

Gender-Based Targeting in Meta Ads and Google Ads: Strategic Advantage or Algorithmic Bias in Customised Advertising Campaigns

Dr. T. SathyaPriya¹, K. Agilmathi², V. Rahul³

¹Assistant Professor, Department of Management Studies, Coimbatore Institute of Technology, ²Student, Department of Management Studies, Coimbatore Institute of Technology

Emails: sathyapriya.t@cit.edu.in¹, agilmathi3112@gmail.com², 2403717663121037@cit.edu.in³

ABSTRACT

The proliferation of data-driven advertising platforms has fundamentally reshaped how brands communicate with consumers. Among the many targeting parameters available to marketers today, gender remains one of the most extensively used yet most contested variables. This paper critically examines gender-based targeting within Meta Ads (formerly Facebook Ads) and Google Ads, investigating whether such practices constitute a legitimate strategic instrument for campaign customisation or represent a deeper structural problem rooted in algorithmic bias. Drawing on an interdisciplinary body of literature that spans digital marketing, behavioural economics, machine learning ethics, and consumer psychology, this study constructs a conceptual framework that maps the tension between advertising effectiveness and ethical accountability. The paper analyses documented cases of gender-discriminatory ad delivery, evaluates the regulatory landscape in the European Union and the United States, and explores how platform-level algorithms can perpetuate or even amplify gender-based inequalities beyond the intentions of individual advertisers. Findings suggest that gender targeting, when deployed strategically and transparently, can yield measurable improvements in click-through rates, conversion efficiency, and return on ad spend; however, the same mechanisms, left unchecked, routinely produce biased outcomes in contexts such as employment, credit, and housing advertising. The paper concludes by proposing a dual-lens framework that empowers practitioners to leverage gender-aware strategies while adhering to ethical design principles. Recommendations are offered for advertisers, platform architects, and policymakers to ensure that personalisation does not become a vehicle for systematic discrimination.

Keywords: *gender targeting, Meta Ads, Google Ads, algorithmic bias, digital advertising, programmatic targeting, ad personalisation, ethical marketing, platform accountability*

1. INTRODUCTION

The modern advertising landscape is almost unrecognisable compared to its pre-digital predecessor. Where a television commercial once broadcast a single message to a demographically undifferentiated mass audience, today's programmatic systems can identify, segment, and serve a uniquely tailored advertisement to a specific individual within milliseconds of a webpage loading. The twin pillars of this transformation are Meta Ads and

Google Ads — the two most dominant digital advertising ecosystems in the world, collectively commanding roughly 48 per cent of global digital ad spend as of 2023 (Statista, 2024). Between them, these platforms have pioneered targeting technologies that allow brands to specify with extraordinary granularity who should see an advertisement: by location, age, purchase history, browsing behaviour, declared interests, and — crucially for this paper — gender.

Gender-based targeting has become a standard tool in the digital marketer's toolkit. Its intuitive appeal is not difficult to understand. Decades of consumer research have established that gender influences purchasing behaviour, brand preference, information processing style, and sensitivity to advertising tone (Meyers-Levy & Sternthal, 1991; Darley & Smith, 1995). If those differences are real — and much evidence suggests they are, albeit in complex, non-binary, and culturally situated ways — then tailoring advertising to account for them is merely an extension of good marketing logic. A luxury skincare brand running separate campaigns for different gender audiences, an automobile manufacturer using different creative assets for men and women, or a streaming service recommending genre categories by audience profile would, on this view, simply be practising intelligent segmentation.

Yet this tidy logic collapses the moment targeting moves beyond aesthetic preference into opportunity allocation. A landmark study by Lambrecht and Tucker (2019) demonstrated that even gender-neutral job advertisements on Facebook were algorithmically delivered in a skewed manner: advertisements for science, technology, engineering, and mathematics (STEM) careers reached a substantially higher proportion of men than women, not because any advertiser specified a gender filter, but because the platform's own optimisation algorithms — trained on historical engagement data — produced discriminatory outcomes by default. This finding, replicated and extended in subsequent research (Ali et al., 2019; Speicher et al., 2018), fundamentally shifts the debate. The question is no longer merely whether advertisers choose to target by gender, but whether the platforms themselves introduce bias regardless of intent.

This paper engages both dimensions of the problem. Section 2 reviews the theoretical and empirical literature on gender targeting in digital advertising. Section 3 examines the mechanics of Meta and Google advertising platforms, identifying the specific features and algorithmic processes through which gender-based differentiation operates. Section 4 assesses the strategic case for gender targeting, presenting evidence of effectiveness across campaign types and industries. Section 5 investigates the documented harms of algorithmic gender bias, drawing on academic studies, regulatory investigations, and legal proceedings. Section 6 surveys the regulatory landscape, with particular attention to developments in the European Union under the Digital Services Act and in the United States. Section 7 develops a conceptual framework — the Strategic-Ethical Targeting Continuum — that situates gender targeting practices along a spectrum from clearly beneficial to manifestly harmful. Section 8 offers practical recommendations for marketers, platform designers, and regulators. The paper concludes with reflections on the future of gender-aware advertising in an era of intensifying scrutiny.

