

Synthetic Influence: How AI-Generated Content is Reshaping Consumer Trust and Brand Authenticity in Entrepreneurial and Sustainable Contexts

OKACHA Sanae

Doctor, Sidi Mohamed Ben Abdellah University Fez -Morocco.

Correspondence address

Sidi Mohamed Ben Abdellah University Fez -Morocco

Cite this article

OKACHA, S. (2026). Synthetic Influence: How AI-Generated Content is Reshaping Consumer Trust and Brand Authenticity in Entrepreneurial and Sustainable Contexts. International Journal of Economics and Management Sciences, Volume 5, Issue 3 (2026), pp. 236-262.

Submitted: 14/06/2026

Accepted: 07/07/2026

Abstract:

The growing use of generative AI in marketing has introduced the concept of 'synthetic influence' persuasive signals created by algorithms to simulate human-like authenticity, thereby shaping consumer perceptions and behaviors toward brands. This article presents a conceptual framework examining how AI-generated content alters consumer trust and brand authenticity in digital markets, with particular relevance for entrepreneurs and sustainable brand development. The framework comprises three layers: the Generation Layer (how companies produce AI-generated content), the Perception Layer (how consumers interpret authenticity signals), and the Outcome Layer (the resulting impact on trust formation or erosion). A central mechanism is the “Authenticity Dilution Effect”, whereby an oversaturation of AI-generated authenticity signals within a market gradually erodes their perceived credibility. In response, the framework offers strategic guidance including managing AI intensity, preserving human-created “content anchors”, and timing AI deployment to optimize trust-building conditions. For entrepreneurs and SMEs, the framework highlights that authentic differentiation can serve as a competitive advantage over larger firms, but that using AI to simulate authenticity risks neutralizing this edge. Sustainable brands face compounded challenges, as consumers already sceptical of greenwashing become even more critical when evaluating AI-generated sustainability claims. Theoretically, the framework reconceptualizes authenticity as a socially constructed phenomenon vulnerable to algorithmic distortion, contributing to both marketing and management theory. Practically, it offers actionable guidance for practitioners navigating marketing in an AI-enhanced environment.

Keywords: Synthetic Influence, AI-Generated Content, Consumer Trust, Brand Authenticity, Digital Marketing.

Introduction

An entirely new shift is occurring in the marketing environment due to generative AI. Such innovations as GPT-4, Claude, Gemini, and even marketing-specific generative AI solutions can create written texts, images, video, and customized messages like never before. However, what sets modern generative AI apart from earlier iterations of automation technology is not only its efficiency but its capability of replicating creative human-like output and generating content that can evoke human emotions (Aggarwal et al., 2024; Desveaud & Pavone, 2025). In other words, generative AI can produce stories that feel personal, images that look genuine, and messaging that strikes a chord with specific consumer demographics – and do it all without any creative effort on behalf of humans. The scale of this shift is immense. Based on the findings provided by the latest industry analysis, almost 35% of marketers employ generative AI in their work (Han, 2024). Even more importantly, consumers themselves encounter increasing amounts of AI-generated content, which is likely why research shows that people do not differentiate AI-created content from human-made content without the latter being disclosed (Wortel et al., 2024). It has given rise to an unparalleled advantage for the company - the ability to create unique, emotional content at no marginal cost. Yet it has also given rise to a paradox for the firm. While it leverages more and more of generative AI technology for greater engagement and conversion rates, it is putting at risk one of the most crucial aspects behind its relationship with consumers - namely, trust and credibility. The application of generative AI for content generation has been fast-tracked by companies. Yet it has failed to appreciate the implications of its actions on consumer perception and trust in brands. The predominant view in marketing has been an optimized view of the process. If a company is able to generate content quicker and cheaper using AI technologies than relying on humans while getting similar results to humans, it would be a win-win (Davenport et al., 2020).

The problem is not that AI content fails to perform well in the short term. Instead, the problem is that the way success is measured, such as by looking at click-through rates, engagement, and conversion rates, does not measure erosion of trust and authenticity behind the scenes. A customer may click an advertisement created via AI or engage with an AI-created social media post, but these actions alone do not equate to building trust or fostering loyalty in brands. In fact, studies have found that once customers realize the content they have seen was generated via AI, their trust and purchase intentions decrease considerably (Zsigmondová et al., 2021; Nuttavuthisit & Thøgersen, 2017).

The real question managers face, then, is not whether to incorporate AI into marketing campaigns this is becoming increasingly difficult to avoid due to competitive pressure but how to incorporate AI while preserving consumer trust and brand authenticity.

Prior research related to AI applications in marketing has mostly revolved around two key issues: (1) Is the performance of AI-created content equal to that of human-produced content? and (2) What are the impediments and enablers of AI marketing tool adoption? (Kim & Song, 2020; Yang & Battocchio, 2021). These topics have proven valuable to shed light on immediate dynamics related to AI use in marketing, but they fail to account for the long-term consequences of synthesized influence on consumer trust. There appears to be a lack of clarity about the ways in which AI-generated content influences consumer perceptions of brands and their willingness to trust them. For instance, how do consumers assess the authenticity of content produced by

algorithms that mimic human writing style? What happens when an organization relies solely on AI to create its content strategy? How does the proliferation of algorithmic content impact the overall credibility of authenticity signals? None of these issues has received any substantial attention in prior research. Moreover, while studies have looked at technology adoption theories (Davenport et al., 2020) and authentic brands research (Beverland & Farrelly, 2010), such theories have not specifically considered how algorithmic content influences consumer trust and brand authenticity. This paper attempts to fill this gap by constructing a framework that incorporates the knowledge gained from previous studies and introduces an entirely new phenomenon known as synthetic influence that will require new theories and strategies.

This paper fills the identified gap by presenting a theoretical framework to explain how AI-generated content transforms consumer trust and brand authenticity. It will attempt to go beyond the adoption and measurement of effectiveness, and instead consider the perception and influence of algorithms on consumer behavior.

This framework becomes imperative in light of the following recent happenings: (i) studies show that about 35% of marketers are now using AI in generating marketing content; (ii) consumers have become aware of the existence of AI-generated content; and (iii) regulators in the EU and US jurisdictions have begun mandating the disclosure of AI use in marketing, with the regulation expected to be fully operationalized in 2026. These trends emphasize the immediate need for firms to adopt effective strategies for dealing with synthetic influence.

This research paper has three contributions towards marketing and management literature:

Firstly, it introduces the concept of synthetic influence in marketing through a detailed definition of the concept as an attempt by organizations to create influence based on algorithmic generation of authentic-looking persuasive communication. Second, the paper proposes a conceptual framework consisting of three layers (Generation, Perception, and Outcome) that helps understand the underlying processes by which artificial intelligence-generated messages influence consumer perception about brand authenticity. An important mechanism in this process is identified, called the Authenticity Dilution Effect, according to which excessive production of AI-generated authenticity cues results in message saturation and reduced credibility. Third, the paper draws out some strategic implications for companies attempting to deal with synthetic influence, especially regarding entrepreneurs, SMEs, and sustainable brands.

1. Methodology

1.1 Research Design and Approach

The conceptual framework of this paper was developed through a review of the existing literature and logical reasoning. This sort of approach to researching and writing follows guidelines provided in marketing that suggest concept development as well as theoretical foundation instead of collecting empirical data (Jaakkola, 2020).

1.2 Literature Identification and Synthesis

In the construction of the framework development process, a comprehensive literature review was done across three main areas: (1) Literature in Marketing and Consumer Behavior: emphasizing trust, authenticity, consumer perceptions, and brands; (2) Literature in Artificial Intelligence and Technologies: emphasizing generative AI technology, algorithmic persuasion, and AI in marketing; (3) Literature on Signaling Theory and Information Asymmetry:

highlighting how consumers make sense of signals and form judgments. The literature search was conducted across four major academic databases: Scopus, Web of Science, Google Scholar, and ScienceDirect. The search was limited to peer-reviewed journal articles, conference proceedings, and academic book chapters published between 2000 and 2025 in English. Gray literature, non-peer-reviewed sources, and opinion pieces were excluded. The initial search yielded 312 potentially relevant articles. After removing duplicates ($n = 47$), screening titles and abstracts for relevance ($n = 265$), and conducting full-text review ($n = 89$), a final set of 41 articles was retained for the theoretical synthesis. The screening process prioritized studies that explicitly addressed consumer trust, brand authenticity, AI-generated content, or algorithmic persuasion in marketing contexts. Studies focusing solely on AI technical architecture or non-marketing applications of generative AI were excluded.

