
THE ETHICS OF DEEPFAKE MARKETING: CONSUMER TRUST IN THE ERA OF SYNTHETIC MEDIA

***A. Amala, Dr. Chakkaravarthy Kumaresan R.**

Assistant Professor, Department of MBA CARE College of Engineering, Tiruchirappalli,
Tamil Nadu, India.

Article Received: 03 July 2026

***Corresponding Author: A. Amala**

Article Revised: 23 July 2026

Assistant Professor, Department of MBA CARE College of Engineering, Tiruchirappalli,
Tamil Nadu, India.

Published on: 10 August 2026

DOI: <https://doi-doi.org/101555/ijrpa.2648>

ABSTRACT

The proliferation of generative artificial intelligence has made synthetic media—commonly termed “deepfakes”—technically trivial to produce and increasingly common in commercial communication. Brands now use AI-generated spokespeople, synthetic voice clones, digitally de-aged or resurrected celebrities, and hyper-personalized synthetic video advertisements to engage consumers at scale. While these technologies offer marketers unprecedented creative flexibility and cost efficiency, they also raise fundamental ethical questions about honesty, consent, autonomy, and the erosion of the perceptual boundary between the real and the fabricated. This paper develops an integrative conceptual framework for understanding how deepfake marketing practices affect consumer trust. Drawing on trust theory, the Elaboration Likelihood Model, media richness theory, and emerging AI ethics scholarship, the paper proposes a multi-dimensional model in which disclosure transparency, perceived manipulative intent, source credibility, and regulatory context jointly determine whether synthetic media marketing sustains or corrodes consumer trust. The paper further examines the current regulatory landscape, offers illustrative hypothetical scenarios, and concludes with practical recommendations for marketers, platforms, and policymakers, alongside directions for future empirical research.

KEYWORDS: deepfake marketing, synthetic media, consumer trust, AI ethics, digital advertising, disclosure, generative AI

1. INTRODUCTION

Synthetic media technologies—broadly referred to as deepfakes—use generative artificial intelligence to fabricate or manipulate audio, video, and images so convincingly that the output is difficult to distinguish from authentic recordings. Originally associated with disinformation and non-consensual imagery, this technology has migrated rapidly into the commercial sphere. Marketers now deploy AI-generated brand ambassadors, synthetic voice clones for personalized customer outreach, digitally de-aged actors, and even AI recreations of deceased celebrities to endorse products. What was once a niche technical curiosity has become a mainstream marketing tool, supported by a growing ecosystem of generative AI platforms marketed explicitly to advertising agencies and brand teams.

This paper treats the phrase “deepfake marketing” as an umbrella term covering any commercial communication in which synthetic media is used to represent a person, voice, or event in a way that a reasonable consumer could mistake for authentic, unmediated reality. This includes synthetic spokespeople, voice cloning, digital resurrection of public figures, hyper-personalized AI-generated video ads, and the manipulation of authentic footage to alter statements or actions.

The central problem this paper addresses is not whether synthetic media can be used in marketing—that question has already been settled by market practice—but under what conditions its use is ethically defensible and what its use does to the foundation of the marketer-consumer relationship: trust. Trust is the currency of persuasion. Advertising has always involved a degree of stylization, aspiration, and selective presentation, and consumers have long understood advertising as a persuasive genre rather than a neutral information source. Deepfake technology, however, does something qualitatively different: it does not merely dramatize reality, it can fabricate the appearance of a real, specific, identifiable act of communication that never took place. This paper argues that this qualitative shift demands a distinct ethical and conceptual vocabulary, one this paper aims to supply.

The paper proceeds as follows. Section 2 defines key concepts and surveys current commercial applications of synthetic media. Section 3 establishes the theoretical foundations drawn from trust theory, persuasion research, and media theory. Section 4 catalogues contemporary use cases. Section 5 develops the core ethical analysis across five dimensions: deception and authenticity, consent and likeness rights, manipulation and autonomy, misinformation and social trust erosion, and impacts on creative labor. Section 6 proposes an integrative conceptual model of consumer trust in synthetic media marketing. Section 7 reviews the regulatory and self-regulatory landscape. Section 8 presents illustrative

hypothetical scenarios to ground the framework. Section 9 offers recommendations, and Section 10 outlines directions for future research before the paper concludes in Section 11.

