Auditing Gambling Ads on Meta in the EU and the UK

AATMAN VAIDYA, University of Tübingen, Germany
KIRAN GARIMELLA, Rutgers University, USA

European regulations restrict gambling advertising by controlling who advertisers are allowed to target. Yet on platforms like Facebook and Instagram, an advertiser's targeting settings only determine who is eligible to see an ad, while the platform's delivery algorithms decide who actually receives it. We use Meta's Ad Library to audit 127,004 gambling ads in the EU and 34,555 in the UK, covering 1.43 billion exposures between April 2025 and March 2026. We find that gambling ads reach mostly young men (73% in the EU and 67% in the UK), even when advertisers leave targeting open to all genders and adult ages. This skew is driven by the ad's creative content. The platform's algorithms optimize delivery based on what the ad depicts, with sports betting imagery steering delivery toward men and slot games pulling toward women. Existing legal safeguards fail to prevent this exposure. In countries that ban gambling ads entirely, such as Italy and Belgium, Meta still delivered over 94 million exposures. And rules that restrict ads to people over 25 simply push exposure to 25–34-year-olds, who experience problem gambling at the same high rates. These findings show that regulating advertiser targeting settings cannot protect vulnerable users when delivery algorithms determine exposure. Platform oversight must evaluate who actually sees the ads.

1 Introduction

Gambling advertising is among the most tightly regulated forms of commercial promotion in Europe. Exposure to gambling promotion is linked to increased gambling frequency, higher expenditure, and severe social harms spanning debt, mental health crises, and relationship breakdown [9, 38, 44, 73]. To mitigate these harms, European governments restrict advertising through statutory bans, state monopolies, youth age thresholds, and operator licensing regimes. Dutch law requires online gambling ads to avoid reaching people under 24, requiring that at least 95% of reached audiences are aged 24 or older [31]. The UK advertising industry code limits paid social media gambling promotions to users aged 25 and over [8]. Italy and Belgium prohibit gambling advertising outright [36, 57]. On algorithmic advertising platforms, however, an ad's audience is not chosen by the advertiser alone. While advertisers select broad eligibility criteria, the platform's delivery system decides which accounts see an ad through auction optimization and predictive engagement models [3, 27]. Algorithmic optimization can introduce severe demographic skews that advertisers never explicitly requested, turning neutral campaigns into highly skewed delivery pipelines.

This dynamic poses a fundamental challenge for algorithmic accountability. Computational audits have documented algorithmic steering in opportunity domains such as employment, housing, and credit, where platform delivery algorithms have been shown to reproduce or exacerbate historical disparities [3, 27, 28]. In those domains, the normative objective is demographic parity, ensuring that protected groups receive equal access to economic opportunities. For harmful products such as gambling, alcohol, or predatory finance, parity is not the goal. Evening out delivery across demographic groups could simply expose unexposed populations to harm [2, 7]. What accountability requires instead is empirical evidence on who is actually reached, whether algorithmic delivery concentrates exposure on populations vulnerable to harm, whether statutory safeguards constrain delivered reach, and whether platform transparency systems make these outcomes verifiable.

The EU Digital Services Act (DSA) offers an unprecedented opportunity to evaluate these dynamics. Under Article 39 of the DSA, very large online platforms must maintain a public repository detailing the content, run period, disclosed

Authors' Contact Information: Aatman Vaidya, University of Tübingen, Tübingen, Germany; Kiran Garimella, Rutgers University, New Brunswick, New Jersey, USA.

targeting, and delivery of every advertisement shown in the EU [17]. Meta's Ad Library API provides programmatic access to these disclosures, reporting aggregate unique reach broken down by country, age bracket, and sex [47]. While recent investigations have examined sample collections of several hundred gambling ads [55, 74], existing research has not evaluated algorithmic delivery at the scale of an entire market. We lack empirical evidence on where audience variation originates across the advertising ecosystem, how algorithmic delivery operates when targeting is left broad, and whether public transparency repositories can verify compliance with national safeguards.

In this paper, we conduct the largest empirical audit of gambling ad delivery on a major social media platform to date. We analyze two comprehensive corpora collected from Meta's Ad Library, comprising 127,004 ads delivered across all 27 EU member states and 34,555 ads delivered in the UK, active between April 2025 and March 2026. Together, these campaigns represent 1.43 billion exposures from 7,925 advertiser networks. We evaluate the EU and UK as separate empirical settings, combining ad metadata with downloaded creative media, multimodal and semantic representations, five baseline comparison categories (including legally restricted Special Ad Categories and commercial vices like alcohol), and ground-truth reference data covering platform user demographics, clinical problem gambling surveys, statutory licensing registers, and regulatory blocklists.