Some of the keywords used for this literature review include: "consumer trust," "brand authenticity," "AI-generated content," "synthetic media," "algorithmic persuasion," "generative AI marketing," and "authenticity perceptions."

1.3 Concept Development and Framework Construction

The synthetic influence construct emerged through an iterative process of refinement, drawing on insights from three theoretical domains :

1. **Trust Theory** (Chaudhuri & Holbrook, 2001; Erdem & Swait, 2004): consists of understanding how consumers form trust judgments and the role of credibility signals.
2. **Authenticity Theory** (Beverland & Farrelly, 2010; Grayson & Martinec, 2004): Understanding authenticity as a perceptual construct and the signals consumers use to evaluate authenticity.
3. **Signaling Theory** (Spence, 1973): Understanding how firms use signals to communicate information and how consumers interpret these signals.

The three-tiered theoretical construct (Generation, Perception, Outcome) is a result of deductive reasoning, with the focus on the important phases involved in the process of influence as well as the significant processes within each of these phases. One of the important theoretical constructs identified is the Authenticity Dilution Effect, which results from saturation effects.

1.4 Proposition Development and Strategic Implications

The strategic implications have been drawn out from the framework itself through logical deduction about the manner in which firms should modify their strategy based on the mechanisms. The strategic implications for entrepreneurs, SMEs, and sustainable brands have been drawn specifically from the framework.

1.5 Methodological Rigor and Transparency

To enhance the reproducibility of this conceptual research, the framework development process adhered to established guidelines for conceptual papers in marketing (Jaakkola, 2020). The literature synthesis followed a structured approach: (1) systematic identification of relevant literature across multiple databases with documented inclusion and exclusion criteria; (2) thematic coding of retained articles to identify recurring constructs, mechanisms, and theoretical gaps; (3) iterative refinement of the synthetic influence construct through constant comparison between emerging themes and existing theoretical frameworks; and (4) logical deduction to construct the three-layer framework and derive strategic propositions. While conceptual research does not involve empirical data collection, transparency in the literature

synthesis process strengthens the credibility and traceability of the theoretical contributions. The complete list of retained articles is available from the author upon request.

1.6 Limitations of the Conceptual Approach

Since this is a theoretical study, this paper proposes ideas and theories that would need further empirical testing. This framework was developed using logic and theory rather than the actual gathering of empirical data. These propositions must be verified via experiments, questionnaires, and field investigations to prove the hypothesized causal links and processes.

2. Conceptual Foundations

2.1 Consumer Trust as an Economic and Strategic Variable

Trust is one of the essential constructs in the field of marketing and consumer behavior, yet, in many cases, the nature of trust seems somewhat abstract and intuitive. From the economic viewpoint, trust is the tool which decreases uncertainty and lowers transaction costs involved in business transactions (Erdem & Swait, 2004). A trusting consumer does not need to put much effort into evaluating a brand's offers; he or she will be ready to enter transactions with minimum information gathering and be less critical about occasional mistakes. In other words, trust is both a psychological condition and an economic resource for brands. The products of trusted brands tend to sell at a premium, and the overall value of a trusted brand tends to be higher, as it enjoys higher lifetime value of customers, lower costs of customer acquisition thanks to positive word-of-mouth recommendations, and other benefits (Poulis et al., 2026). Loss of trust, on the other hand, may quickly destroy any brand's worth.

When it comes to digital marketing, the notion of trust gets an expanded definition. The reason why digital platforms are associated with informational asymmetry – i.e., lack of consumers' ability to independently check the truthfulness of a brand or its product characteristics before purchasing – is why trust is particularly important here. To make up their minds on whether a brand is trustworthy or not, consumers will have to look for signals – such as the brand's reputation, content quality, etc.

2.2 Brand Authenticity as a Perceptual Construct

Indeed, as pointed out in contemporary studies, the nature of authenticity is inherently perceptual in essence (Beverland & Farrelly, 2010; Grayson & Martinec, 2004). As suggested by some authors, there may be an attempt to establish some kind of objectivity when referring to authenticity through heritage, practices, or values. In terms of consumer psychology, however, the perceptual approach will prove more useful since consumer behavior is shaped by perceptions rather than the objective truth.

As mentioned above, synthetic influence has nothing to do with whether the content was created by humans or AI. Authenticity cannot be established according to the method of creation, as in both cases the content might be considered authentic by customers provided that such content satisfies their needs and expectations. Authenticity serves as a signal in the mind of consumers. Consumers will read signs such as transparency in practices, integrity of stated value systems with action, proof of human touch, and recognition of limitations as signs of authenticity. In the presence of these signs, consumers can assume that the brand is an authentic one. However, the absence or presence of contradictory signs indicates manipulation. Although trust and authenticity are separate concepts, they are closely intertwined when it comes to consumer perceptions. As a matter of fact, authenticity serves as an indicator of a company's

trustworthiness since it shows the consumer that the brand is sincere, real, and aligned with their values. However, consumers' perceptions of lack of authenticity can damage a company's credibility despite its trustworthiness. When speaking about the concept of synthetic influence, it becomes difficult to establish the connection between signals that the brand sends and its actual characteristics. The reason for that is that AI algorithms can create signals of authenticity even when the brand lacks such an attribute. Thus, one can state that the essence of the concept of synthetic influence lies in the fact that signals are intentionally created by the brand in order to make its audience perceive the brand's credibility in a certain way.

2.3 AI-Generated Content in Marketing: From Automation to Simulation

The use of generative AI is different from other types of automation that have been used in marketing before. Before generative AI became popular, marketing automation solutions like email marketing software and programmatic advertising tools were simply tools that enabled marketers to automate the distribution and targeting of content but not its creation, it remained entirely human-created. The difference between traditional and generative forms of automation lies in what is automated through such processes. While in traditional automation techniques it is only actions that are executed automatically, in the case of generative AI, it is the creation process that can be automated completely. It means that the content does not necessarily originate from a human mind anymore.

Uses of generative AI in marketing include providing personalized product recommendations, customizing email campaigns, generating social media content, writing advertising copy, and creating synthetic images or videos. All of these involve AI being trained on extensive amounts of data to learn how to generate material that optimizes a particular metric such as engagement, conversion, click-through rate. The result is highly optimized content that statistically maximizes the chance to influence behavior but may lack any of the authenticity or human judgment associated with genuinely human-written material.

Terminology Clarification: This paper uses several related terms that should be clearly distinguished. Table 1 provides a structured comparison of synthetic influence with adjacent constructs.

Table 1: Distinguishing Synthetic Influence from Related Constructs

Construct	Definition	Focus	Key Distinction from Synthetic Influence
AI-Generated Content	Any content created by artificial intelligence systems	Output	Umbrella category; not all AI content simulates authenticity.
Synthetic Media	AI-generated media that simulates human-created content (deepfakes, synthetic images)	Format	Focuses on media type, not persuasive intent.

Algorithmic Persuasion	Use of algorithms to influence consumer attitudes or behaviors	Persuasion	Broader category; does not require authenticity simulation.
Automated Marketing	Technology-driven automation of marketing tasks	Efficiency	Focuses on process automation, not authenticity signals.
Digital Influence	Online strategies to shape consumer perceptions and behaviors	Strategy	General category; includes human and AI-driven tactics.
Synthetic Influence	Algorithmically generated persuasive signals that simulate human authenticity	Authenticity Simulation	Unique combination: algorithmic generation + authenticity simulation + persuasive intent.