2. LITERATURE REVIEW

2.1 GENDER AND CONSUMER BEHAVIOUR: FOUNDATIONAL PERSPECTIVES

Consumer behaviour scholarship has long recognised gender as a meaningful organising variable, though the nature and extent of its influence have been debated extensively. Early work by Meyers-Levy and Sternthal (1991) introduced the selectivity hypothesis, which posits that women engage in more elaborate, detail-oriented information processing than men, who tend to adopt a more heuristic, schema-based approach. This framework, however influential, has been criticised for overstating gender differences and for implicitly reinforcing essentialist assumptions (Fischer & Arnold, 1994). Subsequent research has foregrounded intersectionality, demonstrating that gender's effect on consumption is mediated by age, ethnicity, socioeconomic status, and cultural context (Crenshaw, 1991; Bettany et al., 2010).

More recent work has embraced a social constructionist view, treating gender not as a stable binary attribute but as a performative identity that individuals enact through consumption practices (Butler, 1990; Üstüner & Holt, 2010). This perspective carries significant implications for advertising: if gender is performed rather than fixed, then targeting systems premised on binary classification risk both distorting the identity of the audiences they claim to serve and narrowing the cultural space within which those audiences can express themselves.

Despite these theoretical complexities, the practical marketing literature continues to treat gender as a useful heuristic for campaign planning. Meta's own Business Help Centre (2023) explicitly lists gender among its core audience targeting parameters, alongside age and location, treating it as an unproblematic dimension of consumer identity. Google similarly enables gender-based bid adjustments and audience exclusions across its Display and Search networks. The gap between the sophistication of academic gender theory and the comparative crudeness of platform implementations is itself a source of the tensions this paper seeks to address.

2.2 PROGRAMMATIC ADVERTISING AND ALGORITHMIC DECISION-MAKING

Programmatic advertising refers to the automated purchase and placement of digital ads using algorithmic systems that operate in real-time bidding (RTB) environments (Busch, 2016). Within milliseconds of a user loading a webpage or app, an auction takes place in which advertisers bid for the right to serve an impression. The winning bid is determined by a combination of the advertiser's stated maximum price, the probability of a user clicking or converting (estimated by the platform's predictive model), and various quality signals. Gender, age, and other demographic variables feed into both the advertiser-side targeting specifications and the platform-side delivery optimisation.

The algorithmic nature of programmatic advertising introduces a layer of opacity that distinguishes it from traditional media buying. When a television network sells advertising time on a programme with a predominantly female audience, the gender orientation of the placement is explicit and visible. When a Meta algorithm decides — on the basis of countless inferred signals — that a particular user profile is unlikely to engage with a STEM job advertisement, that decision is invisible, unannounced, and potentially unintentional. Diakopoulos (2016) uses the term 'algorithmic accountability' to describe the challenge of making such decisions legible and contestable, and this challenge is central to the bias debate.

Barocas and Selbst (2016) identify several pathways through which machine learning systems can generate discriminatory outcomes even in the absence of explicit discriminatory intent. Training data may reflect historical inequalities; proxies for neutral variables may correlate with protected characteristics; and optimisation objectives may align reward functions with behavioural patterns that diverge across demographic groups. All three mechanisms are plausibly active in gender-targeting contexts, as the empirical literature examined in Section 5 suggests.

2.3 STRATEGIC EFFECTIVENESS OF GENDER TARGETING

A substantial body of research supports the strategic value of gender-based message customisation. Putrevu (2001) found that women responded more favourably to comparative advertising that featured detailed, attribute-based arguments, while men showed stronger positive responses to non-comparative, image-oriented appeals — consistent with the selectivity hypothesis. Darley and Smith (1995) demonstrated that advertising congruence with gender-role expectations improved brand attitude in product categories traditionally associated with specific genders.

In digital contexts, Jansen et al. (2013) analysed click-through patterns across a large sample of sponsored search results and found systematic gender differences in keyword engagement, suggesting that tailored bidding strategies could improve campaign efficiency. Tang et al. (2015) used large-scale A/B testing data from a social media platform to show that gender-segmented creative assets consistently outperformed gender-neutral versions on engagement metrics including likes, shares, and conversion rates.

However, the case for gender targeting is more nuanced in markets characterised by product-gender incongruence. When women purchase automotive products or men purchase skincare, campaigns calibrated to majority-gender audiences may actively exclude profitable minority-gender customers (Moss et al., 2006). This suggests that gender targeting is not a universally superior strategy but a conditional one whose value depends on the match between product category, brand identity, and the gender distribution of the actual purchase population.

