AI IN ADVERTISING

GOVERNANCE, REGULATION AND OTHER TROUBLES

BY PROFESSOR STEVEN SIDLEY
JBS, UNIVERSITY OF JOHANNESBURG, JULY 4TH, 2025
WITH CONTRIBUTIONS FROM KONRAD SHEK

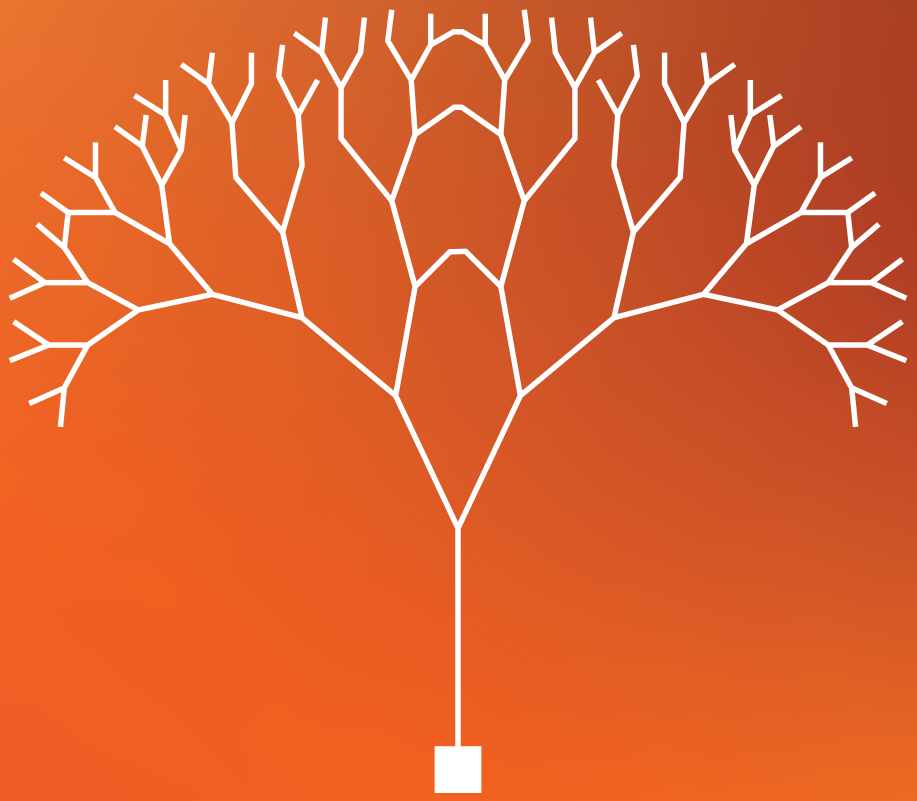

ICAS GLOBAL THINK TANK

Disclaimer

The AI Think Pieces developed by the ICAS Global Think Tank are authored by a diverse group of contributors, each bringing their unique perspectives, disciplinary expertise, and thematic focus. While each piece reflects the views and insights of its respective author(s), all contributions undergo a rigorous peer review process involving subject matter experts of the ICAS Global Think Tank. The analyses are grounded, where possible, in extensive desk research, existing studies, and available evidence.

Given the rapidly evolving nature of AI and of the technologies that use AI, these Think Pieces are not intended to provide definitive answers but rather to stimulate critical discussion, deepen understanding, and highlight areas for further inquiry. As the field progresses and knowledge gaps are identified, future Think Pieces will build on prior work and, where feasible, contribute to new research to advance evidence-based understanding.

About the ICAS Global Think Tank

On 18 September 2024, ICAS launched the Global Think Tank, a new platform dedicated to fostering self-regulatory engagement, critical thinking, research, and open, honest conversations. This multi-stakeholder initiative brings together advertising standards bodies, industry leaders, academic institutions, and other key stakeholders to address core issues, research self-regulatory best practices, and advance responsible advertising. Its goal is to generate high-quality insights and amplify our collective impact.

This is the first topic in the AI-focussed series of reports, with particular emphasis on transparency and the question of whether, when and how AI-generated ads should be labelled.

The insights are designed to help ICAS, self-regulatory organisations (SROs), industry stakeholders, and policymakers navigate the challenges and opportunities of AI in advertising while upholding ethical standards and fostering consumer trust.

About the author

Steven Sidley is Professor of Practice at JBS, University of Johannesburg in the fields of artificial intelligence and in crypto technologies and economies.

He is an award-winning and multi-shortlisted author of seven books (both fiction and non-fiction). His latest book It's Mine: How the Blockchain is Redefining Ownership was published in 2023.

Professor Steven Sidley was educated at the University of the Witwatersrand and the University of California, Los Angeles where he received his Msc in Computer Science. During his 17 years in LA he was one of the first wave of videogame developers and subsequently built and sold one of California's premier computer animation companies.

He returned to South Africa where he directed TV commercials and went on to co-found multiple technology companies, including Prism Payment Technologies which was publicly listed on the JSE, and Agilitude, which was sold to PwC.

He has held worldwide Group Chief Technology Officer positions at Anglo American and Altech/Altron and was CTO on one of SA's first blockchain projects, Project UBU.

Sidley is a widely published columnist and is currently a partner at Bridge Capital, a corporate advisory.

CONTENT
[AI IN ADVERTISING]

Prologue

This 'Think Piece' represents the first in a series of reports designed to provide a comprehensive, 360-degree view of the opportunities and challenges AI presents across policy, practice, and public understanding across our industry.

In this report I provide a broad landscape view of AI and advertising, while also delving into to some of the richer and occasionally more contested areas (such as labelling, targeting and dark patterns). While offering facts, opinions and recommendations, the purpose of this report is to foster engagement and discussion in the face of one of the most transformative technologies of our time.

Future reports will unpack other dimensions in greater detail, continuing to build a nuanced, evidence-informed body of work that supports thoughtful debate, decision-making and regulatory guidance.

1. Background and Context

As AI-fuelled applications start to emerge in every nook and cranny of commercial endeavour, the advertising industry finds itself faced with an urgent need to understand its benefits and threats, its hype and reality, in order to offer guidance, education, and policy/regulatory recommendations to its industry members in pursuance of a timely well-considered response to the explosive arrival of this transformative new technology.

AI has proliferated across the advertising ecosystem with unprecedented speed and breadth. It promises to transform every aspect of how advertisements are conceived, created, targeted, and delivered. The impact extends from creative processes to media buying, audience analysis, and campaign optimisation.

In the short space of time from the launch of ChatGPT in November 2022 (which dealt with text only), we have seen 'multimodal' AI systems that can generate photorealistic images from text prompts, create video content without cameras, write persuasive copy and target audiences with previously unimaginable precision. These capabilities offer tremendous opportunities for advertising professionals – reducing costs, creating novel and compelling content, accelerating production timelines, and enabling personalisation at scale.

But the rapid integration of AI into advertising raises questions about ethics – including responsibility, transparency and trust. The industry must navigate complex challenges related to authenticity, privacy, consent, and the potential for manipulation. When advertisements can be generated in seconds rather than days, when deepfakes become indistinguishable from reality, when provenance and authenticity is questioned by consumers, and when targeting becomes so precise it borders on invasive, we must pause to consider the implications.