(1) AI-generated content - Content produced using any AI technology. This is an umbrella term. Not all AI-generated content constitutes synthetic influence; only content designed to simulate human authenticity for persuasive purposes qualifies.

(2) Synthetic media - AI-generated media that simulates human-made content, whether images, videos, or audio. Synthetic media becomes synthetic influence when it is deployed with persuasive intent and simulates authenticity.

(3) Synthetic influence - Persuasive signals algorithmically produced to mimic human authenticity for consumer influence purposes. This is the subject of this paper. What distinguishes synthetic influence is the convergence of three elements: algorithmic generation, authenticity simulation, and persuasive intent.

(4) Authenticity signals - Any cues that consumers perceive as authentic and use as signals of human involvement, whether real or not. For the purposes of this paper, the focus is on AI-generated authenticity signals.

3. The Synthetic Influence Management Framework

3.1 Definition of Synthetic Influence

We define synthetic influence as follows:

- **Synthetic Influence** = Algorithmically generated persuasive signals that simulate human authenticity to influence consumer perception and behavior.

These components of the definition need further elaboration as well. To begin with, synthetic influence relies on computer-based algorithms for its creation; in other words, it refers to content that is automatically generated and not created manually through an individual's efforts of creativity. This does not mean that algorithms work without human input; people determine their goals and provide data and parameters to optimize the process. Yet it is the algorithm itself that creates content.

Another important aspect of synthetic influence is that it uses certain persuasion cues. Persuasion cues refer to some specific cues used in influencing consumers' perceptions or behaviors. These cues might take various forms, ranging from emotional appeals and testimonials to appeals to values or demonstration of product advantages. Third, simulation of

human authenticity is key to synthetic influence. It is not enough for the algorithm simply to create persuasive content – it must create content which simulates authenticity in human communication. This is what sets it apart from more traditional forms of advertising and marketing communication, where the latter is generally known to be explicitly persuasive communication on behalf of the brand. The purpose of synthetic influence is to create content that simulates authentic human communication. Fourth, synthetic influence aims to exert influence on consumer decision-making. The ultimate goal of synthetic influence is not only the creation of the content, but influencing the perception of the brand as well as the consumer's decision-making, brand loyalty and/or behavior in some way. This point about distinguishing synthetic from traditional influence is crucial. Whereas the traditional form of influence in marketing occurs through persuasion made by the brand, where it makes a statement about its products and consumers assess this statement, synthetic influence operates via implicit persuasion.

3.2 The Synthetic Influence Management Framework: Three Integrated Layers

The framework comprises three integrated layers that explain how synthetic influence operates and affects consumer trust and brand authenticity:




Figure 1 : The Synthetic Influence Management Framework

3.2.1 The Generation Layer: How Firms Produce AI-Generated Content

The Generation Layer encompasses the processes through which firms create AI-generated content. This layer is characterized by firm control and optimization. The firm selects the AI tools, defines the objectives (e.g., maximize engagement, increase conversion), selects or creates the training data, and sets the parameters that guide content generation. Key characteristics of the Generation Layer include:

- **Personalization at Scale.** Generative AI enables firms to create highly personalized content for individual consumers or micro-segments, without the cost of human customization. An e-commerce platform can generate personalized product recommendations for millions of users simultaneously; a marketing automation system can generate customized email

messages for each recipient based on their browsing history, purchase behavior, and demographic profile.

- **Emotional Mimicry.** Generative AI systems are trained on large datasets of human-authored content, including content designed to evoke emotional responses. Through this training, AI systems learn to generate content that mimics human emotional expression. An AI system might generate a social media post that includes emotional language, personal anecdotes, or appeals to shared values all designed to create the appearance of authentic human communication.
- **Narrative Construction.** AI systems can generate coherent narratives that connect product features to consumer benefits, that position the brand within a larger story, or that create a sense of brand personality. These narratives are constructed algorithmically, but they can feel authentic and compelling to consumers who encounter them.
- **Optimization for Engagement.** The Generation Layer is fundamentally driven by optimization. AI systems are trained to maximize specific metrics click-through rates, engagement, conversion rates. This optimization creates a systematic bias toward content that is emotionally engaging, personally relevant, and persuasive, regardless of whether that content is truthful or aligned with the brand's actual values.

The optimization for engagement creates systematic biases in AI-generated content. For example, if an AI system is trained to maximize click-through rates, it will learn to generate emotionally provocative content, sensationalized claims, or content that exploits cognitive biases. If optimized for conversion rates, it will learn to generate content that emphasizes benefits while downplaying limitations or risks. These optimization objectives can lead to content that is engaging and persuasive, but potentially misleading or manipulative. This creates an ethical tension: the same optimization that makes AI-generated content effective at influencing consumer behavior may also make it deceptive.

3.2.2 The Perception Layer: How Consumers Interpret Authenticity Signals

The Perception Layer encompasses how consumers encounter AI-generated content and form judgments about its authenticity and credibility. This layer is characterized by consumer agency and interpretation. Consumers do not passively receive persuasive signals; they actively interpret those signals based on their prior knowledge, expectations, and cognitive heuristics.

Key characteristics of the Perception Layer include:

Authenticity Evaluation. When consumers encounter content, they evaluate whether it appears authentic. This evaluation is based on multiple cues: Does the content feel genuine? Does it align with the brand's established identity? Does it appear to reflect genuine human experience or expertise? Are there signals of human involvement (e.g., personal anecdotes, acknowledgment of limitations)? Is the content transparent about its source and methods?

Sustainability Sensitivity. For brands making environmental or social claims, an additional evaluation dimension operates in the Perception Layer. Consumers exposed to sustainability messaging have developed heightened skepticism due to widespread greenwashing the practice of making exaggerated or false environmental claims. This pre-existing skepticism creates an amplified scrutiny effect: when consumers encounter AI-generated sustainability content, they apply more rigorous authenticity evaluation standards than they would for conventional product claims. The interaction between greenwashing awareness and AI-generated content creates

what we term compounded skepticism a multiplicative rather than additive increase in consumer doubt. Consumers do not merely question whether the sustainability claim is true; they simultaneously question whether the communication itself is authentically human. This dual-layer skepticism makes sustainable brands particularly vulnerable to the Authenticity Dilution Effect, as their core value proposition depends on perceived genuine commitment rather than mere product functionality.

Credibility Assessment. Consumers assess the credibility of content based on perceived source credibility, message quality, and consistency with prior information. If content appears to come from a credible source and is consistent with prior information, consumers are more likely to believe it. If content appears inconsistent or comes from a questionable source, consumers may discount it.

Relatability and Connection. Consumers evaluate whether they can relate to the content and whether it creates a sense of connection with the brand. Content that feels personally relevant, that acknowledges consumer concerns or values, and that creates a sense of shared understanding is more likely to be perceived as authentic and to influence consumer behavior.

Recognition of Artificiality. An important dynamic in the Perception Layer is the extent to which consumers recognize that content is AI-generated. Research suggests that explicit disclosure of AI generation can significantly reduce consumer trust and purchase intentions (Zsigmondová et al., 2021). However, many consumers may not recognize AI-generated content as such, particularly if it is well-crafted and aligned with consumer expectations. This creates an asymmetry: consumers who do not recognize AI generation may perceive the content as authentic, while consumers who do recognize it may perceive it as inauthentic.

Recent research on human-AI interaction provides additional insight into how consumers process AI-generated content. Studies suggest that consumers evaluate AI-generated content through an algorithmic appreciation lens: when the identity of the content creator is unknown, consumers may attribute higher credibility to AI-generated content due to perceptions of machine objectivity (Sundar, 2008). However, this effect reverses when AI authorship is disclosed, activating what has been termed algorithmic aversion a tendency to discount AI-generated outputs even when they are objectively accurate (Dietvorst et al., 2015). The Perception Layer thus operates under conditions of asymmetric information: when AI generation is undisclosed, consumers may apply heuristic processing that favors AI content; when disclosed, they may apply systematic scrutiny that discounts it.