2.4 ALGORITHMIC BIAS AND GENDER DISCRIMINATION IN ADVERTISING

The most consequential strand of recent literature concerns not the strategic deployment of gender targeting by advertisers but the inadvertent introduction of gender bias by platform algorithms. Ali et al. (2019) conducted a seminal experiment in which they placed gender-neutral advertisements for jobs and credit products on Facebook and documented that delivery was systematically skewed by gender: logging ads were delivered predominantly to men, while supermarket cashier ads reached mostly women. These outcomes occurred without any gender targeting instruction from the advertisers and were attributable to the platform's delivery optimisation, which maximised engagement by matching ad content to audiences whose historical behaviour was statistically associated with similar content.

Speicher et al. (2018) extended this analysis to show that even when advertisers explicitly prohibited gender discrimination in their targeting specifications, the ad auction mechanism itself could produce discriminatory delivery outcomes, because ad delivery is shaped by competition for audience attention among many advertisers simultaneously, not merely by any single advertiser's preferences. This systemic character of algorithmic bias is important: it means that the problem cannot be fully addressed at the level of individual advertiser behaviour and requires platform-level architectural changes.

Google has faced analogous criticism. A 2015 investigation by the American Civil Liberties Union found that Google's ad serving system was significantly more likely to show advertisements for high-paying executive recruitment to men than to women across equivalent searches (ACLU, 2015). Subsequent research by Lambrecht and Tucker (2019) confirmed this pattern and offered a market-competition explanation: female users are, on average, more valuable audiences for a wide range of advertisers and thus face higher competition in the ad auction, making them more expensive to reach, which deters some advertisers — particularly those with lower budgets — from reaching them.

3. PLATFORM MECHANICS: HOW GENDER TARGETING WORKS

3.1 META ADS: TARGETING ARCHITECTURE

Meta's advertising platform offers gender as a first-order demographic targeting parameter within its Audience Definition tool. Advertisers may select 'Men', 'Women', or 'All', with these categories mapped to the gender information that users have declared in their profile settings. As of 2023, Meta offers users the option to specify a custom gender in their profile, but the advertising platform operationally collapses this to a male/female binary for targeting purposes, a simplification that has attracted criticism from LGBTQ+ advocacy organisations (GLAAD, 2022).

Beyond declared gender, Meta's Detailed Targeting system enables advertisers to target audiences based on interests, behaviours, and life events that are themselves correlated with gender. An advertiser wishing to reach predominantly female audiences without explicitly specifying a gender filter could, for example, target users interested in particular magazine brands, parenting topics, or lifestyle categories whose audiences skew heavily female. This practice — sometimes called indirect or proxy gender targeting — raises its own set of concerns, as it can achieve demographic discrimination through facially neutral means.

Meta's Advantage+ audience product, launched in 2022, introduces a further layer of algorithmic autonomy, allowing the platform to expand or override advertiser-defined audience segments when its machine learning systems predict that doing so will improve campaign performance. This means an advertiser who has specified a female audience may find their ads served to male users if Meta's algorithm predicts sufficiently high engagement. The inverse is also possible. The advertiser has, in effect, delegated part of the gender-targeting decision to the platform.

3.2 GOOGLE ADS: GENDER TARGETING CAPABILITIES

Google Ads offers gender targeting across its Display Network and YouTube advertising products, and permits gender-based bid adjustments in Smart Shopping and Performance Max campaigns. In Search campaigns, gender targeting is available as a bid modifier rather than an absolute exclusion, meaning advertisers can increase or decrease the amount they bid for impressions served to users of a given gender without entirely removing those users from reach.

Google infers the gender of users who have not explicitly declared it through a combination of search query analysis, browsing behaviour, app usage patterns, and demographic modelling. The inference engine is probabilistic — it assigns a gender likelihood score rather than a definitive classification — and users who fall below a confidence threshold are classified as 'Unknown'. This probabilistic architecture means that gender-based targeting on Google

is inherently approximate, with a meaningful proportion of impressions delivered to users whose inferred gender may not match their actual gender identity.

Google's Performance Max campaigns represent an evolution analogous to Meta's Advantage+ product: the advertiser specifies a creative asset set and a conversion goal, and the algorithm determines audience selection, placement, and bid levels autonomously. Gender, in this context, becomes one of many signals the algorithm may or may not activate in pursuit of the conversion objective, and the advertiser has limited visibility into whether or how it is being used.

3.3 THE ROLE OF DELIVERY OPTIMISATION

Perhaps the most critical and least understood aspect of gender targeting in both platforms is the distinction between targeting (whom the advertiser specifies) and delivery (whom the algorithm actually serves). As multiple studies have shown, these two sets can diverge significantly. Platform delivery algorithms are optimised primarily for engagement and conversion outcomes — clicks, installs, purchases, form completions — and they achieve this by identifying within the target audience those users most statistically likely to perform the desired action.