This Think Piece aims to explore these tensions at the intersection of AI and advertising. It seeks to provide a framework for understanding not just what is possible, but what is responsible in terms of regulatory guidelines. The goal is not to hinder innovation but to ensure that as we embrace these powerful new tools, we do so in ways that maintain consumer trust, respect human dignity, and uphold the integrity of the advertising profession. The industry has supported these pillars in the past, through multiple technological changes; they now need to be strengthened against the onslaught of a powerful and fast-moving new technology.

The approach taken in this document balances technological understanding with ethical considerations. It draws on emerging global regulations, industry best practices, and ethical frameworks and cases to provide insights for advertisers, agencies, platforms, and policymakers. Rather than offering simplistic prescriptions, this Think Piece aims to foster informed dialogue about how the advertising industry can harness AI's potential while mitigating its risks.

Given the dizzying speed of innovation in the field we have a narrow window of opportunity to establish norms, standards, and guidelines that will shape its implementation for years to come. This document represents a contribution to that crucial conversation and will provide some of the raw material that will be used in the formulation of future self-regulatory policy recommendations.

2. A Matter of Definition

Dreams of AI have been around for thousands of years, captured comprehensively in writing, poetry, art and other creative endeavours and philosophical musings. But it only began its real trajectory in the 1950s, when mathematician Alan Turing, and others, turned their attention to the possibility of computers being programmed to exhibit real intelligence.

Attempts at creating AIs gained solid traction this century with the perfect storm of rapidly decreasing hardware, communication and storage costs, coupled with increasingly effective formalisms being developed out of fields like mathematics, statistics and neuroscience.

Today the term is widely bandied about, but there is little understanding as to what *distinguishes good ol' software from AI software*, even amongst professionals.

The understanding of the difference between the two is critical for the question of guidance and policy to the advertising industry, so without trying to dig too deeply into the technical weeds we offer two separate observations.

The first is that AI software learns over time. Hence the term machine learning, which is the mathematical foundation on which most modern AI sits. Sometimes that learning stops after it has been sufficiently trained (like the facial recognition software on your phone) and in other circumstances the learning never stops; the current AI arms race between OpenAI and others is based on this premise. Good ol' software does not learn – it is simply a set of rules. This 'learning' differentiator is unprecedented – humans have never built a technology that can learn autonomously or even semi-autonomously. We do not really understand the implications of that.

The second more critical difference is that traditional software takes a known input (such as a keyboard entry or photograph or temperature reading) and feeds it to a known set of rules called a deterministic algorithm (the computer program) to produce an output which is not necessarily known until the program is run.

AI turns this relationship on its head. It takes known input and known output and unleashes a computer to try billions of things until there is an 'acceptable' (or optimised) match between input and output. There is no known and deterministic 'program' sitting between input and output. This description obviously skips over many details, but the principle is profound, and it represents a new type of problem-solving ecosystem.

Worse, we do not know how in detail the software produces its output. It cannot be easily interpreted, because the original training step has executed too many (up to billions) of different steps and trials and combinations before the finish line – far too many to be unravelled after the fact. This is the problem of modern AI – we can't really peer too deeply under the hood; it is too complicated. This has all sorts of implications for us, like the well-known 'hallucinations', overconfidence and inherent bias, which we will come to later.

There is another matter of definition which needs to be understood. The public zeitgeist is largely focused on GenAI (of which Chat GPT is the prime example). But there are other types of machine learning that do not seek to generate new content but rather just look for patterns in existing data to optimise or make predictions. These too are finding widespread adoption in advertising. Taken together, we have, inter alia, the following applications of AI in advertising:

- Creative generation leverages GenAI to produce videos, audio, images, and copy
- Audience targeting utilises deep learning to identify potential customers based on complex behavioural patterns
- Predictive analytics uses machine learning to forecast consumer behaviour and campaign performance
- Campaign optimisation employs reinforcement learning to continuously improve ad performance

In summary: AI learns and gets better with time and with the ingesting of new data (often acting autonomously within specific domains like car sensors), whereas good ol' software doesn't. Secondly, different AI approaches are applicable in different segments of the advertising value chain. Some of these have ethics, regulatory, privacy and policy implications for the industry, others do not. This paper will concentrate on the former.

3. Ethics and Alignment and AI – A Broad View

In 1950, Isaac Asimov published a book of stories called I, Robot. In the story titled Runaround he defined three laws which were proposed to be applied to robots. Two of the three had to do with alignment of humans and robots – humans should not be harmed by action or inaction of robots, and robots must remain subservient to humans (the final one was that robots should protect their own existence).

Alignment of robots with human ethics and goals had long been debated well before publication of this short story, usually within the ambit of philosophy and other similar ruminations. But Asimov's short story brought the matter to the public consciousness. Where it remains (with considerably more urgency) as the technology has accelerated.

Regulation and policy discussions about the use of these technologies within advertising, intersect with discussions that have happened in other vertical sectors, as well as in the public policy sphere over the past many decades. But the advertising industry has a differently nuanced set of concerns that relate primarily to consumer trust (and harm), responsibility, prejudice, intellectual property and privacy, whereas in other magisteria different concerns may intrude, including national security and other geopolitical matters.

There is a considerable body of academic and policy-related research into general ethical concerns that are beyond the scope of this paper. However, we will list the major topics that have been deeply discussed and exercised in the specific areas of machine learning and GenAI.

They are:

- **Training Data Bias** Machine learning systems reflect the data used to train them. Biased, incomplete, or problematic training data leads to systems that perpetuate or amplify those problems. This raises questions about data collection practices, representation, and the historical biases encoded in existing datasets.
- **Algorithmic Bias** Even with unbiased training data, machine learning algorithms can develop discriminatory patterns through proxy variables or mathematical properties of the learning process. Detecting and mitigating these biases requires ongoing vigilance.

- **Privacy Considerations** The massive data requirements of modern AI systems create tension with individual privacy rights, particularly where the data is related to people, which may violate laws in certain jurisdictions e.g. the GDPR in the EU. The collection, storage, and processing of personal information raise concerns about consent, data security, and the potential for surveillance.
- **Explainability vs Performance** Many of the most powerful AI techniques produce results that are difficult or impossible to explain in human terms. This 'black box' nature creates challenges for accountability, trust, and regulatory compliance.
- **Provenance** There are many areas of human endeavour in which provenance, especially human provenance, represents a core value. This includes creative endeavours, news content, human communication, etc. The ability for AI to spoof humans creates a challenge to provenance and authenticity, and so 'proof-of-humanity' mechanisms are gaining traction.

While it is not a great leap to see that there is a considerable overlap of these issues with the matters of concern to the advertising industry, there are more specific and differently nuanced matters relating to consumer trust and transparency and consent that are differently 'weighted' in non-advertising AI ethics and alignment initiatives, which will covered in greater depth in this document.

As countries, institutions and enterprises have wrestled with these issues, we have seen the early adoption of regulations and/or guidelines, albeit quite fractured and not yet globally harmonised (and perhaps never to be so).