Individual Differences in Authenticity Evaluation. Consumer evaluation of authenticity is not uniform across all individuals. Research suggests important individual differences that moderate the effectiveness of synthetic influence, creating distinct consumer segments with varying susceptibility to AI-generated authenticity signals.

AI Literacy. Consumers with higher AI literacy are more likely to recognize AI-generated content and to be skeptical of it (Han, 2024). These consumers possess greater knowledge of AI capabilities and limitations, enabling them to identify stylistic patterns characteristic of machine-generated text. High AI literacy consumers function as an early warning system within their social networks, potentially accelerating the Authenticity Dilution Effect by raising awareness among less knowledgeable peers. For firms, high-AI-literacy segments represent a trust-risk population where synthetic influence may backfire if AI use is detected.

Brand Familiarity. Consumers familiar with a brand develop mental models of its authentic voice, communication style, and values. This familiarity enables them to detect deviations that signal AI-generated content, such as shifts in tone, vocabulary, or emotional expression that depart from the brand's established identity. For brands with strong, consistent identities developed over time, AI-generated content risks creating dissonance that familiar consumers perceive as inauthentic. Conversely, new brands without established consumer familiarity may face lower detection risk but also lack the trust reservoir to absorb discovery.

Generational Differences. Younger consumers (Gen Z, younger Millennials) have grown up in digitally saturated environments and may exhibit greater acceptance of AI-generated content as a normalized aspect of brand communication. However, this acceptance is conditional: younger consumers also demonstrate high sensitivity to inauthenticity and may penalize brands that fail to disclose AI use, particularly when authenticity is a stated brand value. Older consumers, by contrast, may exhibit stronger preferences for human-authored content and greater skepticism toward AI-generated communications, reflecting lower baseline trust in algorithmic systems.

Trust Propensity. Consumers vary in their dispositional trust the general tendency to trust or distrust others and institutions. High trust propensity consumers may accept AI-generated content at face value, processing authenticity signals heuristically rather than systematically. Low trust propensity consumers scrutinize content more carefully, activating persuasion knowledge and counter-arguing when they suspect manipulative intent. This individual difference interacts with disclosure: low trust propensity consumers may respond positively to transparent AI disclosure as a signal of brand honesty, while high trust propensity consumers may not require such reassurance.

These individual differences suggest that synthetic influence is not uniformly effective across all consumer segments. Firms should tailor their AI intensity and transparency strategies based on target audience characteristics. For example, brands targeting high-AI-literacy or low-trust-propensity consumers should favor greater transparency and lower AI intensity, while brands targeting low-AI-literacy segments may have greater latitude for AI-generated content though this raises ethical considerations regarding vulnerable consumer populations.

3.2.3 The Outcome Layer: Trust Formation or Erosion

The Outcome Layer encompasses the consequences of synthetic influence for consumer trust and brand authenticity. This layer is characterized by the formation or erosion of trust, and the resulting behavioral consequences.

- **Trust Formation.** When consumers perceive AI-generated content as authentic, and when that content is consistent with their values and expectations, synthetic influence may contribute to trust formation. A consumer might read an AI-generated testimonial or review that resonates with their concerns, and this might increase their trust in the brand and their likelihood of purchase. From the consumer's perspective, they have encountered authentic content that has influenced their perception; they may not recognize that the content is AI-generated.
- **Trust Erosion.** Conversely, when consumers recognize that content is AI-generated, or when they discover that content, they believed to be authentic was actually AI-generated, trust may erode rapidly. The discovery that a brand has used AI to simulate authenticity can be perceived as deceptive, creating a sense of betrayal. This erosion of trust can be

particularly severe because it affects not only the specific content that was revealed to be AI-generated, but also consumer perceptions of the brand more broadly. If a brand has used AI to simulate authenticity in one domain, consumers may become suspicious of all the brand's communications.

- **Behavioral Consequences.** The formation or erosion of trust has direct behavioral consequences. Consumers with high trust are more likely to purchase, more likely to pay premium prices, more likely to engage in repeat purchases, and more likely to recommend the brand to others. Conversely, consumers with low trust or eroded trust are more likely to defect to competitors, less likely to engage in repeat purchases, and more likely to share negative information about the brand.

3.3 The Authenticity Dilution Effect: A Critical Mechanism

A key component of the model is the notion of the Authenticity Dilution Effect, which reflects our hypothesis regarding how companies will use AI-generated content to scale up engagement and conversion through systematic scaling.

The Authenticity Dilution Effect occurs as follows: When a company first starts using AI-generated content, in the early days of AI, AI-generated content will appear to be authentic and will thus influence consumer behavior. The messages created by AI are effective because they are persuasive, engaging, and emotional. Moreover, consumers will not question the authenticity of AI-generated content.

When other firms follow suit and AI becomes increasingly commonplace, there may occur a situation where consumers start recognizing AI-generated content because they encounter it frequently from various brands and channels. This may lead consumers to grow more skeptical about the authenticity of any given piece of content. In addition, the sheer abundance of AI-generated authenticity may begin to dilute the effectiveness of authenticity.

It results in a paradoxical situation whereby, while firms increase the volume of synthetic content to ensure maximum engagement and conversions, they risk rendering the use of synthetic influence ineffective collectively. The authenticity cues that helped in generating maximum impact while AI content was scarce may end up being ineffective as the latter grows ubiquitous due to the development of consumer “authenticity skepticism” – i.e., skepticism toward all content, based on the understanding that this content might well be artificially generated.

What the Authenticity Dilution Effect implies goes beyond mere consumer skepticism. The dilution effect has its roots in the oversaturation of authenticity cues – i.e., in a phenomenon similar to the “wolf boy cry” effect, whereby false alarms render even real alarms powerless to generate an adequate response. In other words, if authenticity cues are rare and credible, they convey the message of authentic content. If authenticity cues are abundant and many of them are artificial, cues will lose credibility altogether.

The Authenticity Dilution Effect would imply a collective action problem faced by companies. While each company would have an interest in increasing the extent to which artificial intelligence is used to generate its content in order to increase engagement and conversion, the cumulative effect of all companies doing so is expected to be detrimental to the effectiveness of artificial intelligence influence of all firms. There would be a “race to the bottom,” as

companies might find themselves forced to continually increase the usage of artificial intelligence in their content generation efforts.

3.3.1 Temporal Dynamics of the Authenticity Dilution Effect

The Authenticity Dilution Effect is predicted to operate on a temporal dimension that is important for strategic planning. The effect is not instantaneous; it may unfold over time as consumer exposure to AI-generated content accumulates and as consumer awareness increases. Several temporal dynamics are noteworthy:

- **Lag Effect** : We theorize that there is typically a lag between when firms begin scaling AI-generated content and when consumers become aware and skeptical. This lag creates a window of opportunity for early adopters, but also creates risk if consumer awareness increases rapidly. Early research suggests this lag may be 6-18 months, depending on market visibility and media coverage. Firms that adopt AI-generated content early may benefit from high effectiveness while consumer skepticism is still low. However, as consumer awareness increases, the effectiveness of synthetic influence may decline, creating pressure to adjust strategy.
- **Threshold Effects** : The effect may not be linear; there may be threshold points where consumer skepticism suddenly increases. For example, a critical incident (e.g., a high-profile case of undisclosed AI use, regulatory action, or media exposé) may cause consumer skepticism to jump suddenly. Similarly, when a critical mass of consumers become aware of AI use (e.g., 30-40% of consumers), skepticism may increase more rapidly. Understanding these threshold effects is important for strategic planning, as they suggest that firms should not assume linear relationships between AI adoption and effectiveness.
- **Reversibility** : The effect may be partially reversible if firms reduce their use of AI-generated content or increase transparency. However, rebuilding authenticity is typically slower than eroding it. Once consumer trust is damaged, it may take months or years to rebuild, even with significant changes in strategy. This suggests that firms should be cautious about aggressive AI adoption, as the costs of reversal may be high.
- **Cyclical Patterns** : As consumer skepticism increases and firms reduce AI use, authenticity signals may regain credibility, creating cyclical patterns in the effectiveness of synthetic influence. Early adopters may benefit from high effectiveness, but as saturation increases, effectiveness may decline. As firms reduce AI use in response to declining effectiveness, authenticity signals may regain credibility, creating opportunities for new cycles of adoption. Understanding these cyclical patterns can help firms anticipate market dynamics and adjust their strategies accordingly.