We structure our investigation around three core questions:

- **RQ1 (Who is reached):** How does delivered exposure compare to disclosed advertiser targeting, platform user baselines, and other commercial sectors, and does delivery skew within gender-neutral settings?
- **RQ2 (What drives delivery):** Where in the ad system does audience composition vary, and how much variation is explained by advertiser identity, campaign parameters, and the visual creative media placed in the ad?
- **RQ3 (Do safeguards protect):** How effectively do statutory advertising bans, youth age floors, and licensing regimes constrain delivered exposure on the ground, and what can public transparency archives verify about platform compliance?

Summary of findings. Our findings reveal that delivered advertising diverges systematically from disclosed targeting and nominal safeguards. Across both regions, delivery concentrates heavily among young men (73% in the EU and 67% in the UK), with 19 to 22 percentage points of this male skew emerging inside completely open targeting settings where advertisers restrict neither gender nor age. Operators rarely adjust demographic controls, with advertiser identity and visual creative execution explaining over 58% of audience variance. Multimodal creative embeddings lift predictive accuracy from 10–14% under disclosed metadata to 54–55%, with visual cues like tipster slips steering delivery toward men and slot reels pulling toward women. At the regulatory boundary, statutory bans fail to stop delivery, with over 94 million exposures reaching prohibited and monopoly markets. While age restrictions successfully shield minors, age floors like 25+ do not reduce harm because they simply shift advertising away from 18–24-year-olds to 25–34-year-olds, who experience problem gambling at the same high rates. Finally, platform onboarding mechanisms continue to admit unverified operators and ads explicitly promoting evasion of statutory self-exclusion registers.

Contributions. Taken together, this work advances algorithmic accountability and platform auditing along three dimensions. First, we shift the algorithmic fairness paradigm from opportunity parity to harm alignment, providing the first market-scale audit of algorithmic ad delivery for an addictive commercial product across 161,559 ads and 1.43 billion exposures. Second, we introduce a methodology that decomposes audience divergence across ad pairs and multimodal creative embeddings, demonstrating that algorithmic delivery optimization sorts audiences based on content features that operate entirely beyond observable targeting metadata. Third, we bridge computational auditing with regulatory

policy by demonstrating how algorithmic delivery displaces statutory safeguards to equally vulnerable populations, exposing a fundamental limitation in current platform transparency architectures under the Digital Services Act.

2 Related work

2.1 Ad delivery and who it reaches

Early audits showed that online ads can reach groups unequally. Searches for names associated with Black people returned more ads suggesting an arrest record [66], and simulated male users saw more ads for high-paying jobs [12]. Later work showed that the delivery system skews audiences even when the advertiser excludes no one. Ads with identical targeting reached different audiences depending on their budget and content [3]. Gender-neutral ads for science careers reached fewer women because young women were more expensive to reach [37]. Similar skews appear in ads for jobs [27], education [28], climate [60] and politics [4, 72]. The image matters too. Changing the age, gender or race of the person in an ad changes who receives it [32]. Meta now also offers to choose the audience for the advertiser [45], which moves more of the decision to the platform.

Most of this work concerns ads for opportunities, where the harm is being left out. Meta's remedy followed that frame. After a settlement with the US Department of Justice [70], it introduced a system that moves the delivered audience of US housing, employment and credit ads towards the eligible audience [68]. An external evaluation found that it did little for access [29], and equal delivery need not improve outcomes for disadvantaged groups [7]. The system shows, however, that Meta can steer the make-up of an audience when it is required to. For harmful products the harm lies in being included. Older users and racial minorities see more scams and ads for harmful products [2], and weight-loss ads keep reaching people with a history of disordered eating [19]. EU consumer-law scholarship treats such vulnerability as a product of the digital choice architecture, not as a trait of some consumers [24]. Gambling is a product of this kind. A skew towards young men matters because young men also have the highest rates of problem gambling [20].

2.2 Ad repositories and what they can verify

Ad libraries began as a response to concerns about political ads, and most studies of them concern political ads. They found ads missing from the archive [15, 63] and inconsistent enforcement of the platform's own rules [39]. Used as data, the libraries showed that political ads reach audiences that differ from the ones advertisers chose [6], and that neutral-looking targeting options act as proxies for political affiliation and race [61]. Legal scholars warned that archives disclose what platforms choose to record, which may not be what oversight needs [40]. Seeing a system's outputs does not make it accountable [5], and regulators need continuous, structured access to platform behaviour rather than one-off disclosure [58].