If historical data indicates that a certain category of content is more likely to generate clicks from men than from women, the delivery algorithm will skew towards men even within an audience explicitly defined to include both. The result is that stated targeting parameters represent a ceiling rather than a distribution — the advertiser says who may receive the ad, but the algorithm decides, within that pool, who actually does. This distinction is crucial for understanding how bias can emerge from systems designed with neutral intent.

4. THE STRATEGIC CASE FOR GENDER-BASED TARGETING

4.1 EFFICIENCY GAINS IN CONSUMER PRODUCT ADVERTISING

In categories where purchase behaviour is demonstrably gendered — personal care, fashion, certain food segments, and toy retail, among others — gender targeting can meaningfully improve campaign efficiency by concentrating ad spend on the audiences most likely to convert. A study of Facebook Advertising return on ad spend (ROAS) across 200 consumer brands found that gender-targeted campaigns in personal care achieved an average ROAS 31 per cent higher than gender-neutral campaigns in the same category, attributable primarily to reduced spend on low-probability-conversion impressions (Moorman, 2020).

This efficiency argument is especially compelling for small and medium-sized enterprises (SMEs) with constrained advertising budgets. For a brand that cannot afford to advertise broadly, gender targeting offers a relatively low-cost method of focusing resources. It also enables testing and iteration: a brand can run parallel male-targeted and female-targeted campaigns, measure differential response, and allocate subsequent budget to whichever segment proves more commercially productive.

4.2 CREATIVE RELEVANCE AND MESSAGE RESONANCE

Beyond efficiency, gender targeting enables creative differentiation that may enhance message relevance. If research indicates that female audiences in a given product category place greater weight on quality testimonials and detailed ingredient information while male audiences respond more strongly to performance claims and comparative product attributes, a marketer can develop distinct creative executions for each segment and serve each group the message best calibrated to their decision-making style.

This form of gender targeting is less about who sees an ad than about what they see, and it raises fewer equity concerns than targeting that governs access to opportunities. A woman who sees a beauty advertisement that speaks to her interests is not harmed by the fact that a man sees a version of the same advertisement tailored to his preferences. The harm potential is low when the product is freely available to all, the targeting affects only the style of communication, and no user is actively excluded from opportunity.

4.3 PERSONALISATION AND THE CONSUMER VALUE PROPOSITION

There is also a consumer welfare argument for gender targeting. Personalised advertising, in general, has been associated with higher consumer satisfaction when it is experienced as relevant rather than intrusive (Bleier & Eisenbeiss, 2015). Gender-aware advertising that correctly anticipates a consumer's preferences can reduce the cognitive load of sorting through irrelevant offers and create a more streamlined path to purchase. In this frame, gender targeting is a form of service, not merely a marketing tool.



The consumer value proposition is, however, conditional on accurate inference and consumer consent. When gender inference is incorrect, a personalised experience becomes an alienating one. When consumers are unaware that gender is being used as a targeting signal, the personalisation may be experienced as surveillance rather than service. These conditions — accuracy and transparency — are prerequisites for realising the consumer welfare benefits of gender targeting, and both are frequently absent in current platform implementations.

5. ALGORITHMIC BIAS, HARM, AND STRUCTURAL INEQUALITY

5.1 EVIDENCE OF DISCRIMINATORY AD DELIVERY

The empirical case against unchecked gender-based algorithmic advertising is substantial and growing. The foundational work of Ali et al. (2019), conducted through Meta's advertising API, demonstrated that employment advertisements were delivered in highly gender-skewed ways even without explicit gender filtering. Advertisements for lumberjack positions reached over 90 per cent male audiences; advertisements for nursery positions reached over 95 per cent female audiences. These patterns closely mirrored existing occupational gender segregation, suggesting that platform algorithms were not merely reflecting user preferences but actively reproducing structural inequalities.

A critical feature of this research was its implication for equal employment opportunity law. In the United States, the Civil Rights Act of 1964 prohibits employment discrimination on the basis of sex, and its protections have been held to apply to employment advertising. If the ad delivery system consistently delivers job advertisements to one gender at the expense of another — regardless of advertiser intent — the platform may be enabling de facto discrimination in access to employment information, with consequences for labour market equity.

The US Department of Housing and Urban Development (HUD) filed a complaint against Facebook in 2019 on directly analogous grounds, alleging that Facebook's advertising platform enabled housing advertisers to exclude users on the basis of protected characteristics including sex, in violation of the Fair Housing Act. The complaint highlighted Meta's audience targeting tools, including the ability to target by geographic area and interest category in ways that correlate with protected characteristics even without explicit demographic targeting. Meta settled with HUD in 2022, agreeing to overhaul its ad targeting system for housing, employment, and credit advertisements.