3.4 Full Framework Dynamics: The Causal Chain

The complete framework can be represented as a causal chain:

AI Generation → Synthetic Signals → Consumer Perception → Trust Formation/Erosion → Behavioral Consequences → Feedback Loop




Figure 2 : Causal Chain of Synthetic Influence with Moderating Factors

- **Stage 1: AI Generation.** Firms use generative AI to produce content designed to maximize engagement, conversion, or other business objectives. This content is personalized, emotionally resonant, and designed to simulate human authenticity.
- **Stage 2: Synthetic Signals.** The AI-generated content creates persuasive signals emotional language, social proof, appeals to shared values, demonstrations of product benefits. These signals are designed to influence consumer perception.
- **Stage 3: Consumer Perception.** Consumers encounter these signals and interpret them. They evaluate whether the content appears authentic, whether it is credible, and whether it is relevant to their needs and values. At this stage, consumers may or may not recognize that the content is AI-generated.
- **Stage 4: Trust Formation or Erosion.** Based on their perception of the content, consumers form judgments about the brand's trustworthiness. If they perceive the content as authentic and aligned with their values, trust may increase. If they perceive the content as inauthentic or manipulative, or if they discover that content, they believed to be authentic was AI-generated, trust may erode.
- **Stage 5: Behavioral Consequences.** Changes in trust lead to changes in consumer behavior. Increased trust leads to increased purchase likelihood, higher customer lifetime value, and positive word-of-mouth. Decreased trust leads to reduced purchase likelihood, customer defection, and negative word-of-mouth.
- **Stage 6: Feedback Loop.** The behavioral consequences feed back into the system. If synthetic influence is effective in increasing sales and engagement, firms have incentive to scale their use of AI-generated content. If synthetic influence becomes less effective due to authenticity dilution, firms may need to adjust their strategies. Additionally, as more

consumers experience AI-generated content and become aware of its prevalence, consumer skepticism increases, which affects how consumers perceive and interpret synthetic signals.

3.4.1 Moderating Factors in the Causal Chain

While the causal chain provides a useful framework for understanding synthetic influence, it is important to recognize that numerous factors may moderate each link in the chain:

- **Stage 2-3 Transition (Synthetic Signals → Consumer Perception):** Content quality is a key moderator. High-quality AI-generated content is more likely to be perceived as authentic than low-quality content. Similarly, alignment with brand voice and values moderates whether signals are perceived as authentic. Content that is inconsistent with the brand's established voice or values will be perceived as less authentic.
- **Stage 3-4 Transition (Consumer Perception → Trust Formation/Erosion):** Brand reputation is a key moderator. Consumers with high prior trust in a brand are more likely to perceive AI-generated content as authentic. Conversely, consumers with low prior trust may be more skeptical. Additionally, consumer AI literacy moderates this relationship: consumers with higher AI literacy are more likely to recognize AI-generated content and to discount it. For sustainable brands specifically, prior exposure to greenwashing functions as an additional moderator. Consumers who have encountered greenwashing in the past apply heightened scrutiny to all sustainability communications, and discovering that such communications were AI-generated amplifies distrust disproportionately. This interaction effect means that sustainable brands face a steeper trust penalty from undisclosed AI use than conventional brands, as the deception is perceived as compounding both the sustainability claim and the communication's authenticity are called into question simultaneously.
- **Stage 4-5 Transition (Trust Formation/Erosion → Behavioral Consequences):** Disclosure of AI use is a critical moderator. Transparent disclosure of AI use may help prevent trust erosion even when consumers recognize AI generation. Conversely, discovery of undisclosed AI use can cause severe trust erosion. Research suggests that proactive disclosure is more effective than reactive disclosure after discovery. This aligns with emerging literature on AI transparency and explainability (Felzmann et al., 2019; Shin & Park, 2019), which argues that algorithmic transparency fosters user trust by enabling consumers to understand how and why AI systems generate specific content. However, transparency also entails risks: disclosure may activate persuasion knowledge (Friestad & Wright, 1994), reducing content effectiveness even when consumers appreciate the honesty. The net effect of disclosure thus depends on contextual factors including brand reputation, consumer AI literacy, and the nature of the content being disclosed.
- **Stage 5-6 Feedback Loop (Behavioral Consequences → System Adjustment):** Consumer skepticism moderates the feedback loop. As consumer skepticism increases, the feedback loop may become stronger, creating accelerating erosion of synthetic influence effectiveness. This creates a self-reinforcing cycle where increased AI use leads to increased skepticism, which leads to reduced effectiveness, which incentivizes further AI use to maintain effectiveness. Understanding these moderating relationships is critical for developing effective strategies.

4. Strategic And Managerial Applications

4.1 Strategic Positioning Choices: AI Intensity and Transparency

Firms face a fundamental strategic choice regarding their approach to AI-generated content: the level of AI intensity (the proportion of their content that is AI-generated) and the level of transparency (the extent to which they disclose their use of AI).

These two dimensions create a strategic matrix with four distinct positioning options:

- **High AI Intensity, High Transparency.** In this positioning, firms use AI extensively for content generation but are transparent about this use. They might explicitly label AI-generated content, disclose their use of AI in their marketing practices, or provide explanations of how AI is used. This positioning prioritizes efficiency and scale, but sacrifices the authenticity simulation that is central to synthetic influence. The advantage is that it may help avoid the trust erosion that can result from discovery of undisclosed AI use. The disadvantage is that it may reduce the persuasive effectiveness of synthetic influence.
- **High AI Intensity, Low Transparency.** In this positioning, firms use AI extensively for content generation but do not disclose this use. They attempt to create the appearance of authentic human communication while actually using AI. This positioning may maximize the short-term effectiveness of synthetic influence, but creates significant risk of trust erosion if consumers discover the undisclosed AI use. This positioning is increasingly problematic from both ethical and legal perspectives, as regulators in various jurisdictions are beginning to require disclosure of AI use in marketing.
- **Low AI Intensity, High Transparency.** In this positioning, firms use AI selectively and are transparent about this use. They might use AI for specific tasks (e.g., generating product descriptions or initial drafts) while maintaining significant human involvement and oversight. This positioning may maintain authenticity while benefiting from AI efficiency. It is likely to become increasingly common as firms seek to balance efficiency gains with authenticity preservation.
- **Low AI Intensity, Low Transparency.** In this positioning, firms use AI minimally and do not emphasize their use of AI. This positioning is essentially a reversion to traditional content creation practices. It may be appropriate for brands that prioritize authenticity and human connection, but it sacrifices the efficiency and scale benefits of AI.




Figure 3 : Strategic Positioning Matrix: AI Intensity and Transparency

4.1.1 Selecting the Optimal Positioning

The optimal positioning depends on several firm and market characteristics:

- **Brand Positioning.** Brands positioned as premium, authentic, or human-centered should favor Low AI Intensity, High Transparency. These brands have built their value proposition on authenticity and human connection; high AI intensity would undermine this positioning. Conversely, brands positioned as innovative or technology-forward may be able to support High AI Intensity, High Transparency. These brands can position AI use as an innovation advantage.
- **Competitive Environment.** In highly competitive markets where differentiation is difficult, firms may need High AI Intensity to maintain efficiency and cost competitiveness. In markets where authenticity is a key differentiator, Low AI Intensity may be preferable to maintain competitive advantage.
- **Consumer Expectations.** In B2B markets or among professional audiences, consumers may expect and accept AI use if disclosed. In B2C markets with emotional connection, consumers may prefer human-authored content. Understanding target audience expectations is critical.
- **Regulatory Environment.** In jurisdictions with strict AI disclosure requirements (e.g., EU, California), Low Transparency positioning becomes legally risky. In jurisdictions with minimal regulation, firms have more flexibility.
- **Resource Constraints.** Firms with limited content creation resources may need High AI Intensity regardless of strategic preference. In this case, High Transparency becomes more important to maintain trust despite high AI intensity.
- **Consumer Segment Characteristics.** The optimal positioning is also contingent on the target audience's AI literacy, generational composition, and trust propensity. Brands serving high-AI-literacy audiences should adopt greater transparency to preempt consumer detection and backlash. Brands targeting younger, digitally native consumers may operate with moderate AI intensity, provided transparency is maintained. Brands serving older or low-trust-propensity consumers should minimize AI intensity in trust-critical communications, as these segments exhibit stronger preferences for human-authored content and greater skepticism toward algorithmic persuasion.