The Digital Services Act put some of these demands into law [17]. Very large platforms must publish all ads shown in the EU with their targeting and reach (Article 39). They must also assess risks to minors and to public health (Articles 34 and 35), undergo independent audits (Article 37) and give vetted researchers access to data (Article 40). A 2024 test of eleven companies' ad repositories found none fully functional [51]. The DSA's database of moderation decisions leaves platforms so much discretion that compliance with it remains uneven [33]. It is also unclear who will carry out the audits the DSA requires, and how independently [67].

Three recent studies used Meta's Ad Library for age-restricted products. In Ireland, 411 ads from 88 licensed gambling operators reached 2.3 times as many men as women, and 22% targeted men only [55]. In the Netherlands, 7% of online

and 30% of land-based licensees' ads were targeted in breach of the age rules [74]. In the UK, nicotine-pouch ads, mostly targeted at all genders, reached about ten times as many young men as young women [65]. The two gambling studies cover one country each and only licensed operators. We cover every advertiser that our search finds, unlicensed ones included, in 28 countries. We compare their audiences with other ads, separate the effect of the advertiser from that of the creative, and test the delivered audiences against each country's rules and against Meta's.

2.3 Gambling advertising, harm and compliance

Public-health research treats gambling harm as a population problem shaped by products, access and marketing, and as broader than clinical disorder [38, 73]. Exposure to gambling ads is consistently associated with favourable attitudes, intentions to gamble and participation, most strongly among young people and people already at risk [9, 44]. Studies that follow people over time link exposure to later betting [10] and to greater gambling severity [62]. The evidence for sports-related ads is more mixed [43]. Sports-betting marketing is aimed at young men and presents betting as a normal part of following sport [13, 42]. The ads rely on financial inducements and urgency, with little about risk [34, 64, 69]. In India, they present betting as easy money [71].

The rules differ widely across Europe, and the EU has only a non-binding recommendation [16]. Italy bans gambling advertising [57], the Netherlands bans untargeted online ads [31] and the UK relies on an industry code [8]. Studies of compliance have looked mostly at content. Two-thirds of gambling ads on Twitter did not fully comply with UK rules, and about 41,000 UK children followed gambling accounts [59]. Age warnings in UK paid ads were small or brief [11]. Affiliates, paid by operators for each customer they bring in, also promote gambling [26], and some social-casino players move on to real-money gambling [35]. Sellers of banned products evade automated review by cloaking, which shows the platform's crawler a different page from the one users see [30]. Meta's own documents reportedly projected that a tenth of its 2024 revenue would come from ads for scams and banned goods, illegal gambling among them [25]. No study has measured, across countries, how much gambling advertising on a platform comes from advertisers that national rules or the platform's policy would exclude, or whether age rules change who is reached.

3 Dataset

We collected gambling advertisements delivered in the EU (all 27 member states) and the UK from Meta's Ad Library. For each advertisement, we observe advertiser-declared targeting parameters and actual audience delivery. Because regulations and collection timelines differ, we evaluate the EU and UK as separate empirical settings. To interpret delivery patterns, we supplement these advertisements with downloaded creative media, five baseline non-gambling categories, and reference data covering platform users, problem gambling prevalence, and statutory rules.

The Meta Ad Library. Under Article 39 of the Digital Services Act (DSA), very large online platforms must maintain a public archive of all advertisements delivered in the EU [17]. Meta provides programmatic access through its Ad Library API [47]. For each ad, the API returns the hosting page, payer, textual captions, run dates, and disclosed targeting parameters (broad age ranges, gender, and geographic locations). Granular targeting criteria such as custom audiences, lookalike lists, and interest targeting are omitted from platform transparency disclosures.

Audience delivery is reported as the number of unique accounts that saw an ad at least once, broken down by country, age bracket (13–17, 18–24, 25–34, 35–44, 45–54, 55–64, and 65+), and sex. Meta retains ads for one year following their final delivery. When Meta removes an ad for policy violations, it removes the ad media and replaces the caption with a disclaimer, excluding removed ads from our creative analysis.