Among them are:

- **The European Approach** The European Union has taken the most comprehensive regulatory stance through its AI Act (August 2024), which adopts a risk-based framework categorizing AI applications according to their potential harm. The regulation prohibits certain 'unacceptable risk' applications while imposing stringent requirements on 'high-risk' systems, including those used in targeted advertising. The EU approach emphasises precaution, fundamental rights protection, and human oversight with a governing body, the European Artificial Intelligence Board, to ensure compliance.
- **The US Approach** The United States has favoured a more sectoral and principles-based approach to AI regulation. Rather than comprehensive legislation, the US has relied on existing regulatory authorities addressing specific domains (healthcare, financial services, etc.) supplemented by voluntary guidelines and executive orders. And in keeping with a gradual devolvement of responsibilities to states in the current political climate, we have the Colorado AI Act of 2024, with others in-process. In general, US regulation prioritises innovation while addressing particular risks as they emerge.
- **The Chinese Approach** China has developed a regulatory framework that balances innovation with national security and social stability concerns. Not unsurprising, given their priorities. Chinese regulations focus heavily on algorithm recommendation systems, particularly those used in information services and e-commerce, with specific provisions regarding user profiling and automated decision-making.
- **International Standards** Organisations such as the IEEE, ISO, and NIST (engineering and cybersecurity standards organisations) have developed voluntary technical standards for AI development and deployment. While lacking enforcement mechanisms, these standards provide frameworks for responsible AI design that may inform both regulation and industry self-governance. We also note some of the broader recent initiatives such as the March 2024 UN resolution titled 'Seizing the opportunities of safe, secure and trustworthy artificial intelligence systems for sustainable development' and the September 2024 Framework Convention on AI adopted by the Council of Europe.

With this context in mind we now focus on our industry, advertising.

4. Ethics in Advertising

by Konrad Shek

The four core principles of Legal, Decent, Honest and Truthful are at the heart of advertising self-regulatory bodies and the regulation of advertising content standards around the world. To clarify: 'Legal' means complying with all relevant laws and regulations; 'Decent' refers to respecting community standards of taste and appropriateness; 'Honest' means creating no false impressions through either claims or presentation; and 'Truthful' requires all factual assertions to be accurate and verifiable. These principles have shown remarkable resilience despite the technological shifts in the media landscape and the delivery of advertising.

These principles also represent what it means to be ethical in advertising. Fundamentally, ethics in advertising is grounded in the idea that advertising should not deceive or mislead consumers. Nor should it hide or omit important information. It should also reflect current societal norms – in other words be respectful to others and not display content that others would find offensive. Finally, it should not exploit consumers, in particular minors and those who are in a vulnerable state.

The opportunities presented by the latest technological developments have led the advertising industry to prioritise their investment in AI tools and capabilities. However, the industry still needs to be mindful of the risks and challenges outlined above when using AI, otherwise it could put brands at risk of breaching consumer trust, potentially attracting sanctions from advertising self-regulatory bodies and/or consumer protection authorities. It is important to balance these considerations to avoid an unnecessary tension between commercial interests and consumer trust.

Ethics in advertising directly impacts consumer trust, which can drive quantifiable business value. According to one survey, nearly 80% of millennial customers said they were driven to purchase a product because it was a trusted brand name.[1] According to research by the Institute of Practitioners in Advertising, the relationship between building trust in a brand and achieving greater profit has strengthened considerably.[2] As Keith Weed, former CMO for Unilever and ex-President of the Advertising Association in the UK, would often say "Without trust, a brand is just a product, and its advertising is just noise". In the AI era, maintaining this trust becomes even more critical as consumers grow increasingly concerned about authenticity and manipulation.

When applied to AI in advertising, these foundational principles require adaptation and extension. 'Legal' now encompasses emerging AI regulations and data protection laws. 'Decent' must address how AI might unintentionally generate inappropriate content reflecting biases in training data. 'Honest' may require transparency about AI's role in content creation itself. 'Truthful' faces new challenges when AI generates synthetic yet plausible content that may be indistinguishable from reality. The following additional principles specifically address these new dimensions while maintaining alignment with advertising's ethical foundations.

Guiding the use of AI according to ethical principles and norms therefore becomes a societal and economic imperative especially when statutory regulation may lag technological developments.

These principles exist along a maturity spectrum. Organisations should aim to progress from basic compliance (preventing clear ethical violations), through systematic management (consistent ethical processes), to optimised practice (where ethics drives innovation).

Below outlines additional ethical principles that an advertiser, advertising agency – or AI developer for that matter – should consider instilling in its organisation to guide its use of AI. Implementing these principles would help ensure that a human-centred ethical AI practice is put at the heart of an organisation to prevent or, at the very minimum, minimise the risk of harm or unethical behaviour.

The first principle concerns transparency and the need for appropriate disclosures. It is important that organisations are transparent about the use of AI especially when producing content, or even advice. Clients need to be clearly informed when AI-generated content or recommendations are being provided, and to what extent humans have been involved or provided oversight. Furthermore, the capabilities and limitations of the tools should be communicated to clients to manage expectations and minimise risk. Thought must be given to whether it would be appropriate to disclose the use of AI in an advert to the consumer especially if there is a risk it could otherwise mislead because of the way the AI is utilised.

The second principle requires respect for intellectual property and the need for proper attribution. For image generation, advertisers need to be vigilant about unintentional reproduction of copyrighted visual styles or trademark infringement. The advertiser or advertising agency should have clear policies and processes in place that mitigate the risks of plagiarism or unauthorised use of copyrighted material in AI-generated content.

The third principle requires robust fact-checking and verification processes to ensure the accuracy of AI-generated text and to prevent the spread of misinformation or disinformation. As mentioned earlier, large language models are prone to hallucinations; hence, appropriate human oversight and review of the output should be put in place to identify and correct potential errors, biases, or inconsistencies in the output.

The fourth principle focuses on diversity, inclusion, and non-discrimination. As such, the generative AI system should be trained or grounded with diverse and inclusive data to promote diversity, inclusion, and non-discrimination in the content it generates. Further steps could be taken to review the output to mitigate unforeseen biases and ensure fair representation of different groups and perspectives to avoid perpetuating harmful stereotypes, discrimination, or marginalisation of protected groups. A practical approach includes regular diversity audits of AI outputs, measuring relevant representation across demographics, and continually refining input parameters when biases are detected.

The fifth and final principle relates to ethical content and social responsibility. To that end advertisers or advertising agencies should already have their own internal guidelines and review processes to ensure that AI-generated content or advice aligns with ethical standards and positive societal values. The AI system should be grounded to avoid generating content that promotes or encourages illegal, unethical, or harmful behaviours, such as hate speech, violence, or the exploitation of vulnerable groups. It is also important to consider the potential societal impact of the AI-generated content and strive to promote positive messaging i.e. sustainable consumption, and responsible advertising practices.

These principles should be supported by robust governance frameworks, including regular audits, employee training, and clear processes for addressing ethical concerns or violations. To ensure accountability, organisations should establish measurable KPIs, such as Disparate Impact Ratios, Bias Detection Rates, Compliance Rates etc, for their ethical AI practice. Implementation of these principles may vary based on organisation size and the type of AI deployed and specific use cases. Ultimately, an organisational culture that prioritises ethical AI practices and encourages open dialogue and continuous improvement must be fostered.

1]
https://www.statista.com/statistics/1384112/consumers-who-buy-products-due-to-brand-trust-by-generation-usa/
[2] https://adassoc.org.uk/our-work/the-value-of-trust-report/

5. AI Regulation in Advertising – where are we now?

by Konrad Shek

At present, no jurisdictions have enacted legislation that specifically addresses AI-generated or AI-assisted content in advertising. This is because advertising is primarily regulated through existing consumer protection laws which prohibit misleading or deceptive advertising. These are supplemented by media regulation that mandate decency and the protection of minors and vulnerable consumers. Typically, this is complemented by sectoral specific rules for products such as alcohol, tobacco and pharmaceuticals among others. On top of this are national self-regulatory codes that help maintain and enforce advertising standards. Consequently, AI-generated content is subject to the same advertising rules as any other content, and advertisers are liable for any marketing communications they publish.