2. Conceptual Background: Defining Synthetic Media and Deepfake Marketing

2.1 From Deepfakes to Synthetic Media

The term “deepfake” originated from a portmanteau of “deep learning” and “fake,” describing AI-generated face-swap videos that emerged in online communities in the mid-2010s. As the underlying generative adversarial network and diffusion model architectures matured, the scope of what could be synthesized expanded far beyond face-swapping to include fully synthetic voices, photorealistic human avatars with no underlying human referent, and text-to-video systems capable of generating entire scenes from a written prompt. The industry increasingly prefers the more neutral term “synthetic media” to capture this broader category, reserving “deepfake” for the subset of synthetic content that depicts a specific real person doing or saying something they did not actually do or say.

For the purposes of this paper, deepfake marketing is defined as the commercial use of AI-generated or AI-manipulated audio, image, or video content that portrays a real or realistic human likeness in a manner that could reasonably be mistaken for an authentic, unaltered recording, deployed for the purpose of promoting a brand, product, or service.

2.2 A Taxonomy of Commercial Synthetic Media

It is useful to distinguish several categories that fall under the deepfake marketing umbrella, since they carry different ethical weight:

- Wholly synthetic spokespeople: AI-generated humans with no real-world referent, used as virtual brand ambassadors or product presenters.
- Likeness-based synthetic endorsers: AI recreations of a specific living celebrity, generated with that person's authorization and typically compensated under a likeness-licensing agreement.
- Posthumous digital resurrection: AI recreations of deceased public figures, authorized by an estate, used to endorse or appear in advertising.
- Voice cloning and localization: synthetic replication of a real speaker's voice, often used to translate or personalize spoken advertising content at scale.
- Hyper-personalized synthetic video: mass-customized video advertisements in which a synthetic presenter's script, tone, or even appearance is algorithmically adjusted per viewer.

- Manipulated authentic footage: real footage of a real person that has been altered to change words, context, or apparent endorsement, typically without that person's consent.
- The first five categories can, in principle, be conducted transparently and consensually; the sixth is inherently non-consensual and forms the clearest case of unethical practice. Much of the ethical complexity in this paper concerns the first five categories, where consent and disclosure exist on a spectrum rather than a binary.

3. Theoretical Foundations

3.1 Trust Theory

Organizational trust research conceptualizes trust as a willingness to be vulnerable to another party based on positive expectations of that party's intentions and behavior. This conceptualization identifies three antecedents of perceived trustworthiness: ability, benevolence, and integrity. Applied to marketing communication, integrity is the most directly implicated antecedent when synthetic media is involved, because integrity concerns whether the communicator adheres to principles the audience finds acceptable, including honesty about the nature of the communication itself. When a consumer discovers that a seemingly authentic endorsement was synthetic and undisclosed, the injury is not merely to the specific claim being made but to the perceived integrity of the communicating brand as a whole, an injury that can generalize to distrust of that brand's future communications.

3.2 The Elaboration Likelihood Model and Source Credibility

Dual-process theories of persuasion distinguish between a central route, in which audiences scrutinize argument quality, and a peripheral route, in which audiences rely on heuristic cues such as source attractiveness or credibility. Synthetic spokespeople are, by construction, peripheral-route stimuli optimized for likeability and credibility signals without the constraints of a real human's availability, aging, or personal conduct risk. This creates an asymmetry: synthetic endorsers can be engineered to maximize peripheral persuasive cues in ways unconstrained by authentic personal history, which raises the ethical question of whether such engineered credibility constitutes a legitimate persuasive appeal or a manipulative exploitation of heuristic processing.

3.3 Media Richness and the Authenticity Heuristic

Media richness theory holds that richer channels—those closer to face-to-face interaction—carry more social presence and are processed with greater perceived authenticity. Video and voice have historically served as high-trust channels precisely because they were difficult to fabricate convincingly; audiences developed an implicit authenticity heuristic in which

“seeing is believing.” Generative AI directly undermines this heuristic by making the richest media channels the easiest to falsify. This paper terms this the authenticity heuristic breakdown: as synthetic video and voice approach perceptual indistinguishability from authentic recordings, the very channels that consumers have relied upon most heavily for credibility judgments become the least reliable indicators of truth.