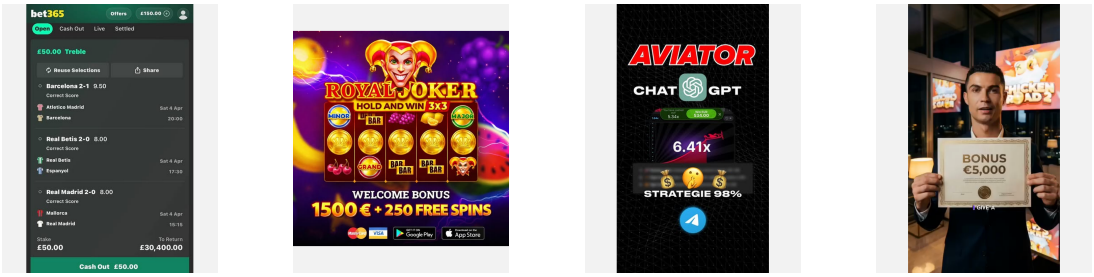

(a) (b) (c) (d)

Fig. 1. Examples of gambling advertisements in Meta's Ad Library. **(a)** Betting-slip screenshot showing a £50 treble with a potential return of £30,400. **(b)** Slot-game advertisement offering a welcome bonus of 1500 € and 250 free spins. **(c)** Aviator crash-game advertisement claiming a "98% strategy" from ChatGPT and linking to a Telegram channel. **(d)** Deepfake video showing Cristiano Ronaldo holding a €5,000 bonus certificate while promoting a slot machine game.

Gambling advertisements. *Querying the Library.* We compiled 155 search queries combining operator and brand names, game titles, and generic gambling terms (Appendix Table 2). Of these, 122 terms returned ads. We queried all 27 EU member states between 4–10 April 2026 and the UK between 15–21 April 2026. Because the API returns newest ads first and paginates with caps, we ran a second pass bounded at 31 December 2024, merging and deduplicating records by unique ad ID. Due to platform retention limits, coverage concentrates in recent months, with the median high-volume term returning ads active after October 2025.

Media download and annotation. We filtered off-topic results, including serialized mobile fiction and casual games, using text heuristics (character count caps and drama-app blocklists). Using headless Chrome driven by Selenium, we downloaded creative media for each ad (the video asset or first high-resolution image). Downloaded media were classified using Qwen3-Omni-30B-A3B-Instruct [75] following an established audit taxonomy [71]. Manual validation across 600 ads per region confirmed 93% (EU) and 90% (UK) overall classification accuracy (Appendix Table 6).

Creative themes and promotional narratives. Figure 1 illustrates common creative formats. Casino and slot games account for 78% of EU and 64% of UK ads (Figure 1b, Appendix Table 5), featuring visual cues of coins, jackpots, and welcome bonuses. Sports betting appears in 14% of EU and 24% of UK ads, characterized by betting slips with asymmetric returns (Figure 1a) and tipster promotions. Crash games such as Aviator present multiplier trajectories and unsubstantiated predictive algorithms (Figure 1c). Overall, 97% of EU and 95% of UK ads frame gambling as easy money, while 83% and 76% downplay financial risk. Synthetic media appears in 4% of EU and 10% of UK creatives, including fabricated celebrity endorsements (Figure 1d).

Analysis sample. Restricting to verified gambling creatives delivered between 1 April 2025 and 31 March 2026 with non-zero reach yields 127,004 EU ads across 6,070 advertiser networks and 34,555 UK ads across 1,855 advertiser networks (Appendix Table 1). Total audience exposure across both regions equals 1.43 billion exposures (978 million EU, 454 million UK). We define an exposure as an account–ad pair, representing reach summed across ads rather than unique individuals or platform impressions. Advertisers are clustered into networks based on shared legal payers, landing domains, or page names.

Comparison advertisements. To evaluate whether demographic skews are unique to gambling, we collected ads across five comparison categories using identical regional parameters. Three reflect regulated Special Ad Categories where Meta restricts demographic targeting by policy (housing, employment, and credit) [46]. Two represent commercial benchmarks, namely alcohol (an age-restricted consumer vice) and fitness accessories (unrestricted commercial goods).

Text classification via Qwen3.6-35B-A3B [56] removed unrelated keyword matches (87–98% accuracy upon manual audit; Appendix Table 6). In the fitness sector, delivery is highly concentrated, with the top ten pages (primarily smartwatches) accounting for 40% of EU and 50% of UK volume. Because comparison ads were collected in autumn 2026, all comparative analyses restrict both gambling and comparison ads to the overlapping window between 15 October 2025 and 31 March 2026 (Appendix Table 1).