However, as highlighted earlier in the report, several jurisdictions have introduced rules to regulate AI more broadly. While the EU and China have established what is arguably the most comprehensive regulatory frameworks globally, Brazil is due to finalise its own legislation in 2025. Other countries, such as the UK and Singapore, have adopted an advisory approach, preferring to utilise existing sectoral legislation alongside non-binding guidance or frameworks.

The US presents a somewhat unique case. Whilst no Federal legislation currently exists, individual states have implemented their own laws, with California being one of the more active in enacting AI legislation. Most of California's laws have a particular focus on regulating deepfakes and safeguarding electoral integrity.

This broader AI regulatory landscape has several important implications for the advertising industry:

- **Transparency Requirements**
  Many emerging AI regulations require transparency about AI usage. In some jurisdictions, advertisers may be required to make explicit disclosures when they are utilising deepfakes. Whether this might transform into broad mandatory labelling or watermarking requirements of AI-generated advertisements is not clear at this time.

- **Liability Framework**
  Current consumer regulations maintain that advertisers remain fully responsible for their content, regardless of how it is created. Using AI tools does not absolve brands of responsibility for misleading claims or inappropriate content, even if such issues stem from algorithmic decisions rather than human ones.
- **Data Privacy Considerations**
  Global AI regulations, particularly the EU's AI Act and GDPR, impose strict requirements around data collection and processing. Advertisers using AI for personalised targeting or consumer profiling must ensure compliance with these heightened data protection standards, potentially limiting certain targeting practices.
- **Prohibited AI Applications**
  Some jurisdictions, such as the EU, explicitly prohibit certain AI applications, such as those using subliminal techniques or exploiting vulnerabilities of specific groups. Advertisers must carefully review these prohibited categories to ensure compliance.
- **Sectoral Impact**
  Highly regulated sectors (financial services, healthcare, etc.) may face additional scrutiny. AI-generated advertising in these sectors will likely need to meet both general AI regulations and sector-specific requirements.
- **Cross-Border Compliance Challenges**
  The fragmented regulatory landscape (especially between EU regulations, US state laws, and advisory frameworks elsewhere) creates significant compliance challenges for global advertising campaigns, potentially requiring region-specific content adjustments. In the absence of advertising-specific AI regulations, brands and agencies must navigate a complex regulatory landscape. This requires them to proactively interpret broader AI frameworks and apply them to their advertising activities, whilst potentially preparing for more targeted regulation that might emerge as the technology evolves. This proactive approach to compliance will help mitigate risk in what remains an evolving legal environment.

6. Content Creation and a History of Technology Tools

Those readers old enough to have been in advertising content creation in the early to mid '90s will certainly remember the arrival of Photoshop. While other technologies have been deployed into advertising for much longer, Photoshop was an astonishing tool. It almost immediately overshadowed early, more cumbersome and labour-intensive image creation and editing techniques, many of them manual or semi-manual.

Readers may also remember the considerable concern that swirled around the ungoverned use of Photoshop and other image editing tools. How was the consuming public to know whether a photo was original or manipulated? Should they be informed? If not, would trust-in-advertising be a casualty? Was there any intellectual property infringement in the creation of the image? Would there be legal hazard?

There were, in fact, a number of global attempts to regulate image manipulation (by law, regulation, self-regulation or policy). Many of the more formal attempts lagged the introduction of the technology by decades (regulation and innovation do not live in similar time zones). Most of these attempts were codified in the 2010s. They have been largely unsuccessful; enforcement is too complex and the market too fractured. Image editing technology is now a commoditised tool in advertising production, used liberally and with only rare and high-profile examples of sanction (like the banning of a L'Oreal ad 2011 and a SkinnyTan ad in 2019 by the ASA in the UK). Content producers constantly push the boundaries of image editing guidelines and policies in their search for novelty and attention; disclosure is rare.

While Photoshop is an apposite and fairly recent example of a powerful technology entering the advertising content production space, there is a much longer history, dating back to process improvements to typography and lithography at the beginning of the 20th century to the widespread use of improved rendering tools like airbrushing in the mid-20th Century and green screen in the 1970s.

But the arrival of digitisation in the 1990s was the big inflection point. Beyond Photoshop-style image manipulation, an expanded transformation of content production began with the arrival of digital video and sound which completely transformed TV commercials in terms of budget, time, and editing flexibility. Also, desktop publishing which facilitated much faster layout as well as zero-cost experimentation, and the digital audio studio which enabled an infinite range of synthetic sounds to meld and co-exist with traditional analogue sound. Content research and location scouting all benefited from the Internet, 3D animation and VFX (video effects) technologies vastly expanded the canvas of visual possibilities. Cameras have plummeted in price while expanding in functionality (partially because of the rise of smartphones with embedded camera capabilities).

All of these technologies have been absorbed enthusiastically into the advertising content creation toolset. Like Photoshop, some of these created some consternation within the industry, ranging from job losses to quality concerns to the critical matter of consumer trust.

One could be tempted to argue that AI-fuelled content tools are just another in a long line of precedents that will aid the creative and logistical process of producing advertising content, and that we should expect that it will quickly find its place in a warm corner of the existing toolbox. But I submit that 'this time is different'. AI's breadth and reach are unprecedented (including copy and music, thus far a purely human endeavour). Moreover, as discussed previously in this document, some of the AI tools are not fully understood even by its creators, making it a 'squishy' target for regulators. Its effect on the industry must be more carefully considered than a fancy editing tool or clever multilayer compositor. The recent release of Veo 3 from Google, which allows the creation of super-realistic videos replete with talking human characters, sound effects and music is one startling recent example of how quickly AI technology is moving (as evidenced by the blizzard of amateur-produced videos, which are often indistinguishable from expensive commercial productions, at a fraction of the cost).

Consumers are already sceptical of what they see and hear onscreen (where most information consumption is now ingested) – from brazen deepfakes to digital hucksterism to invasions of privacy to more subtle deceits and deformations of truth. AI can supercharge all of these toxins. Some will be inadvertently created by innocent participants; some will be deliberately exploited by bad actors.

AI is not just a new technology. For the advertising industry it is potentially much more than that, and the matter of policy consideration is paramount, as we will discuss in the upcoming chapters.

7. To Label or not to Label

7.1 Introduction

At the heart of advertising ethics lies the concept of consumer trust, legality, honesty and truthfulness, with trust sitting at its core. Effective advertising depends on maintaining a relationship of trust with consumers, who must believe that the information they receive is generally reliable and that advertisers will not mislead, exaggerate or omit important information.

These principles support informed choice – the principle that consumers should have access to relevant information to make decisions aligned with their own interests and values. This requires not just truthfulness but also transparency about material aspects of products and services that are being marketed.

This is all common cause, accepted by consumers and producers alike. But before we get into the subject of labelling, we need to address a more vexed question. To what extent do consumers trust or distrust ads created by GenAI, either in whole or part? And do advertisers have a proper understanding of consumer attitudes?