3.4 The Uncanny Valley and Consumer Discomfort

Robotics research identified a nonlinear relationship between the human-likeness of an artificial agent and observers' affective response: as artificial humanlikeness increases, affinity rises, then drops sharply into discomfort as the agent approaches but does not perfectly achieve realism, before rising again once realism is complete. This uncanny valley effect is relevant to synthetic spokespeople, whose perceived trustworthiness may depend not only on disclosure but on where a given rendering falls along the realism spectrum. Ethically salient marketing effectiveness and consumer comfort are therefore both functions of technical execution quality, complicating any simple rule that treats all synthetic media identically.

3.5 Synthesis

Taken together, these frameworks suggest that consumer trust in deepfake marketing is not a single variable but an emergent property of at least four interacting factors: the integrity signal sent by disclosure practices, the persuasive route being exploited, the erosion of the authenticity heuristic across the category as a whole, and the affective response to the specific execution. Section 6 formalizes these interactions into a conceptual model.

4. Contemporary Applications of Deepfakes in Marketing

Commercial adoption of synthetic media has accelerated across several domains. Virtual influencers—computer-generated personas with their own social media presence, personality, and brand partnerships—have moved from novelty to a recognized marketing category, appealing to brands seeking full creative control and freedom from the reputational risk of human talent. Voice cloning is used to localize advertising and customer service scripts into dozens of languages while preserving a single brand voice, and to generate personalized outbound messages at a scale no human voice actor could match. De-aging and posthumous recreation technology has allowed brands and estates to feature historical or deceased celebrities in new campaigns, raising distinctive consent questions addressed in Section 5.2. Retailers and platforms increasingly experiment with hyper-personalized synthetic video, in which a single script is delivered by an AI presenter whose tone, pacing, or even apparent

identity is varied by viewer segment. Finally, some political and issue-advocacy actors have begun applying deepfake techniques adapted from marketing workflows to advocacy messaging, blurring the line between commercial and political synthetic persuasion.

Each of these applications carries a different trust profile. Fully disclosed virtual influencers, whose synthetic nature is a known and even celebrated part of their brand identity, present relatively low deception risk. Undisclosed hyper-personalization and manipulated authentic footage present the highest risk, since they exploit the audience's default assumption that video content is an unmediated record of an actual event.

5. Ethical Dimensions of Deepfake Marketing

5.1 Deception and Authenticity

The most immediate ethical concern is deception: does the communication cause a reasonable consumer to form a false belief about what actually occurred or who is actually endorsing a product? Classical advertising ethics has long tolerated puffery and stylization on the theory that consumers apply a discount to obviously aspirational claims. Deepfake marketing complicates this theory because synthetic media is designed precisely to defeat the perceptual cues consumers use to identify stylization. A synthetic video that looks exactly like an authentic recording does not signal “this is marketing artifice” in the way a cartoon mascot or an obviously staged commercial does. The ethical weight of a given execution therefore turns heavily on disclosure: clearly labeled synthetic content preserves the consumer's ability to apply an appropriate discount, while undisclosed synthetic content removes that ability entirely.

5.2 Consent and Likeness Rights

A second dimension concerns whose likeness is used and under what authority. Living celebrities can, in principle, negotiate specific contractual terms governing how their AI likeness may be used, for how long, and in what contexts—terms that a real-time human endorsement could never specify with comparable precision, since a cloned likeness can be redeployed indefinitely across contexts the original person never anticipated. Posthumous digital resurrection raises a further layer of concern: the deceased person cannot consent in real time to how their likeness is used in new campaigns, and estate-level authorization, while legally sufficient in most jurisdictions, may not fully substitute for the individual's own contemporaneous judgment about brand alignment. A distinct and unambiguously unethical category involves the use of an ordinary person's likeness, image, or voice without any consent at all—for example, scraping social media images to generate synthetic

testimonials—which implicates both consent violations and, in many jurisdictions, existing rights of publicity.