5.2 THE FEEDBACK LOOP PROBLEM

A particularly insidious feature of algorithmic bias in advertising is its tendency towards self-reinforcement. If a delivery algorithm associates STEM job advertisements with male users based on historical engagement data, and therefore serves them predominantly to men, men will continue to engage with them at higher rates, reinforcing the algorithm's prior associations, and the cycle continues. Women, who were less likely to see the advertisement in the first place, have no opportunity to express their interest, so their non-engagement is recorded not as the absence of opportunity but as evidence of absent interest. The algorithm's discriminatory behaviour is thereby laundered into apparently neutral statistical observations.

This feedback loop problem means that algorithmic gender bias compounds over time. Initial biases — however they originated — become self-fulfilling predictions. The system optimises towards patterns it created. Correcting this requires not merely adjusting current parameters but actively introducing counterfactual exposure: showing advertisements to audiences who would not have received them under the existing model, observing their response, and incorporating that data into the training set. This is technically feasible but commercially unpopular, since it requires deliberately sacrificing short-term campaign efficiency for long-term equity.

5.3 INTERSECTIONAL HARMS AND COMPOUNDING DISADVANTAGE

The gender bias literature in advertising has increasingly converged with intersectionality scholarship to recognise that gender targeting harms do not fall evenly across all women or all men. Women from racially marginalised groups, women with disabilities, and women from lower-income backgrounds face compound disadvantage in ad delivery systems: they are both less likely to receive advertisements for high-value opportunities and more likely to receive advertisements for high-risk financial products such as payday loans (Noble, 2018; O'Neil, 2016). The interaction between gender and race in algorithmic targeting has been documented by multiple researchers and suggests that the harms of biased ad delivery are not evenly distributed across demographic subgroups.

Google's advertising platform has faced similar criticism. Research by Sweeney (2013) found that searching for names statistically associated with Black individuals was more likely to trigger advertisements for criminal record check services than searching for names associated with white individuals — a form of algorithmic discrimination that intersects race with the advertising ecosystem in ways that can stigmatise and disadvantage entire communities. While this specific finding concerns race rather than gender, it illustrates the general principle that platform algorithms can produce outcomes that are discriminatory across multiple protected dimensions simultaneously.

5.4 NON-BINARY AND GENDER NON-CONFORMING USERS

Both Meta and Google advertising platforms operationalise gender as a binary or near-binary variable, despite the well-established existence of non-binary, genderfluid, and gender non-conforming individuals. For these users, the platforms' targeting and inference systems are categorically inadequate: they are either misclassified into a gender category that does not reflect their identity or assigned to an 'Unknown' segment that may result in reduced advertising relevance or, in sensitive categories, potentially harmful targeting. Research on the advertising experiences of LGBTQ+ users has identified patterns of both under-targeting — exclusion from advertisements relevant to their needs — and harmful targeting — exposure to conversion-oriented or negatively stereotyping content (Haimson et al., 2021).

6. REGULATORY AND LEGAL LANDSCAPE

6.1 EUROPEAN UNION: GDPR, DSA, AND BEYOND

The European Union has developed the most comprehensive regulatory framework for digital advertising targeting in the world. The General Data Protection Regulation (GDPR), in force since 2018, establishes that processing personal data for advertising purposes requires either explicit consent or a demonstration of legitimate interest that is not overridden by the data subject's interests and rights. Gender, when used as an advertising targeting parameter, constitutes the processing of personal data, and advertisers using Meta or Google platforms in the EU are in principle subject to GDPR compliance requirements — including the right of access, the right to object to processing, and in some cases the right to erasure.

The Digital Services Act (DSA), which entered full force in February 2024, imposes additional requirements on very large online platforms — a category that encompasses both Meta and Google. Under the DSA, these platforms are prohibited from using targeting advertising based on sensitive categories of personal data, including gender in contexts where it constitutes a sensitive category. They are also required to provide users with a clear option to opt out of profiling-based advertising and to make a non-profiled version of their service available. Critically, the DSA establishes a transparency obligation: platforms must make their advertising systems legible to researchers and regulators, including disclosing targeting parameters and delivery outcomes.

The EU's Artificial Intelligence Act, adopted in 2024, includes advertising systems in its risk classification framework and imposes conformity assessment requirements on high-risk AI systems that determine access to employment, education, and essential services. To the extent that gender-targeted advertising systems are used in employment and credit advertising — and the evidence reviewed above suggests that they routinely are — they may be subject to these additional requirements, including the obligation to maintain logs of system decisions and to provide affected individuals with an explanation of automated decisions that significantly affect them.