While there is considerable research into consumer attitudes to AI in general, data with respect to consumer attitudes to the more specific topic of *AI in advertising* is less robust. One of the most widely quoted recent studies was carried out by Yahoo, in collaboration with Publicis Media in April 2024, titled ***Trust Through Transparency: The Future of AI and Advertising.***

The data points to a large disparity in AI sentiment between consumers and advertisers. While 77% of advertisers view AI positively, only 38% of consumers share this sentiment. Interestingly, 53% of consumers didn't know whether AI was used in advertising and couldn't contribute to positive or negative comment. This last cohort is important; they are persons with unformed opinions, and potentially persons for whom trust can be won before it is lost.

Other studies around AI and advertising gave differing results depending on whether the question was about brand design, or product image, or content generation, or brand ambassadors, or media placement. An ***international survey*** from YouGov from June 2024 provides these metrics:

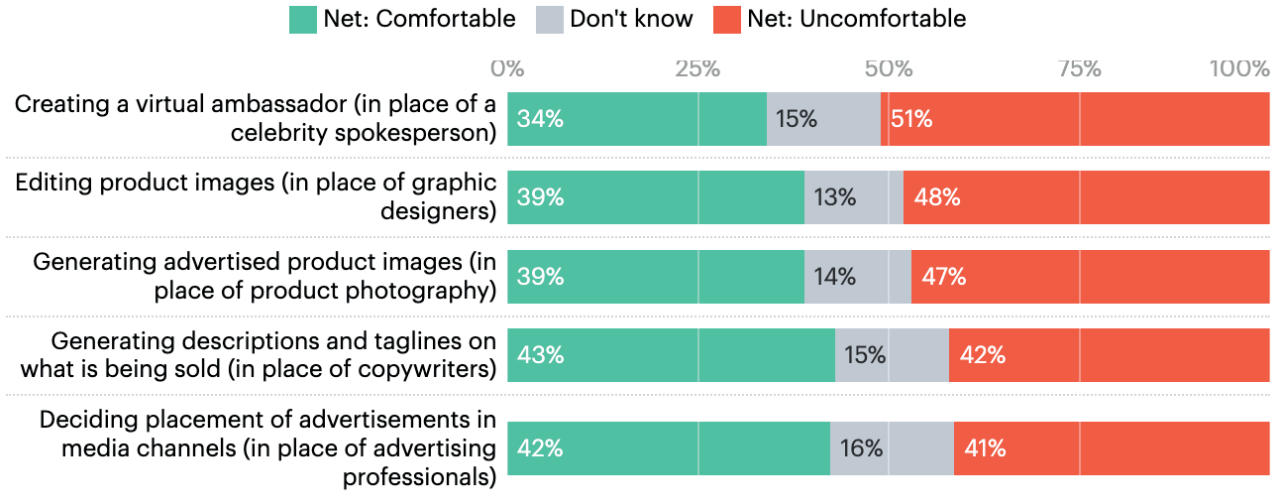

Note: Figures are rounded to the nearest whole number, and may not add up to 100%.

YouGov

YouGov Surveys: January 2024 •

One can argue these metrics, or point to somewhat different studies, but there is really one overriding message here. It is that a minority of consumers are completely comfortable with the intrusion of AI into advertising, in all its guises. It is true that the word 'comfortable' may be a little hard to define, but it certainly points to a less-than-ideal state of affairs; lack of comfort can easily migrate to lack of trust.

Of course, trust is not the only metric when considering labelling. There is consumer autonomy (allowing user choice whether to engage with labelled advertising), there is potentially valuable information about human craftsmanship and authenticity, there is education and awareness of GenAI, and there is labelling precedent for other technologies, so why not for GenAI?

But I submit that the preservation of trust between advertiser and consumer is paramount. So, we return to the Yahoo/Publicis study for perhaps the most important survey result of all. When respondents were presented with (and noticed) AI disclosure in an advertisement there was a 73% lift in ad trustworthiness, and a 96% lift in overall trust for the company (I note for completeness sake that there are other studies which contradict this figure, referenced in our companion Think Piece *Beyond Simple Labelling* by Konrad Shek).

This is a powerful pair of data points. It says that disclosure increases trust. However, as we will see, that is not the full story.

7.2 What is on the label?

Labelling and similar forms of disclosure have a long and proud history both within and outside advertising. While there were codified consumer protections dating back to Greek and Roman history, the first labelling laws in modern advertising started with the US-promulgated Pure Food and Drug Act (1906), to the expanded by the FTC Act (1914) to combat deceptive advertising, to the health warnings on cigarette packets (1965), laws and regulations and policies have blossomed in numerous jurisdictions globally. Some of these were directly tied to health safety concerns, others to ensure statutory legal disclosures, others to avoid deliberate deception, others to protect against liability, and others to bolster consumer trust.

What has historically appeared on the labels (including audio labels) has also been a matter of contestation. To what extent is disclosure required? One may take the position 'as much as possible', but this is not realistic. There is a scarcity of real estate for disclosure (on physical labels, audio segments, screen time) and scarcity means cost. Labelling is not free.

There is also the question of user resistance. A label indicating 'this pack of cigarettes is not good for your health' is not that difficult to absorb, but food labelling, as anyone who has tried to penetrate the dense tiny-font reams of data on the side of a food pack will attest, has far less efficacy. Ditto the 'terms and conditions apply' audio blather (as some experience it) at the end of an audio ad, which serves a legal purpose but generally just frustrates the consumer. Finally, there is the question of disclosure blindness (‘banner blindness’ online or simply ignoring, say, food labels).

Which brings us to AI-generated content for ads. What then should be said?

7.3 When to label – a plethora of challenges

It would be a simple matter to mandate a policy that required advertisers to simply proclaim 'this advertising was generated by AI'. But 100% GenAI created advertisements are rare. Consider the following and consider whether they should be labelled:

- The visuals are generated, but not the audio or music or copy
- A single visual element is AI-generated, while the rest of the advertisement is traditionally produced
- 49% of the advertisement is AI-generated
- GenAI is used for concept creation and creative approval, whereas the final product is recreated with other techniques
- Machine learning software is used in production, but not GenAI

And then there is this:

Many production tools are now primarily software driven, from image editing to VFX (video effects) to animation to edit studios to digital capture. The developers of the software (such as Adobe) do not provide (and could never be expected to provide) technical information as to how the code underlying each UI button, slider, dial and knob is engineered. Some will be 'good ol' software', some will be (and in the case of some packages, already are) stochastic generation driven by GenAI software. Some will be a mixture. The gnarly code details of production tools are likely to be partially, or even completely opaque to the user.

How then is the producer or advertiser expected to comply with labelling mandates if they do not know whether AI was embedded in their toolsets? Not only that, but how is the regulator able to monitor compliance? It simply is not practical.

There is also the question of cost, mentioned previously. Big advertisers and producers may be able to afford compliance regimes for even light self-regulation. But given the decentralisation of the industry into boutique advertising companies, and small production houses, there is little chance of imposing a broad solution that is universally affordable.

Finally, there is media distribution. Prior to the Internet when TV networks, major newspapers, syndicated radio stations and national magazines ruled, it was the era of 'big' campaigns – loud and visible. This made advertising standards compliance not only easier to monitor, but more of a public sensation when sanctions were applied. Now that our access to information has exploded into countless small shards, standards become harder to enforce.

Without broad and effective policing, regulation (even self-regulation) loses its teeth. Of course, mass media still lives on (the BBC, CNN, New York Times, etc.), and there will always be broad-based advertising for large-audience programming like sport. But the fraying of the edges of wide-reach advertising under the onslaught of narrowcasting has changed the calculus of advertising regulation and policy. GenAI and its potential to lower costs, shorten production cycles and expand creative possibilities can only exacerbate the problem.