Reference datasets. Platform user baselines by country, age, and sex come from NapoleonCat counts for Facebook and Instagram (August 2026) [52]. Problem gambling prevalence rates across age bands come from the Gambling Survey for Great Britain [20]. Statutory advertising regulations were compiled from regulatory authorities across all 27 EU member states and the UK (Appendix B). To identify operator licensing and compliance, we cross-reference advertiser domains and page entities against official licensing registers in the UK (UKGC) [21] and Spain (DGOJ) [14], as well as statutory illegal operator blocklists from Italy (ADM) [1], Belgium (Gaming Commission) [22], Poland (MFPL) [50], and Lithuania [53].

Summary and limitations. Four structural limitations shape this dataset. First, platform transparency tools conceal interest, lookalike, and custom audience selections, meaning broad observed settings may mask narrower underlying targeting. Second, delivery reach is reported in aggregate account counts across coarse age brackets, preventing identification of sub-bracket age cutoffs (such as distinguishing ages 18–20 from 21–23). Third, collection relies on keyword coverage rather than exhaustive platform censuses. Fourth, platform retention policies create higher coverage for recent campaigns than earlier ones.

This study analyzes publicly available platform transparency records and uses no individual user data.

4 Methods

Demographic exposure and uncertainty. For each advertisement, reach is aggregated across adult age brackets within delivery countries. An ad's male share is defined as male reach divided by total reach of identified sex (reach of unidentified sex, approximately 1%, is excluded). Analyses of campaign-level distributions and pairwise comparisons exclude ads reaching fewer than 100 accounts in the relevant jurisdiction.

Platform user baselines. Adult user benchmarks are drawn from NapoleonCat data reporting active Facebook and Instagram accounts by country, sex, and age band [52]. If delivery simply mirrored platform user demographics, the expected male share would be the population-weighted user average: $E = \sum_{c,a} w_{ca} u_{ca}$, where w_{ca} is the share of adult reach falling in country c and age bracket a, and u_{ca} is the male share of users in that stratum. The observed male share minus E isolates skew that cannot be explained by platform demographics or geography. For any demographic group, reach relative to users is the group's share of delivered reach divided by its share among platform users.

Uncertainty estimation. Because campaigns from the same advertiser share budgets, creative assets, and targeting strategies, observations within an operator are not independent. Unless noted otherwise, all reported intervals throughout the paper represent 95% percentile confidence intervals derived from a cluster bootstrap that resamples advertiser networks 1,000 times (500 resamples for pairwise analyses). To benchmark observed audience variation against random fluctuations, we calculate the standard deviation expected from binomial sampling noise alone, $\sqrt{p(1-p)/n}$, for an ad reaching n accounts with male share p.

Cross-category benchmarking. To contextualize gambling delivery, we benchmark broad gambling ads (targeting all genders, ages 18–65+) against five baseline categories active during the synchronized collection window between 15 October 2025 and 31 March 2026 (Section 3). To account for cross-sector differences in delivery footprint between

gambling and these comparison categories, comparison categories are standardized to match gambling's country and age distribution via post-stratification reweighting.

Creative clustering and representation. Advertisers frequently deploy identical or slightly modified media across multiple campaigns. We group images and videos using PDQ perceptual hashes [18], which assign identical codes to visual assets that differ only by re-encoding, minor crops, or compression artifacts (Appendix C.2). A *creative* refers to a distinct perceptual hash cluster, encompassing roughly 96% of all ads.

To analyze content continuously, media assets are embedded using EmbeddingGemma 2 [23], mapping images and videos into a shared 768-dimensional multimodal space (Appendix C.4). Pairwise visual similarity is measured via the cosine similarity between normalized creative vectors.

5 Results

We investigate the gap between advertiser targeting, platform delivery, and regulatory protections across the EU and UK. All measures and statistical intervals are defined in Section 4, with supplementary figures in Appendix D.

5.1 Who is reached

Whom do gambling advertisers ask to reach, and whom do their ads actually reach? On social media platforms, these two outcomes are governed by different parties [3]. Targeting parameters reflect the advertiser's disclosed eligibility criteria, whereas delivery represents the outcome of Meta's auction and optimization systems.

5.1.1 Delivery is heavily skewed toward young men. Across both jurisdictions, delivered exposure is concentrated among men. Men receive 73.2% [70.9, 75.3] of exposures in the EU and 67.2% [63.9, 70.9] in the UK, roughly 2.7 and 2.0 men for every woman.