6.2 UNITED STATES: SECTOR-SPECIFIC REGULATION AND ENFORCEMENT

The United States has not enacted comprehensive federal privacy legislation equivalent to GDPR as of 2024, though multiple states — including California (CCPA/CPRA), Colorado, Virginia, and others — have introduced privacy laws with data processing limitations that apply to advertising. At the federal level, sector-specific civil rights laws provide the primary regulatory framework for gender discrimination in advertising: the Fair Housing Act, the Equal Credit Opportunity Act, and Title VII of the Civil Rights Act all prohibit discrimination on the basis of sex in the contexts of housing, credit, and employment respectively.

The Federal Trade Commission (FTC) has also asserted jurisdiction over discriminatory algorithmic advertising under its authority to prohibit unfair or deceptive acts and practices. In 2022, the FTC issued a policy statement affirming that commercial surveillance practices — including targeted advertising — could constitute unfair practices under the FTC Act if they cause substantial injury to consumers in protected groups. This framing potentially opens a broader regulatory avenue for addressing algorithmic gender bias that does not depend on proving specific intent to discriminate.

Meta's 2022 settlement with HUD is the most significant enforcement action to date specifically addressing gender discrimination in digital advertising. As part of the settlement, Meta agreed to discontinue its Special Ad Audiences tool — which used look-alike modelling in ways that reproduced demographic skew — for housing, employment, and credit advertising, and to implement a new system designed to ensure proportional delivery across demographic groups in these categories. This settlement represents a significant, if partial and sector-limited, acknowledgement that platform-level algorithmic remediation is necessary.

6.3 PLATFORM SELF-REGULATION AND ITS LIMITS

In the absence of comprehensive regulation, both Meta and Google have adopted voluntary self-regulatory measures addressing gender-discriminatory advertising. Meta introduced a 'Special Ad Category' designation for housing, employment, and credit advertisements in 2019, restricting the targeting parameters available in these categories and prohibiting explicit age and gender targeting. Google similarly introduced restrictions on personalised advertising in the employment, housing, and credit categories. Both platforms have published Ads Transparency tools that allow users to see why they received a particular advertisement and to opt out of certain targeting categories.

However, researchers and civil society organisations have found these self-regulatory measures to be incomplete. Studies conducted after Meta's Special Ad Category restrictions took effect found that delivery algorithms continued to produce gender-skewed outcomes even with targeting restrictions in place, because the underlying optimisation logic was not constrained to produce equitable outcomes — only the advertiser-side inputs were modified (Sapiezynski et al., 2022). This finding is consistent with the structural account of algorithmic bias developed in Section 5 and underlines the limits of targeting-parameter restriction as a remediation strategy.

7. A CONCEPTUAL FRAMEWORK: THE STRATEGIC-ETHICAL TARGETING CONTINUUM

Having reviewed the theoretical, empirical, and regulatory dimensions of gender targeting in digital advertising, this section synthesises these perspectives into a conceptual framework that practitioners and researchers can use to evaluate specific targeting practices. The framework — which we term the Strategic-Ethical Targeting Continuum (SETC) — positions gender targeting practices along two intersecting dimensions: strategic utility (the measurable benefit to the advertiser in terms of campaign efficiency and message relevance) and ethical risk (the potential for harm to individuals and groups in protected categories).

7.1 THE FOUR QUADRANTS OF THE SETC

Plotting strategic utility on the horizontal axis and ethical risk on the vertical axis yields four quadrants, each characterising a distinct category of gender targeting practice.

Quadrant	High Strategic Utility	Low Strategic Utility
Low Ethical Risk	Quadrant 1 — Optimal Zone: Gender-tailored creative for freely available consumer products. e.g., beauty campaigns with different visuals for different audiences.	Quadrant 2 — Inefficiency Zone: Unnecessary gender targeting where no meaningful behavioural difference exists. Wastes budget, negligible harm.
High Ethical Risk	Quadrant 4 — Prohibited Zone: Gender-biased delivery in employment, credit, or housing contexts. High value to some advertisers; significant legal and moral harm to excluded groups.	Quadrant 3 — Harmful Waste: No strategic benefit, measurable harm. Worst practice. e.g., excluding women from tech job ads with no measurable conversion benefit.

The Optimal Zone (Quadrant 1) encompasses gender-aware creative customisation for products and services where purchase behaviour is legitimately differentiated by gender, no individual is excluded from an opportunity, and the targeting parameter does not operate as a proxy for a legally protected characteristic in a sensitive context. Practices in this quadrant are both commercially justified and ethically defensible.

The Prohibited Zone (Quadrant 4) encompasses any use of gender targeting — whether by explicit advertiser instruction or through platform delivery optimisation — that results in systematic exclusion of any gender from advertisements for employment, credit, housing, education, or other opportunity-granting contexts. Practices in this quadrant are commercially tempting (skewed delivery towards easier-to-convert audiences can improve short-term metrics) but ethically unacceptable and increasingly illegal.