At the base of these difficulties is the question of 'consumer harm'. If it is determined that an advertisement using some portion of AI is harmful, one then must be sure of the causation chain. Was it the GenAI art that caused the harm, or was the GenAI part of the advertisement so integrated into the whole that such a determination is impossible? I submit that there is no easy answer to this.

7.4 Do these challenges make AI labelling DOA?

The short answer is no, or at least probably not, or perhaps even not yet.

I submit that there are reasonably contained circumstances that would justify disclosure. For instance, public figure lookalikes or even realistic AI-generated human avatars. Or any attempts to pass off any AI-generated element as 'real', including physical locations, landmarks, animals or products, although even here we run into difficulties. The first is the definition of 'real', and the second is the obvious fact that high-quality 3D renderings face no such restrictions and often can render realistic objects that escape detection. And then of course there is the more fundamental subjective question of whether a fake generated product shot or even a AI-generated human that is perceived as real, causes any harm at all

And what of music, audio and copy?

It seems absurd for there to be forced disclosure of AI-generated music given that elements of widely used digital sound techniques (like sampling or purchased beats or pitch correction) are used without comment. Perhaps it is GenAI composition that should be disclosed, but again, it depends on contribution to the whole.

Voice mimicry seems like a reasonable target of disclosure, particularly if it is publicly recognisable. But anonymous unrecognised voice labelling makes no sense, there have been computer generated voices with us for years, such as on customer service bots.

AI-generated ad copy seems like a reasonable labelling target, but it would seem pedantic to do so and might not have the effect of increasing consumer trust.

All of this leaves us in a quandary. If there are so many exceptions, ifs and buts, perhaps the most rationale approach would be simply to say that 'Some AI-generated elements may have been used in the creation of this advertisement'.

A little hand-wavy, to be sure, but surely better than nothing?

(This issue will be further explored in an upcoming second deep dive Think Piece on this subject from the Global ICAS Think Tank.)

8. Unintended Consequences

The regulatory road is paved with good intentions (to mangle a cliche). Whether arising out of legislation, policy guidance or self-regulation, these dictates, in all cases, are directives meant to balance risks and rewards in the pursuance of a public or private initiatives, activity, business, or system.

In any definition of regulation, one finds references to an 'authority'. Those are the authors of regulation who are assumed to have not only expertise in the matter being regulated, but the wisdom to write regulations which correctly balance any competing interests that may conceivably arise.

This is tautological, but what happens when the matter being regulated itself is ill-understood and unpredictable, even to experts in the field? The regulatory space then becomes a place of risk, because once a regulation is codified, it becomes obstinate, difficult to edit or delete.

And this is where the advertising industry finds itself in contemplating regulations for AI. It is technology that is written on ground that shifts quickly, sometimes within weeks or days. What then are the unintended consequences that might arise from well-intentioned regulations when AI shifts its shape, as has happened regularly over recent decades.

To offer an example: what is currently referred to as 'GenAI' (and which is at the heart of these discussions) arose from a statistical formalism presented in a 2017 Google paper called 'Attention is all you Need'. The entire edifice of novel content being generated by an AI, rests on this formalism, known as 'Transformers', which require a training set on which to feed. It is not only possible, but likely, that this formalism will be entirely replaced by a non-transformer-based architecture in the future. Will that still be under the general rubric of 'GenAI', and if not, does this obviate GenAI regulations?

A simple response to this might be that the underlying architecture is not the issue; when novel content is 'generated', then it must be regulated, at least to some degree. Which leads us back to the quandary we faced in Chapter 2, where we discussed good ol' software, which also has the ability to generate novel content and which is not AI in any sense. Digital tools have contributed to the generation of novel content for decades. The current discussion about GenAI is a matter of scale, quantity, quality and breadth of choice. But it is, in principle, the same as asking a 2015 version of a paint program to generate a background.

However, let's push through this wrinkle and table some of the unintended consequences of AI that need to be mitigated:

- **Overreliance on Human Content** Mandatory labelling might create artificial preferences for human-created content even when AI-generated alternatives are equivalent or superior in quality, accuracy, or relevance.
- **Tactical Avoidance** Strict labelling requirements might incentivise technical workarounds, such as minimal human modification of AI outputs specifically designed to avoid labelling thresholds.
- **Small Business Disadvantages** Compliance costs associated with tracking and disclosing AI use throughout complex production processes might disproportionately burden smaller advertisers without sophisticated workflow systems.
- **Creative Limitation** Creators might avoid beneficial AI tools in certain contexts simply to avoid labelling requirements, potentially limiting creative possibilities and innovation.
- **False Security** Labels might create a false sense that non-labelled content is entirely ‘natural’ or unmanipulated, despite the long history of non-AI manipulation techniques in advertising. This is referred to as implicit truth bias.

- **Geographic Regulatory Arbitrage** Inconsistent global regulations might push AI advertising development and deployment toward jurisdictions with less restrictive rules, potentially undermining the effectiveness of stricter regulations.
- **Reduced Competition** Compliance requirements that demand significant technical and legal resources could entrench dominant market players while creating barriers to entry for new competitors.
- **Formal Compliance vs Substantive Impact** Complex regulatory frameworks might shift focus toward technical compliance rather than meaningful consumer protection, creating 'checkbox' approaches that satisfy legal requirements without addressing underlying concerns.
- **Information Overload** Proliferation of AI-related disclosures across numerous advertisements might contribute to disclosure fatigue, reducing consumer attention to truly important information.

This is a thicket of challenges. How best to respond? How to mitigate regulatory harms?

This answer seems straightforward (which, in practice, it never is). Move slowly, start with a small number of low hanging fruit (like 100% GenAI generated ads), have mechanisms to deregulate and re-regulate often, consult with stakeholders continuously (creating dynamic and frictionless feedback loops), and be prepared to pivot on a dime. And as discussed previously, the thorny issue of whether any harm at all would be inflicted if an 100% AI generated ad was viewed by a consumer without harm arises. Perhaps the harm is simply the lack of disclosure, even if nothing else is deceitful?

This is not a common approach to regulation which normally seeks a final resting place. I argue that it is necessary here.

9. Audience Targeting

Anyone who spends time on the Internet has occasionally had the creepy feeling that ads are being served up minutes after having had an IRL conversation about a related subject. There have been rumours for years about this activity, called 'active listening' and at least one warning from the FTC with regards to a company called Alphonso Software which sold data collection IP to mobile game developers.

It would, of course, be unethical (and perhaps illegal), and most definitely against terms of services of platform providers like the App Store.

But this highlights a conundrum. The attempt to target audiences is as old as advertising itself. It has historically been uncontroversial, as advertisers sought to make the most efficient use of media spend in search of a holy grail of sorts – the perfectly relevant advert. It has also been welcoming for many consumers, allowing them to avoid the distraction of an irrelevant advertisement. But to understand the unique implications of AI targeting, we must first consider its historical context:

- **Mass Media Era (Pre-1990s)** Targeting relied primarily on broad demographic categories and media consumption patterns, with limited personalisation capability.
- **Early Digital Era (1990s-2000s)** Cookie-based tracking enabled more granular targeting based on browsing behaviour and simple profiling techniques.
- **Data Integration Era (2000s-2010s)** Cross-device tracking and data broker ecosystems enabled more comprehensive consumer profiles combining online and offline behaviours.
- **AI Targeting Era (2010s-Present)** Machine learning algorithms now analyse vast datasets to identify patterns, predict behaviours, and optimise targeting across multiple dimensions simultaneously.