5.3 Manipulation and Consumer Autonomy

Beyond factual deception, synthetic media raises manipulation concerns distinct from truth-telling. Manipulation, in the ethical sense used here, refers to influence that bypasses a person's rational agency rather than appealing to it—for instance, by engineering a synthetic spokesperson's facial expressions, vocal cadence, or apparent emotional sincerity to maximize affective response independent of the substantive merits of the product being sold. Because generative systems can be optimized against engagement or conversion metrics at a level of granularity unavailable to human performers, deepfake marketing raises the prospect of persuasion techniques calibrated to exploit specific psychological vulnerabilities of individual viewers, particularly when combined with behavioral data used for hyper-personalization.

5.4 Misinformation Spillover and the Erosion of Generalized Trust

A distinctive feature of deepfake marketing ethics is that its harms are not fully contained to the transaction between a single brand and its customers. Every instance of synthetic media in commercial use, whether disclosed or not, contributes to a general public awareness that video and audio can be fabricated. This contributes to what might be called the liar's dividend: as audiences learn that any video could be fake, bad actors gain a new defense against authentic incriminating recordings, since they can now plausibly claim any real recording is itself a deepfake. Marketing's widespread, mainstream normalization of synthetic media therefore carries a societal externality beyond the immediate advertiser-consumer relationship, contributing to a broader epistemic environment in which the evidentiary value of audiovisual recording is diminished for everyone, including in domains far removed from commerce, such as journalism, law, and political accountability.

5.5 Impact on Creative Labor and Market Fairness

Finally, deepfake marketing raises labor and fairness concerns for the voice actors, models, and performers whose work synthetic systems increasingly substitute for or are trained on. Voice cloning and synthetic likeness generation are frequently trained on datasets that include the recorded work of professional performers, raising questions about compensation, ongoing consent, and the long-term economic displacement of creative professions that have historically supplied the “real” in advertising's promise of authenticity. Industry labor negotiations in adjacent creative sectors have already begun to address AI likeness and voice

rights explicitly, reflecting the practical urgency of this dimension alongside its ethical significance.

6. Toward a Conceptual Model of Consumer Trust in Synthetic Media Marketing

Building on the theoretical foundations in Section 3 and the ethical dimensions in Section 5, this paper proposes an integrative model in which consumer trust outcomes are jointly determined by four constructs: Disclosure Transparency, Perceived Manipulative Intent, Source-Content Congruence, and Regulatory-Normative Context.

6.1 Disclosure Transparency

Disclosure transparency refers to the clarity, prominence, and timing of any signal that content is AI-generated or AI-altered. This construct is proposed as the strongest direct predictor of trust outcomes: prominent, timely, and comprehensible disclosure is expected to preserve trust even when synthetic media is used extensively, while the discovery of undisclosed synthetic content after the fact is expected to produce the sharpest trust penalties, both for the specific campaign and for the brand's broader credibility.

6.2 Perceived Manipulative Intent

Independent of disclosure, consumers form judgments about why synthetic media was used. Synthetic media adopted for accessibility, localization, or creative reasons unrelated to concealing brand attributes is expected to be judged more favorably than synthetic media adopted specifically to simulate an endorsement, emotional reaction, or claim that could not have been authentically obtained. This construct captures the manipulation concern from Section 5.3 as a perceptual variable: what matters for trust is not only whether manipulation occurred in some objective sense, but whether consumers infer manipulative intent from the campaign's design.

6.3 Source-Content Congruence

This construct captures the fit between the synthetic execution and the category or brand context in which it appears. A virtual influencer promoting a technology product may be perceived as congruent and even novel, while an AI-recreated deceased celebrity endorsing an unrelated modern product may be perceived as incongruent, producing discomfort independent of disclosure adequacy. Congruence judgments are expected to interact with the uncanny valley effect described in Section 3.4, such that technically imperfect synthetic executions in low-congruence contexts are especially vulnerable to negative consumer response.

6.4 Regulatory-Normative Context

Finally, consumer trust responses are shaped by the broader regulatory and social-normative environment. In markets with mandatory AI-disclosure regulation and high public awareness of synthetic media, consumers may apply a more forgiving baseline to disclosed synthetic content while penalizing non-disclosure more severely, since non-disclosure in a regulated environment signals an active choice to violate a known norm rather than an oversight in an ambiguous one.