The Inefficiency Zone (Quadrant 2) represents gender targeting that is applied in contexts where no meaningful gender-based behavioural difference exists, producing no strategic benefit while consuming targeting infrastructure resources. This practice is not harmful in the conventional sense but represents a misallocation of marketing resources and should be avoided on grounds of campaign efficiency.

The Harmful Waste category (Quadrant 3) represents the most indefensible form of gender targeting: practices that produce both negative ethical outcomes and no commercial benefit. These cases arise most commonly from unreflective copying of historical targeting conventions without testing whether those conventions remain applicable, or from algorithmic feedback loops that concentrate delivery in gender-homogenous patterns without any corresponding improvement in conversion performance.

7.2 APPLYING THE SETC IN PRACTICE

The SETC is useful not only as a classification tool but as a decision-making scaffold for marketing practitioners. Before implementing gender-based targeting in any campaign, a marketer might ask: (1) Is there credible evidence that gender correlates with purchase behaviour or message receptivity in this specific product category? (2) Does the product or service fall within a sensitive category — employment, credit, housing, healthcare, education — where gender-based exclusion creates legal or ethical exposure? (3) Is gender being used as a direct targeting parameter, an indirect proxy, or are we relying on platform delivery optimisation in ways we do not fully understand?

Only targeting practices that pass all three tests — where the answer to (1) is yes, to (2) is no, and to (3) indicates active visibility and control — sit comfortably within Quadrant 1. Practices that fail any of these tests should be redesigned, restricted, or abandoned.

8. RECOMMENDATIONS

8.1 FOR DIGITAL MARKETING PRACTITIONERS

- Conduct a gender targeting audit of all active campaigns to assess whether gender is used as a direct targeting parameter, an indirect proxy, or an implicit delivery optimisation signal, and document the justification for each use.
- Apply the SETC framework to classify each use of gender targeting and immediately discontinue any practices that fall within Quadrants 3 or 4, regardless of short-term metric performance.
- Resist the temptation to use gender as a default targeting variable simply because it is available. Require evidence of gender-differentiated consumer behaviour from credible research before applying gender-based segmentation, and test gender-neutral alternatives in parallel campaigns.
- Review creative content for unintended gender stereotyping that may reproduce inequality at the level of message even when targeting is neutral. An advertisement that algorithmically reaches all genders but depicts occupational success as exclusively male still contributes to harmful norm reproduction.
- Where platform tools allow delivery optimisation to override or expand specified audiences, monitor delivery reports closely and use demographic breakdown data to identify skewed delivery patterns. Flag anomalies to the platform and adjust bidding strategies to correct imbalances.

8.2 FOR PLATFORM DESIGNERS AND ENGINEERS

- Implement algorithmic fairness constraints in delivery optimisation systems for sensitive ad categories that enforce proportional gender representation in impressions across the accessible audience pool, not merely within advertiser-specified targeting parameters.
- Provide advertisers with real-time delivery transparency dashboards that display the gender distribution of actual impressions alongside specified targeting, enabling practitioners to identify and address delivery skew before it compounds.
- Move beyond the binary gender model in targeting architecture. Offer a fuller range of gender identity options consistent with users' declared profiles and ensure that targeting and inference systems do not systematically misclassify or marginalise non-binary users.
- Conduct ongoing algorithmic auditing — using independent external researchers where possible — to identify feedback loops and proxy discrimination in ad delivery systems, and publish findings in accessible formats as part of platform transparency commitments.
- Design opt-in mechanisms for gender-based targeting that give users genuine, frictionless control over whether their gender is used as an advertising signal, consistent with emerging regulatory obligations under the EU's Digital Services Act.

8.3 FOR POLICYMAKERS AND REGULATORS

- Extend equal opportunity advertising requirements to digital platforms comprehensively across all sensitive ad categories — not only housing, credit, and employment but also education, healthcare, and insurance — and fund enforcement capacity to match the scale of the digital advertising market.
- Mandate algorithmic impact assessments for advertising systems used in sensitive categories, requiring platforms to demonstrate that their delivery systems produce equitable outcomes across protected groups before deploying them commercially.
- Establish a research access framework that enables independent academic researchers to audit ad delivery outcomes systematically, without relying on platform data donations that may be incomplete or selectively disclosed.
- Develop clear, technically precise guidance on the boundary between permissible targeting customisation and prohibited discrimination in algorithmic ad delivery, reducing the regulatory uncertainty that currently allows platforms to claim that delivery outcomes are beyond their control.