This evolution has been characterised by increasing precision, from broad demographic segments to hyper-specific individual targeting based on comprehensive behavioural profiles.

With the arrival of AI, it is no surprise that the issue of targeting and privacy has become more acute. Which brings up somewhat of a paradox. We don't like to be sold things we don't need or want. We also don't want strangers to know too much about us; it feels invasive, and people value their personal interpretations of privacy. This is the core tension of targeting.

In which ways has AI expanded targeting opportunities? AI systems (whether GenAI or non-generative machine learning) vastly outperform traditional targeting software in two areas: pattern recognition and continuous improvement. This facilitates both higher precision in matching content to consumers, and ensures that the targeting is adaptive – continuously aligning with consumer profiles even as they may change.

And so we have:

- Machine learning algorithms that predict consumer receptivity to particular messages based on past behaviours, demographic information, and contextual factors
- AI systems that can identify patterns among existing customers and find others who share similar characteristics, expanding potential audience reach
- AI that can automatically adjust creative elements based on individual viewer characteristics, preferences, and contexts
- AI that helps connect identities across devices and platforms, creating more comprehensive targeting profiles
- Natural language processing that evaluates consumer sentiment in social media and other contexts to inform targeting strategies (sometimes in real-time) more accurately than traditional keyword-based sentiment analysis
- Recommendation engines

Which leads us back to the privacy conundrum, which immediately raises a definitional problem. What is privacy? It is certainly not monolithic (a personal reveal with a best friend has different consequences from that same reveal with a stranger). It changes over time (what I choose to reveal today may not be the same as tomorrow). It is culturally based (expectations of privacy are different across religious, national, and cultural groups). And it changes with context (medical data vs retail).

Individual consumers need the ability to declare their privacy preferences and change them at their discretion (sometimes referred to as 'privacy permission management'). And targeting technologies need to disclose their methodologies to consumers to allow opt-in or opt-out. Between these two disclosures, a shared space exists in which both consumers and targeting technologies can be made to be 'target-aware' and to act accordingly.

However, even if one accepts this as the framework on which to build regulations, problems remain. The incentives are not aligned - it is in the consumers' interest that they manage their privacy preferences; that the benefits are understood and welcomed. But it is not in the advertisers' commercial interest that they disclose their targeting secrets; doing so may reveal internal IP as well as create an open door for consumers to 'opt out'. Appealing to ethics in business will always be an uphill battle against these two imperatives.

Still, there are targeting-related regulatory frameworks that have already been tabled within advertising and even legislated for in general industry in the pre-GenAI era. What is now required is a considerable shoring up of those frameworks and principles to protect consumers from potentially far more surreptitious and surgical harvesting of their personal data (without permission or payment). And secondly, to offer consumers the opportunity to declare their comfort level in the matter of personal data profiling, and to have the technology be filtered through it.

We suspect that, given the aforementioned misalignment of incentives, the relationship between targeting and privacy is not only going to remain tense but will also shapeshift as targeting technologies and personal interpretations of privacy change.

Notwithstanding that, the following approaches (many of which are also well-regulated in some jurisdictions) should be re-examined and possibly strengthened in light of the new AI-fuelled targeting technologies already deployed or under development:

- **Meaningful Transparency** Providing clear, accessible information about how targeting works in practice, not just in technical legal disclosures.
- **Contextual Targeting Alternative** Exploring contextual relevance (based on content being viewed) as an alternative to extensive behavioural tracking when appropriate.
- **Sensitive Category Limitations** Establishing clear boundaries around targeting based on sensitive characteristics (such as health conditions, financial status, emotional states) whether explicitly collected or algorithmically inferred.
- **User Control Mechanisms** Developing intuitive, accessible tools for consumers to understand and control how they are targeted.
- **Accountability Structures** Implementing regular auditing of targeting algorithms for potential discriminatory impacts or other problematic patterns.
- **Proportionality Principle** Ensuring the level of data collection and profiling is proportional to the purpose and consumer benefit of the targeting.

Finally, we need to mention Web3. This is a vision of a future Internet owned entirely by its users, including the data generated by that usage, such as browser journeys. This data will be blockchain-lodged; it technically cannot be used without permission of the user. This would obviously prevent un-permissioned targeting. There are already technology layers being built in support of this vision, but its prospects of wide adoption remain a long way off.

10. Dark Patterns and AI – Fertile Ground for Malfeasance

The matter of deliberate mischief in both advertising and non-advertising content has been around for a while, particularly since the arrival of the interactive Internet and social media.

Surveys and research from organisations like Statista, Byyd, Bannerflow, Datareportal and others indicate that about 70% of global ad spend is now online, some of it originating from non-traditional advertising sources, and some of those ungoverned and untrusted (and about half of it on social media). More pointedly, according to some research metrics, social media ads are now the leading source of brand awareness for the 16-24-year-old age group. Other research indicates offline ads reaching only 3%-4% of the global population, compared to 93% for social media ads (I am not vouching for research methodology here, but indicators are clear).

It is also well reported that the traditional agency business has seen the rise of uncountable micro-agencies who, while likely not representing the larger reputable brands, are a significant force in the larger advertising ecosystem. It is in this underregulated space where AI looms as an avenue for bad actors.

In addition to AI, the power of interactive technologies has dramatically amplified the possible avenues of deceit, which has entrapped many, including this author, who recently paid for a pair of running shoes on a fake Adidas site, stemming from a well-designed online ad.

AI is and will be supercharging the deceit business in ways that we have yet to fully understand (other than a general discomfort that the ground for miscreancy is fertile). 'Dark patterns' – user interface design choices that manipulate or heavily influence users to make certain decisions – represent a particularly concerning area where AI capabilities might be misused in advertising contexts.

In addition to AI's ability to generate text, audio, images and video, there are pattern-recognition and statistical capabilities which can be combined together to malfeasant effect. And so we have a long list of possible incognito deceit that becomes available with misuse of AI:

- **Hyper-Personalised Emotional Manipulation** AI systems capable of detecting emotional states could deliver advertisements precisely when users are most vulnerable to particular persuasive techniques -- capitalising on moments of sadness, boredom, or low self-esteem.
- **Synthetic Social Proof** AI-generated fake reviews, testimonials, or social engagement signals could create artificial impressions of product popularity or effectiveness, exploiting the powerful influence of social proof on decision-making.
- **Misleading Conversational Agents** AI chatbots or voice assistants designed to simulate helpful advisors could subtly steer conversations toward particular purchasing decisions while maintaining the illusion of neutrality.
- **Fake Scarcity and Urgency** AI systems could generate personalised, false signals of limited availability or time constraints based on analysis of individual decision-making patterns and susceptibility to such pressures.
- **Reality Distortion** Advanced image and video generation could create subtly altered depictions of products that systematically exaggerate benefits while appearing authentic and unmanipulated.
- **Cognitive Overload Engineering** AI could help design interfaces that strategically overwhelm cognitive resources at key decision points, increasing reliance on heuristics and emotional responses rather than rational evaluation.
- **Addiction Mechanisms** Machine learning could optimise engagement features to exploit addictive behavioural patterns, maximising time spent with advertising content through variable reward schedules and other psychological techniques.