4.2 Managing Synthetic Influence: Practical Strategies

For firms seeking to use AI-generated content while maintaining consumer trust and brand authenticity, several practical strategies emerge from the framework:

- **Control Volume of AI-Generated Content :** Rather than maximizing the proportion of AI-generated content, firms should carefully control the volume. This means using AI selectively for specific tasks where it adds clear value (e.g., generating product recommendations, personalizing email subject lines) while maintaining human authorship for content that is central to brand identity and consumer relationships (e.g., brand storytelling, values communication, crisis communication).
- **Maintain Human Anchors :** Firms should maintain visible human involvement in their content strategy. This might include featuring real employees, customers, or brand ambassadors; acknowledging human creative input even when AI is involved; or creating content that explicitly reflects human judgment and values. These "human anchors" signal

authenticity and help consumers maintain trust even when AI is involved in content generation.

- **Use AI Selectively for Personalization, Not Simulation :** Rather than using AI to simulate human authenticity, firms should use AI for legitimate personalization purposes tailoring content to individual consumer preferences, providing relevant recommendations, or customizing communication based on consumer behavior. This use of AI is more defensible ethically and is less likely to erode trust if discovered.
- **Transparent Disclosure of AI Use :** When AI is used in content generation, firms should consider transparent disclosure of this use. This might be done through explicit labeling (e.g., "This recommendation was generated by our AI system"), through general statements about AI use in marketing practices, or through detailed explanations of how AI is used and why. While transparent disclosure may reduce short-term effectiveness of synthetic influence, it may help build long-term trust and avoid the trust erosion that can result from discovery of undisclosed AI use.
- **Maintain Content Quality Standards :** AI-generated content should be subject to the same quality standards as human-authored content. This means reviewing AI-generated content for accuracy, appropriateness, alignment with brand values, and potential for misinterpretation. It also means being willing to reject or significantly revise AI-generated content that does not meet these standards, even if doing so reduces efficiency gains.

4.3 Trust Calibration Strategy: Regulating Trust, Not Maximizing Automation

A critical insight from the framework is that firms should focus on calibrating trust rather than maximizing automation. The traditional logic of marketing automation has been to maximize efficiency and scale by automating as much of the marketing function as possible. However, this logic is misaligned with the goal of building long-term consumer trust and brand authenticity. Instead, firms should adopt a trust calibration approach, which involves:

- **Identifying Trust-Critical Functions.** Firms should identify which marketing functions are most critical for building and maintaining consumer trust. These typically include: brand storytelling and values communication, customer service and support, crisis communication, and relationship-building with key customers or influencers. These functions should be prioritized for human involvement and oversight, even if AI could perform them more efficiently.
- **Using AI for Efficiency in Non-Critical Functions.** For marketing functions that are less critical for trust building such as routine content generation, data analysis, or administrative tasks firms can more aggressively adopt AI to gain efficiency benefits.
- **Monitoring Trust Indicators.** Firms should implement systems to monitor consumer trust and authenticity perceptions. This might include regular surveys of consumer trust, analysis of social media sentiment, monitoring of customer reviews and feedback, and analysis of customer retention and lifetime value. These indicators should be used to adjust the balance between AI automation and human involvement.
- **Adjusting Strategy Based on Trust Trends.** If trust indicators decline, firms should be prepared to reduce their reliance on AI-generated content and increase human involvement,

even if this reduces short-term efficiency. Conversely, if trust indicators remain stable or improve, firms can be more aggressive in adopting AI for specific functions.

4.3.1 Metrics and Decision Rules for Trust Calibration

Trust calibration requires specific metrics and decision rules to guide implementation:

- **Trust Index.** Develop a composite trust index based on survey items measuring consumer trust (e.g., "I trust this brand," "This brand is reliable," "I would recommend this brand"). Use a 10-point scale and calculate quarterly. Set a target trust level (e.g., 7.0 on a 10-point scale). Decision rule: If trust declines below 6.5, increase human involvement in content creation.
- **Authenticity Perception.** Measure consumer perception of brand authenticity using items such as "This brand communicates authentically," "I believe this brand is genuine," "This brand's communications feel real." Decision rule: If authenticity perception declines by more than 10% year-over-year, reduce AI intensity.
- **Customer Retention.** Monitor customer retention rates as a proxy for trust. Decision rule: If retention declines by more than 5% compared to prior year, investigate whether AI content changes are responsible.
- **Social Sentiment.** Monitor social media sentiment regarding brand authenticity using sentiment analysis tools. Decision rule: If negative sentiment regarding 'artificial' or 'inauthentic' communication increases by more than 20%, reduce AI use.
- **AI Awareness.** Monitor consumer awareness of AI use through surveys. Decision rule: If awareness increases by more than 30% quarter-over-quarter, increase transparency to prevent trust erosion.

4.4 Implications for Entrepreneurs and Small Businesses

The framework has particular relevance for entrepreneurs and small and medium-sized enterprises (SMEs), who face distinct challenges and opportunities in the context of synthetic influence.

- **Challenge: Limited Trust Reservoir.** Unlike established brands with substantial accumulated consumer trust, startups and SMEs often lack an established trust reservoir. This makes synthetic influence particularly risky for these firms. If consumers discover that a startup has used AI to simulate authenticity, the trust damage may be severe precisely because the startup had limited trust to begin with. A large, established brand might recover from such a discovery; a startup may not.
- **Challenge: Resource Constraints.** SMEs often have limited resources for content creation and marketing. The efficiency gains from AI-generated content are therefore particularly attractive to these firms. However, the framework recommends significant human oversight and quality control, which may be difficult for resource-constrained firms to implement.
- **Opportunity: Authenticity as Competitive Advantage.** Startups and SMEs often compete against larger incumbents by emphasizing authenticity, human connection, and genuine engagement. This authenticity is frequently a key competitive advantage. Using AI to simulate authenticity risks destroying this advantage. Instead, SMEs should consider emphasizing their authentic human involvement as a differentiator.

- **Strategic Recommendation for Entrepreneurs and SMEs:** We recommend that entrepreneurs and SMEs adopt a *Low AI Intensity, High Transparency* positioning. This means: (1) Using AI selectively for non-trust-critical functions (e.g., data analysis, routine content generation, administrative tasks); (2) Maintaining visible human involvement in all trust-critical functions (brand storytelling, customer service, values communication); (3) Being transparent about AI use when it is employed; (4) Building authenticity and human connection as core competitive differentiators; (5) Focusing on quality over volume in content creation.

This approach allows SMEs to benefit from AI efficiency while preserving the authenticity that is often their key competitive advantage.

4.5 Implications for Sustainable Brands

The framework has particular relevance for sustainable brands and brands making environmental, social, or governance (ESG) claims, who face unique challenges in the context of synthetic influence. First challenge is the Heightened Skepticism, as sustainable brands already face heightened consumer skepticism due to concerns about greenwashing the practice of making exaggerated or false environmental claims. This skepticism creates a challenging environment for any marketing communication, including AI-generated content. The second one is Authenticity Criticality about sustainable brands, authenticity is not merely a nice-to-have; it is central to the brand's value proposition. Consumers choose sustainable brands precisely because they believe the brand is genuinely committed to environmental or social values. If consumers perceive the brand's sustainability messaging as inauthentic, the entire value proposition is undermined.