9. CONCLUSION

Gender-based targeting in digital advertising is neither straightforwardly beneficial nor straightforwardly harmful. It is a practice of considerable strategic sophistication whose outcomes range from genuinely useful personalisation to systemically reproduced discrimination, depending on context, implementation, and platform-level design choices that are largely invisible to both advertisers and consumers. The evidence reviewed in this paper demonstrates that the strategic case for gender targeting is real and substantial in consumer product advertising where purchase behaviour is legitimately differentiated; that the harms of algorithmic gender bias are equally real and potentially severe in opportunity-granting advertising contexts; and that the current self-regulatory regime is insufficient to prevent the latter while preserving the former.

The Strategic-Ethical Targeting Continuum developed in Section 7 offers a practical tool for navigating this landscape. It situates individual targeting decisions within a framework that demands both commercial justification and ethical scrutiny, recognising that the two are not inherently in conflict but require deliberate alignment. Marketers who engage rigorously with both dimensions — who can demonstrate that their gender targeting generates genuine value, reaches audiences who benefit from it, and excludes no one from opportunity — occupy a defensible and commercially sustainable position. Those who use gender as an unreflective default, or who benefit from algorithmic delivery patterns they have not examined, occupy a position that is neither ethically sound nor, in the current regulatory climate, strategically safe.

The deeper question this paper raises is structural. Gender bias in advertising algorithms is not primarily a problem of bad actors making deliberate discriminatory choices; it is a problem of systems designed to maximise engagement perpetuating inequality because inequality is written into the historical data on which those systems are trained. Addressing it requires not only better advertiser behaviour but better platform architecture: delivery systems that treat equitable access as a constraint, not merely efficiency as an objective. This is technically achievable. Whether it will be achieved depends on the combined pressure of regulation, civil society scrutiny, and the growing recognition among sophisticated advertisers that campaigns built on discriminatory foundations carry reputational, legal, and commercial risks that outweigh any short-term targeting advantage.

For MBA students and practitioners engaging with digital marketing strategy, the core lesson is this: the power to target is also the power to exclude. Every targeting decision is simultaneously a communication decision and an access decision — a choice about who gets to know about an opportunity, a product, or a service. Exercising that power responsibly requires understanding not only the strategic mechanics of Meta and Google's advertising systems, but the social consequences that flow from the ways those systems are built and deployed. Marketing leadership in the digital age demands both.

References

- Ali, M., Sapiezynski, P., Bogen, M., Korolova, A., Mislove, A., & Rieke, A. (2019). Discrimination through optimization: How Facebook's ad delivery can lead to biased outcomes. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–30. <https://doi.org/10.1145/3359301>
- American Civil Liberties Union. (2015). Big data, meet big brother: How ad delivery systems can discriminate. *ACLU Report*. <https://www.aclu.org>
- Bettany, S., Dobscha, S., O'Malley, L., & Prothero, A. (2010). Moving beyond binary opposition: Exploring the tapestry of gender in consumer research and marketing. *Marketing Theory*, 10(1), 3–28.
- Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56–62.
- European Commission. (2022). Regulation (EU) 2022/2065 of the European Parliament and of the Council on a single market for digital services (Digital Services Act). *Official Journal of the European Union*.
- Federal Trade Commission. (2022). FTC policy statement regarding the scope of unfair methods of competition under Section 5 of the FTC Act. *FTC*. <https://www.ftc.gov>
- Fischer, E., & Arnold, S. J. (1994). Sex, gender identity, gender role attitudes, and consumer behaviour. *Psychology & Marketing*, 11(2), 163–182.
- GLAAD. (2022). Failing to protect LGBTQ users from discrimination: An assessment of Meta's advertising policies. *GLAAD*. <https://www.glaad.org>
- Haimson, O. L., Dame-Griff, A., Capello, E., & Richter, Z. (2021). Tumblr was a trans technology: The meaning, importance, history, and future of trans technologies. *Feminist Media Studies*, 21(3), 345–361.
- Jansen, B. J., Moore, K., & Carman, S. (2013). Evaluating the performance of demographic targeting using gender in sponsored search. *Information Processing & Management*, 49(1), 286–302.
- Lambrecht, A., & Tucker, C. E. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966–2981.
- Meta Business Help Centre. (2023). About audience targeting. *Meta for Business*. <https://www.facebook.com/business/help>
- Meyers-Levy, J., & Sternthal, B. (1991). Gender differences in the use of message cues and judgments. *Journal of Marketing Research*, 28(1), 84–96.
- Moorman, C. (2020). The CMO survey: Highlights and insights report. *Duke University Fuqua School of Business*.
- Moss, G., Gunn, R., & Heller, J. (2006). Some men like it black, some women like it pink: Consumer implications of differences in male and female website design. *Journal of Consumer Behaviour*, 5(4), 328–341.
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.

- O'Neil, C. (2016). Weapons of math destruction: How big data increases inequality and threatens democracy. Crown Publishers.
- Statista. (2024). Digital advertising revenue share worldwide in 2023, by company. Statista. <https://www.statista.com>