We are aware that this list has a large overlap with previous chapters on targeting and unintended consequences. There is, however, a critical difference. The actions described in the list above are a *deliberate* strategy to circumvent regulations, and (in some cases) the law. That is different to a targeting strategy that trades on the grey area of privacy, or unintended consequences of technological overreach.

In contemplating how to defend advertising content from this less-than-exhaustive list of attacks, several issues need to be considered.

The first is that of attribution. Is it a bad human actor who wrote the code that enabled the dark pattern, or is it a consequence of an AI agent? It is not currently possible to distinguish between the two.

Secondly, is there even a bright line between the light art of persuasion and the dark art of manipulation? Probably not.

Thirdly, the definition of manipulation is not the same in all jurisdictions, or even cultures.

Fourthly (and most worrisome), there is little chance that regulators will be able to stay at the cutting edge of AI code generation and agentic AI. Indeed, even AI professionals are worrying that the technology will soon break free of their ken and control.

So, what is there to do? There are measures that can be taken and are already taken by many companies who have been taking ethics seriously for decades, none of them watertight, but at least attempt to hobble and constrain the basic instincts that can be found at the edges of the industry:

- **Algorithmic Auditing** Regular independent audits of AI advertising systems could help identify patterns of manipulation or unfair influence.
- **Design Ethics Guidelines** Industry-specific ethical guidelines for AI design in advertising could establish boundaries for acceptable persuasive techniques.
- **Consumer Education** Empowering consumers to recognise manipulation attempts could provide some protection while regulatory frameworks develop.
- **Technical Countermeasures** Browser extensions and other tools might help users identify potential dark patterns and manipulation attempts.
- **Regulatory Approaches** Expanded definitions of unfair and deceptive practices could be applied to AI-enabled dark patterns, with meaningful enforcement mechanisms.
- **Whistleblower Protections** Encouraging those within the industry to report unethical practices by providing meaningful protection from retaliation.

The potential for AI to enable more sophisticated dark patterns represents one of the most significant ethical challenges in AI advertising. Addressing this challenge requires coordinated effort from industry leaders, technology developers, regulators, and consumer advocates to establish and enforce appropriate boundaries on persuasive techniques.

The stakes are particularly high because widespread exploitation of dark patterns could undermine trust in the broader digital ecosystem and trigger severe regulatory backlash that might constrain beneficial applications of AI in advertising. Proactive attention to these risks represents both an ethical imperative and a practical necessity for the industry's sustainable development.

11. Final Thoughts and Recommendations

Here is a thought experiment: Is there any point in trying to regulate AI in advertising content?

The preceding chapters have covered many instances in which the new technology is being (and will be) brought to bear on the generation of advertising content. AI technology, in some form or another, is likely to be deeply embedded in just about every digital tool used for the ideation, conception, creation, production and distribution of advertising content. Creators and producers, in many cases, will not know or care what technology framework and code is driving the toolset – they will just use those that best facilitate their end vision.

Some consumers may indeed care; trust in the advertising ecosystem is fragile and precious. But I submit that it is not AI that they care about. They care that they are not being deceived or manipulated. Whether the technology doing that is AI or good ol' software, the concerns are exactly the same.

This leads us to questions of whether AI and the ethical arguments that rage around it are at a useful level of abstraction. We reach outside of the advertising industry to illustrate this point, recently used by tech commentator Benedict Evans.

Here is the (slightly edited) section:

"I had a conversation a couple of days ago with a political journalist. He said, well, shouldn't we regulate AIs to make sure it follows our values?

And I gave an example of configuring telco billing systems. Do you think AI regulators need to pass a law to make sure that the company that's helping a big telco reconfigure its billing systems to incorporate a different pricing plan needs to reflect our democratic values? What are you talking about?

What you actually mean is people building AI systems that will decide who gets parole."

(From "Another Podcast: Looking for AI strategies," 05 Apr 2025)

The point is clear – regulating from top down is a blunt catch-all that will produce less-than-useful, wasteful and perhaps absurd regulations. Regulation is more effective when it is directed at very specific scenarios and situations (like whether a 'testimonial' comes from a person or an AI avatar).

Of course, this might simply be moving the problem from the frying pan into the fire – there are myriad 'specific' scenarios that would benefit from regulation, perhaps too many to consider.

However, given that there are already robust self-regulations that have been accepted by the advertising industry, we recommend the following:

1. An analysis of which GenAI capabilities are clearly differentiated from pre-GenAI technologies. This would include AI 'human-like' spoofing (across video, audio and copy), real-time dynamic pattern-based microtargeting based on autonomously-improving machine learning, hyper-real AI video generation (including products and location), non-permissioned consumer data collection for AI training, etc.
2. Comparison of new capabilities to pre-AI codes to see whether any can be easily expanded to accommodate GenAI.
3. 'Common-sense' codes-of-conduct (including but not limited to labelling) for each new discrete AI capability (with industry collaboration and input).
4. 'Common-sense' codes-of-conduct for new instances of AI generation that may arise from time to time (such as agentic 'brief-to-completed ad' applications, which are likely to be built soon).
5. Continuous 'fine-tuning' as AI innovation continues apace (with updates to interim regulations at least several times per year).
6. Establish 'general' global AI regulation landscape monitoring, with frequent report-backs, possibly monthly.
7. Establish AI technology landscape monitoring, with frequent report-backs, possibly monthly.
8. Periodic consumer temperature surveys to establish thresholds of AI trust.
9. Education update sessions with industry.
10. Do not consider technology monitoring – it is impractical.
11. Produce guidelines for advertisers to engage with external producers/creators to exchange ideas around advertising content ethics for AI.
12. Inculcation of flexibility to quickly and frictionlessly drop and add codes to try and match GenAI evolution.

It has not escaped notice that this list of recommendations, implemented, would require considerable resources to analyse, synthesise, manage and maintain on an ongoing basis, and may not be practical or affordable in all cases. Although it is also true that self-regulation does in fact mitigate these problems.

However, GenAI in advertising content will soon move from novelty to commodity across the entire content production chain as the technology seeps into general usage and goal-driven AI continues its dizzying pace. Given the criticality of maintaining consumer trust during this transition period, I believe that the window is small and the matter is urgent.

The following AI research tools were used in the preparation of this report: Anthropic 3.7, ChatGPT Deep Research, DeepSeek DeepThink, Perplexity Deep Research, GrokAI, Gemini Deep Research

Approximately 90% of the words on the page were written by the author. Some lists and bullet points were copied from research.

The author thanks members of the ICAS Global Think Tank AI Working Group for advice, comments and suggestions.

For more information about the ICAS Global Think Tank, please contact Sibylle Stanciu, ICAS Director (sibylle.stanciu@icas.global)

The International Council for Advertising Self-Regulation (ICAS)
https://icas.global

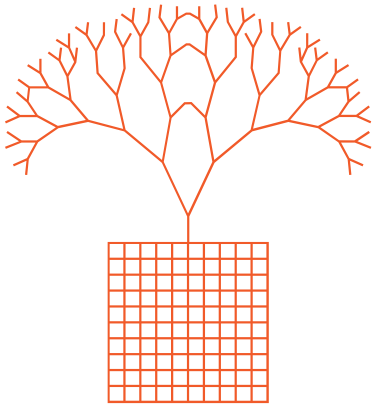

ICAS
GLOBAL
THINK
TANK