A potential risk is compounded skepticism. If consumers discover that a sustainable brand used AI-generated content for its sustainability messaging, the trust damage may be particularly severe. The discovery may lead consumers to question not only the authenticity of the specific content, but also the authenticity of the brand's entire sustainability commitment. This creates what we might call "compounded skepticism" skepticism about both the authenticity of the communication and the authenticity of the underlying commitment.

Connection to the Authenticity Dilution Effect. The Authenticity Dilution Effect poses particular risks for sustainable brands. As AI-generated sustainability claims proliferate across the market, consumers' ability to distinguish genuine environmental commitment from algorithmically generated virtue signaling diminishes. This creates a paradox: individual sustainable brands may feel competitive pressure to adopt AI-generated content to maintain visibility, yet the collective effect of such adoption accelerates the dilution of sustainability authenticity across the entire sector. For brands whose competitive advantage rests on genuine environmental commitment, the Authenticity Dilution Effect threatens to commoditize sustainability claims, making it increasingly difficult for truly committed brands to signal their authenticity credibly. The strategic implication is that sustainable brands should view restraint in AI use not merely as a risk mitigation strategy, but as a potential source of competitive differentiation in an increasingly AI-saturated sustainability marketplace.

The strategic recommendation for sustainable brands is that they should adopt a *Low AI Intensity, High Transparency* positioning with particular emphasis on human involvement in sustainability messaging. This means: (1) Using AI selectively for non-sustainability functions

(e.g., product recommendations, routine administrative tasks); (2) Maintaining strict human oversight of all sustainability-related communications; (3) Being transparent about any AI use in marketing, with particular emphasis on transparency regarding sustainability claims; (4) Emphasizing genuine human commitment to sustainability; (5) Using third-party verification and certification to support sustainability claims; (6) Avoiding AI-generated content for sustainability storytelling, impact reporting, or values communication.

This approach allows sustainable brands to benefit from AI efficiency in non-critical areas while protecting the authenticity that is essential to their value proposition.

5. New Performance Metrics: Beyond Engagement

Traditional marketing metrics click-through rates, engagement, conversion rates do not capture the effects of synthetic influence on consumer trust and brand authenticity. Firms should develop new metrics that directly measure these constructs:

To monitor the effectiveness of their communication strategies, sustainable brands should regularly evaluate consumers' perceptions of authenticity. This can be achieved through survey items such as "This brand communicates authentically with me" or "I believe this brand is genuine and honest." Tracking these perceptions over time allows firms to identify shifts in consumer attitudes and assess whether their communication continues to be perceived as genuine. Another important indicator is trust resilience, which refers to the extent to which consumer trust remains stable when the brand encounters criticism or negative information. Unlike current trust levels, trust resilience reflects a brand's ability to maintain consumer confidence during challenging situations, making it a valuable measure of long-term brand strength.

Firms should also assess the effectiveness of the authenticity signals they employ. For example, when customer testimonials are incorporated into marketing campaigns, their impact on consumer trust and purchase intentions should be evaluated. If these signals no longer enhance credibility because of authenticity dilution, firms can revise their communication strategies to emphasize more effective forms of engagement. Finally, organizations should monitor consumer skepticism toward AI-generated content and synthetic influence. This can be measured through surveys that examine consumers' awareness of AI in marketing, their concerns about AI-generated communications, and their preferences for human-created versus AI-generated content. An increase in skepticism may indicate growing authenticity dilution and highlight the need to adapt communication strategies to preserve consumer trust.

5.1 Measurement Guidance for New Metrics

Measurement guidance for new metrics:

- **Perceived Authenticity.** Use survey items on a 10-point scale:
 - "This brand communicates authentically with me"
 - "I believe this brand is genuine and honest"
 - "This brand's communications feel real and not artificial"

Combine into composite authenticity score (average of items). Track quarterly.

- **Trust Resilience.** Measure trust before and after exposure to negative information or competitive challenge. Calculate resilience as: $(\text{Trust After} / \text{Trust Before}) \times 100$. Example: If trust is 8.0 before challenge and 7.0 after, resilience is 87.5%.

- **Authenticity Signal Effectiveness.** For each authenticity signal (e.g., customer testimonials, founder stories), measure consumer response using items such as "This [signal] increased my trust in the brand" (1-10 scale). Track changes over time to identify which signals are losing effectiveness.
- **Consumer Skepticism Index.** Use survey items:
 - "I am skeptical of AI-generated marketing content"
 - "I worry that brands are using AI to manipulate me"
 - "I can usually tell when content is AI-generated"

Combine into skepticism index and track over time

Conclusion

In this paper, the concept of synthetic influence in marketing has been presented, which involves the use of algorithm-generated persuasion stimuli to give the impression of human authenticity in order to influence consumer perception and behavior. The paper has constructed a conceptual model based on three interrelated components, the Generation Layer (the mechanism through which firms create algorithmic-based stimuli), the Perception Layer (how consumers understand and perceive authenticity cues) and the Outcome Layer (the outcome which includes consumer trust formation/eruption). An important part of the conceptual model is the 'Authenticity Dilution Effect', which suggests that excessive generation of artificial stimuli can lead to a saturation point where artificial authenticity loses credibility. In this case, an interesting collective action problem emerges in which the interests of each individual firm (to increase the use of AI content to achieve marketing goals) diverge from the interests of the collective (where everyone is better off if fewer firms engage in such activities). Strategic implications for organizations that have been suggested through the paper include the necessity to manage AI intensity, retain human anchors for their content, use AI selectively for personalization as opposed to simulations and calibrate levels of trust instead of automating their processes. It has also suggested new measures for performance, emphasizing perceptual aspects of authenticity and resilience as opposed to the more common metrics for engagement. From the perspective of entrepreneurship and SMEs, the framework suggests that in an environment in which AI simulates authenticity, such initiatives may destroy a firm's competitive advantage over other organizations that utilize AI. Sustainable brands will also face new challenges associated with increased skepticism among consumers, especially regarding greenwashing that involves AI-generated content. Contributions of this paper to marketing and management theory include the shift in conceptualizing authenticity as a characteristic of a brand as opposed to a perception of such. With the advent of AI generated content, the latter aspect of authenticity poses new research questions on the part of organizations. Second, it distinguishes the concept of synthetic influence as an independent category that necessitates new theory and strategies to deal with the problem. While past studies have considered the factors related to adopting AI for marketing and the impact of such applications on marketing efficiency, this study is concerned with the issues related to the perception and influence processes which play a major role in shaping brand value. Third, it draws attention to the issue of Authenticity Dilution Effect, which can help explain the problem of how the collective adoption of artificial intelligence content can affect its influence on customers. Fourth, it underscores the importance of trust calibration over automation maximization. The current

logic underlying marketing automation should be questioned since brand value in the AI era is more likely to emerge from the limited adoption of artificial intelligence coupled with appropriate disclosure and human participation in trust-related activities.

This framework has several important limitations that should be acknowledged. First one is conceptual in nature and has not yet been empirically validated. Consequently, its propositions should be tested through future research using methods such as experiments, surveys, and field studies. In addition, the framework was developed primarily within a Western marketing context, which may limit its applicability across different cultures where attitudes toward AI, authenticity, and trust can vary. Another limitation is that the framework reflects current generative AI capabilities; as these technologies continue to evolve, some assumptions and propositions may need to be revised. The framework also does not fully capture individual differences among consumers, such as variations in AI literacy, skepticism, and trust propensity, which may influence responses to synthetic influence. Finally, while the framework highlights important strategic implications, it provides only limited consideration of the ethical challenges associated with AI-driven marketing. Future research should therefore examine these ethical issues and contribute to the development of responsible guidelines for the use of AI in marketing.