6.5 Model Summary

The model predicts that trust is highest when synthetic media is prominently disclosed, used for a purpose consumers judge as legitimate rather than manipulative, executed in a way congruent with the brand and technically clear of uncanny-valley discomfort, and situated within a regulatory environment that normalizes appropriate use while sanctioning concealment. Conversely, trust is most severely damaged not by the mere use of synthetic media, but by the combination of concealment and inferred manipulative intent—the conjunction of low disclosure transparency and high perceived manipulative intent constitutes, in this framework, the highest-risk condition for both the individual brand and the broader marketing ecosystem.

7. Regulatory and Industry Landscape

7.1 Regulatory Approaches

Regulatory responses to synthetic media in commercial contexts have generally taken one of three forms: disclosure mandates, which require labeling of AI-generated or AI-altered content without necessarily restricting its use; risk-based restrictions, which impose heightened obligations on synthetic media applications deemed high-risk, such as biometric-adjacent uses; and sectoral rules targeting specific harms, such as the use of synthetic media in political advertising or in impersonating public officials. Consumer protection regulators in multiple jurisdictions have signaled that existing prohibitions on deceptive advertising practices apply to undisclosed synthetic endorsements without requiring new legislation, treating a fabricated endorsement as a species of false testimonial regardless of the technology used to produce it.

7.2 Platform and Industry Self-Regulation

Major advertising platforms and industry self-regulatory bodies have begun introducing AI-content labeling requirements and policies restricting synthetic depictions of real individuals without consent. Talent and performer unions have negotiated specific contractual protections

governing digital replica use, informed consent, and compensation, reflecting the labor concerns raised in Section 5.5. Advertising industry trade associations have likewise issued voluntary guidelines recommending disclosure of synthetic media, though compliance remains uneven and largely unenforced outside of the largest brands and agencies.

7.3 Gaps and Tensions

Notable gaps remain. Disclosure requirements vary widely in specificity, from general labeling obligations to no requirement at all, creating inconsistent consumer expectations across markets and platforms. Enforcement resources for consumer protection regulators remain limited relative to the scale of AI-generated content now in circulation. And likeness rights frameworks, largely developed before generative AI existed, do not uniformly address scenarios such as posthumous digital resurrection or the training of voice models on publicly available recordings without the speaker's specific advance knowledge. These gaps mean that, in practice, ethical restraint by individual marketing organizations currently does more work than binding law in determining whether a given deepfake marketing execution respects consumer trust and creative labor interests. This gap between the current pace of regulation and the pace of technological deployment is a recurring theme this paper returns to in its recommendations.

8. Illustrative Scenarios

The following hypothetical scenarios are constructed for analytical purposes only and do not describe any actual company, campaign, or individual. They are offered to illustrate how the conceptual model in Section 6 applies in practice.

8.1 Scenario A: The Disclosed Virtual Ambassador

A beauty brand launches a fully synthetic virtual influencer, openly marketed as an AI creation with its own social media biography describing it as computer-generated. The brand uses the character consistently across campaigns and never implies the character is a real person. Applying the model: disclosure transparency is high, perceived manipulative intent is low because the synthetic nature is part of the brand story rather than a concealment, and source-content congruence is high given the novelty-forward beauty and technology category. The model predicts minimal trust erosion and potential trust-neutral or even trust-positive reception as a form of transparent creative innovation.

8.2 Scenario B: The Undisclosed Personalized Testimonial

A financial services company generates thousands of personalized synthetic video testimonials, each showing a different synthetic “customer” describing satisfaction with the

product, without disclosing that no such customers exist. Applying the model: disclosure transparency is at its lowest possible level, perceived manipulative intent is high once discovered, since fabricated testimonials are a category regulators have long treated as a paradigm case of deceptive advertising, and the financial-services context heightens the trust stakes given the sector's reliance on credibility. The model predicts severe trust damage upon discovery, extending beyond the specific campaign to the brand's broader credibility, consistent with the integrity-based mechanism described in Section 3.1.

8.3 Scenario C: The Licensed Posthumous Endorsement

An estate authorizes a technology company to use AI recreation of a deceased celebrity's voice and likeness in a new product campaign, with clear on-screen disclosure that the appearance is an authorized AI recreation. Applying the model: disclosure transparency is high and legal consent exists at the estate level, but source-content congruence and perceived manipulative intent become the pivotal variables—if the product aligns with the celebrity's known public associations, congruence is high and reception is likely favorable; if the product is unrelated or the recreation is used to imply an endorsement view the celebrity never held in life, perceived manipulative intent rises sharply even with full disclosure, illustrating that disclosure alone does not neutralize all ethical concerns identified in Section 5.2.

