforms, how big social media platforms will deal with that, And what they will do next and how they will allow that content, also whether a person will generally go to platforms in the future or simply use voice assistant AI and simply won't see advertising anymore, and that advertising will be simply physical, where a person simply walks.

[Follow-up Question: Hedonistic vs Functional Products]: Question about hedonistic vs functional product advertising. We don't have many B2C many brands, we work more with B2B. But from what others do and what I consume myself, what I see myself, comparing with marketing nuances, the difference is that that cheaper product, which is used according to function for the consumer in advertising is always much more about price, functionality, delivery speed or if next day delivery is available - all sorts of such things. Can look quite cheap, everything goes according to those such ugly ads, where there is neither beauty, nor anything, can be thrown together with AI, can even text be bad, nonsensical, because AI translation was done. Sometimes brands do other markets in local language simply completely trusting AI and it's unclear how what there, but since the product is cheap, it simply through those three seconds two visible value and a person buys it. And where there are luxury high-end products, that much much naturally everything goes through story telling, everything goes through lifestyle, as I say - selling the lifestyle and in that lifestyle simply you would need extremely much budget for all ads commercials to film, everywhere there mountains, cars, there everything else etc. and when you go through that emotion it is really much easier to do with AI, but that it wouldn't be visible, but simply it sells much more. When talking about high-end products, the more high-end the product, the more you can utilize AI, because such a large budget is not needed, Also the cheaper the product, you can utilize AI extremely well, because you don't worry neither about translate, nor there about visuals, they can go extremely ugly and can stamp, because then it doesn't matter, that it is done with AI, because they buy for functionality.