Several avenues for future research emerge from this framework. First, its propositions should be empirically validated through experiments, surveys, and field studies to better understand how synthetic influence shapes consumer trust and perceptions of authenticity. Future studies could, for example, compare consumer responses to content identified as human-authored versus AI-generated. Second, researchers should examine how the framework applies across different contexts, including industries, product categories, market settings (e.g., B2B versus B2C), and consumer segments with varying levels of AI literacy and skepticism. Finally, longitudinal studies are needed to investigate how authenticity dilution evolves over time as AI-generated content becomes more widespread and consumer awareness continues to grow.

This framework also raises several important research questions. First, *at what point does widespread adoption of AI-generated content lead to an irreversible decline in consumers' perceptions of authenticity?* Identifying this threshold would help firms anticipate when authenticity becomes difficult to restore. Second, *do the dynamics of authenticity dilution differ between B2B and B2C markets?* Because business buyers and individual consumers often evaluate information differently, the effects of synthetic influence may vary across these contexts. Finally, *can transparent disclosure of AI use strengthen consumer trust in trust-sensitive markets, or does it consistently reduce the persuasive impact of marketing communications?* Answering this question would clarify whether transparency can serve as a strategic advantage rather than simply fulfilling an ethical or regulatory obligation.

As synthetic influence becomes more prevalent, important ethical and regulatory questions arise. Future research should explore the ethical implications of synthetic influence, the appropriate regulatory frameworks for governing AI-generated content in marketing, and the rights and responsibilities of firms and consumers in the era of synthetic influence. As AI-generated content becomes more sophisticated, consumers and regulators will need tools to detect whether content is AI-generated. Future research should explore the development of

detection technologies and the implications of such technologies for the effectiveness of synthetic influence.

The rise of generative AI in marketing represents a fundamental shift in how brands communicate with consumers. The ability to generate highly personalized, emotionally resonant content at scale creates unprecedented opportunities for firms to influence consumer behavior. However, this same capability creates risks for consumer trust and brand authenticity if not managed carefully.

The framework presented in this paper provides a foundation for understanding these opportunities and risks, and for developing strategies that allow firms to benefit from AI-generated content while maintaining the trust and authenticity that are essential for long-term brand value. The key insight is that firms must move beyond the logic of automation maximization and instead adopt a trust calibration approach that uses AI selectively and transparently, maintains human involvement in trust-critical functions, and continuously monitors consumer perceptions of authenticity and trust.

For Practitioners: We recommend that marketing leaders move beyond the logic of automation maximization and adopt a trust calibration approach that uses AI selectively and transparently, maintains human involvement in trust-critical functions, and continuously monitors consumer perceptions of authenticity and trust. For entrepreneurs and SMEs, we recommend emphasizing authenticity as a competitive differentiator. For sustainable brands, we recommend particular caution with AI-generated sustainability messaging.

For Researchers: We encourage empirical testing of the framework's propositions, exploration of contextual variations across industries and cultures, investigation of the ethical implications of synthetic influence, and exploration of the specific research questions identified above.

For Policymakers: We recommend developing regulatory frameworks that require transparent disclosure of AI use in marketing while allowing firms flexibility in implementation. We particularly recommend attention to the unique risks faced by sustainable brands and the importance of preventing AI-enabled greenwashing.

As the marketing environment keeps changing, the capability of dealing with the problems and opportunities associated with synthetic influence will be the key source of competitive advantage. Companies that create complex frameworks for managing synthetic influence, establish genuine connections with customers, and adjust to increasing customer cynicism will be most likely to succeed in the world of marketing mediated by artificial intelligence.

References:

- Akerlof, G. A. (1970). *The market for 'lemons': Quality uncertainty and the market mechanism*. *Quarterly Journal of Economics*, 84(3), 488–500.
- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). *GEO: Generative Engine Optimization*. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, 5–16. <https://doi.org/10.1145/3637528.3671900>
- Beverland, M. B., & Farrelly, F. J. (2010). *The quest for authenticity in consumption: Consumers' purposive choice of authentic cues to shape experienced outcomes*. *Journal of Consumer Research*, 36(5), 838–856. <https://doi.org/10.1086/615047>
- Chaudhuri, A., & Holbrook, M. B. (2001). *The chain of effects from brand trust and brand affect to brand performance: The role of brand loyalty*. *Journal of Marketing*, 65(2), 81–93. <https://doi.org/10.1509/jmkg.65.2.81.18255>
- Davenport, T. H., Guha, A., Grewal, D., & Bressgott, T. (2020). *How artificial intelligence will change the future of marketing*. *Journal of the Academy of Marketing Science*, 48, 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- Desveaud, K., & Pavone, G. (2025). *Consumer perceptions of AI-generated marketing content*. In *Encyclopedia of Artificial Intelligence in Marketing* (pp. 1–7). Springer. https://doi.org/10.1007/978-3-031-75316-9_94-1
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). *Algorithm aversion: People erroneously avoid algorithms after seeing them err*. *Journal of Experimental Psychology: General*, 144(1), 114–126.
- Erdem, T., & Swait, J. (2004). *Brand credibility, brand consideration, and choice*. *Journal of Consumer Research*, 31(1), 191–198. <https://doi.org/10.1086/383434>
- Felzmann, H., Fosch-Villaronga, E., Lutz, C., & Tamò-Larrieux, A. (2019). *Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns*. *Big Data & Society*, 6(1), 1–14.
- Friestad, M., & Wright, P. (1994). *The persuasion knowledge model: How people cope with persuasion attempts*. *Journal of Consumer Research*, 21(1), 1–31.
- Grayson, K., & Martinec, R. (2004). *Consumer perceptions of iconicity and indexicality and their influence on assessments of authentic market offerings*. *Journal of Consumer Research*, 31(2), 296–312. <https://doi.org/10.1086/422109>
- Han, S. (2024). *Consumer perceptions of AI-generated content and disclaimer in terms of authenticity, deception, and content attribute*. *24th Biennial Conference of the International Telecommunications Society (ITS)*, Seoul, Korea, 23–26 June 2024. <https://hdl.handle.net/10419/302503>
- Jaakkola, E. (2020). *Designing conceptual articles: Four approaches*. *AMS Review*, 10(1), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Kim, J. H., & Song, H. (2020). *The influence of perceived credibility on purchase intention via competence and authenticity*. *International Journal of Hospitality Management*, 88, 102539.
- Nuttavuthisit, K., & Thøgersen, J. (2017). *The importance of consumer trust for the emergence of a market for green products: The case of organic food*. *Journal of Business Ethics*, 140, 323–337. <https://doi.org/10.1007/s10551-015-2690-5>

- Poulis, A., Theodoridis, P., & Zacharatos, T. (2026). *Green branding in the digital era: The role of influencer credibility and greenwashing in shaping brand authenticity, trust and purchase intentions*. *Sustainability*, 18(1), 451. <https://doi.org/10.3390/su18010451>
- Shin, D., & Park, Y. J. (2019). *Role of fairness, accountability, and transparency in algorithmic affordance*. *Computers in Human Behavior*, 98, 277–284.
- Spence, M. (1973). *Job market signaling*. *Quarterly Journal of Economics*, 87(3), 355–374. <https://doi.org/10.2307/1882010>
- Sundar, S. S. (2008). *The MAIN model: A heuristic approach to understanding technology effects on credibility*. In M. J. Metzger & A. J. Flanagin (Eds.), *Digital Media, Youth, and Credibility* (pp. 73–100). MIT Press.
- Wortel, C., Vanwesenbeeck, I., & Tomas, F. (2024). *Made with Artificial Intelligence: The effect of artificial intelligence disclosures in Instagram advertisements on consumer attitudes*. *Emerging Media*, 2(3), 547–570. <https://doi.org/10.1177/27523543241292096>
- Yang, J., & Battocchio, A. F. (2021). *Effects of transparent brand communication on perceived brand authenticity and consumer responses*. *Journal of Product & Brand Management*, 30(8), 1176–1189. <https://doi.org/10.1108/JPBM-03-2020-2803>
- Zsigmondová, A., Zsigmond, T., & Machova, R. (2021). *Theoretical background to the role of trust in marketing*. *SHS Web of Conferences*, 115, 03019. <https://doi.org/10.1051/shsconf/202111503019>