9. Implications and Recommendations

9.1 For Marketers and Brands

- Adopt clear, prominent, and persistent disclosure of synthetic media rather than minimal or buried labeling, treating disclosure as a trust-building asset rather than a legal formality to be minimized.
- Establish internal review processes that assess not only legal consent but perceived manipulative intent and source-content congruence before deploying synthetic spokespeople or personalization at scale.
- Negotiate specific, time-bound, and context-bound likeness agreements with living talent, and apply comparably rigorous estate-consultation standards for posthumous use rather than relying solely on minimal legal sufficiency.

9.2 For Policymakers and Regulators

- Pursue harmonized disclosure standards specifying label prominence, timing, and required content, reducing the current inconsistency across jurisdictions and platforms identified in Section 7.3.
- Extend rights-of-publicity and likeness-protection frameworks explicitly to cover AI voice cloning and digital replica generation, closing gaps left by pre-generative-AI legal frameworks.
- Resource consumer protection enforcement proportionate to the scale of AI-generated commercial content now entering circulation.

9.3 For Platforms and Industry Bodies

- Implement machine-readable provenance standards so that disclosure travels with content across platforms rather than being stripped upon redistribution.
- Extend collectively bargained digital replica protections across the broader creative labor market beyond the specific unions that have already negotiated them.

10. Directions for Future Research

This paper's conceptual model generates several testable propositions for future empirical work. Experimental studies could manipulate disclosure prominence, inferred purpose, and execution realism independently to test their proposed interactive effects on trust and purchase intention, directly testing the model developed in Section 6. Longitudinal research could examine whether repeated exposure to disclosed synthetic media normalizes consumer acceptance over time or instead produces cumulative erosion of the general authenticity heuristic described in Section 3.3. Cross-cultural research could examine whether regulatory-normative context moderates trust responses as the model predicts, comparing markets with strong disclosure mandates to those without. Finally, research on creative labor markets could examine the economic and psychological effects of digital replica proliferation on professional performers, complementing the ethical analysis in Section 5.5 with empirical labor-market evidence.

11. CONCLUSION

Deepfake marketing sits at the intersection of remarkable creative and commercial opportunity and serious ethical risk. This paper has argued that the central ethical stakes are not reducible to a single question of whether synthetic media should be permitted in

advertising, but instead turn on a cluster of interacting factors: how transparently synthetic content is disclosed, what purpose consumers infer behind its use, how congruent the execution is with its brand and category context, and the regulatory-normative environment in which it appears. The integrative model proposed here suggests that consumer trust is most robust when brands treat disclosure as a core creative and ethical commitment rather than a compliance afterthought, and most severely damaged by the combination of concealment and inferred manipulation. As generative AI capabilities continue to advance and the perceptual boundary between authentic and synthetic media continues to narrow, the marketing industry's choices about disclosure, consent, and restraint will shape not only individual brand reputations but the broader public capacity to trust audiovisual communication as such. The stakes of getting this right, in other words, extend well beyond marketing itself.

REFERENCES

1. Chesney, R., & Citron, D. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1820.
2. Federal Trade Commission. (2023). Guidance on the use of AI-generated endorsements and testimonials in advertising. Washington, DC: Federal Trade Commission.
3. Floridi, L. (2018). Artificial intelligence, deepfakes and a future of ectypes. *Philosophy & Technology*, 31(3), 317–321.
4. Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709–734.
5. Mori, M. (1970). The uncanny valley. *Energy*, 7(4), 33–35.
6. Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123–205.
7. Rini, R. (2020). Deepfakes and the epistemic backstop. *Philosophers' Imprint*, 20(24), 1–16.
8. Short, J., Williams, E., & Christie, B. (1976). *The social psychology of telecommunications*. London: Wiley.
9. European Commission. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.
10. Whittaker, L., Kietzmann, T. C., Kietzmann, J., & Dabirian, A. (2020). All around me are synthetic faces: The mad world of AI-generated media. *IT Professional*, 22(5), 90–99.