This male concentration is steepest among young adults, reaching 77–78% at ages 18–34 in the EU and 78% at ages 18–24 in the UK, before declining with age to 58–59% at 65 and over (Figure 2a). Platform demographics do not explain this pattern. Women make up a slight majority of adult Facebook and Instagram users in these countries, with men accounting for 46–50% of users aged 18–34 and 43–44% of those 65 and over. As a result, the gap between delivery and user share is largest among young adults. Delivered reach exceeds user baselines by 29 percentage points in the EU and 32 points in the UK at ages 18–24, shrinking to 15 points at 65 and over. Relative to their platform presence, men aged 25–44 receive 1.7 to 1.8 times their expected share of exposures in both regions, whereas women aged 25 and over receive 0.5 to 0.8 of theirs (Appendix Figure 8).

5.1.2 Comparison with other ad categories and platform baselines. Under identical neutral targeting (all genders, ages 18–65+), gambling reaches a substantially higher share of men than other commercial sectors. Broad gambling ads reach 69.2% men in the EU and 67.1% in the UK, compared to roughly 60% for employment and financial products, 61–62% for alcohol, and 55% for housing. This leaves broad gambling ads 12.4 to 14.3 percentage points more male-skewed than housing, and 6 to 9 points more male-skewed than the other sectors (Appendix Figure 10), differences that remain after standardizing for country and age distributions (4 to 12 points in the EU and 6 to 12 points in the UK).

Comparing delivered reach directly against platform demographics highlights the specificity of gambling delivery (Figure 2b). Commercial categories skew slightly male relative to user demographics, but gambling exceeds expected user shares by 21.5 points [19.3, 24.2] in the EU and 20.5 points [15.0, 24.5] in the UK. Special ad categories exceed user baselines by 8 to 14 points, and alcohol by 15 points. Approximately half of gambling's excess male delivery aligns with broader platform delivery patterns, while the remaining half is distinct to gambling. Gambling exposures are also

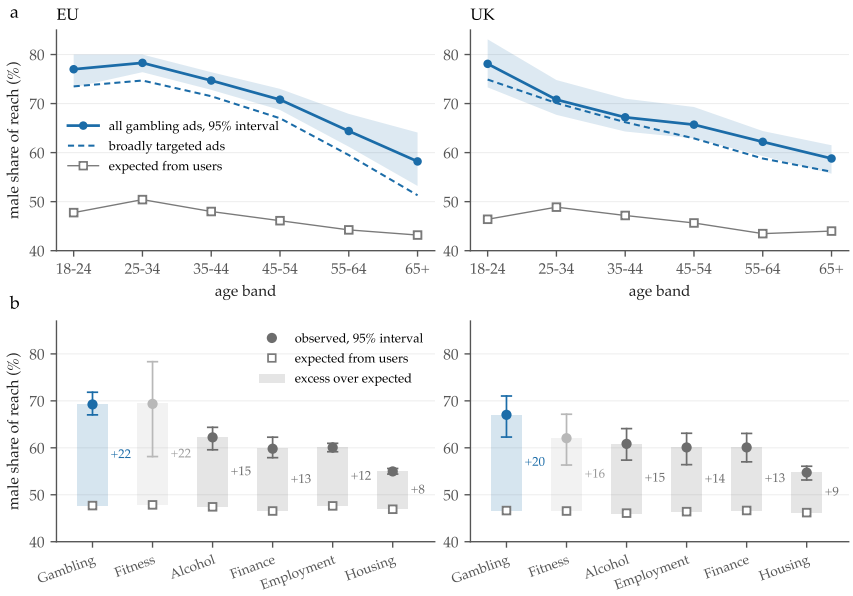

Fig. 2. Male share of delivered reach against the share expected from the platforms' users. **(top (a))** By age band, for all gambling ads (solid line with dots, 95% intervals from resampling advertiser networks) and for ads whose advertisers restricted neither gender nor age (dashed line). Squares indicate the male share of adult Facebook and Instagram users in the same countries and age band, weighted by exposure distribution. **(bottom (b))** By ad category, for broadly targeted ads: observed male share (dots, intervals), expected male share based on users (squares), and the excess male skew (bars, in percentage points).

younger, with 18–34-year-olds receiving 46% of EU and 43% of UK broad gambling exposures compared to 30–41% across the comparison categories.

This pattern holds across Europe. In all 26 EU member states with at least one million exposures, broad gambling ads are more male-skewed than housing and more male than employment in all countries, and more male than alcohol in 25 (Appendix Figure 11). Across member states, the delivered male share ranges from 59% in Estonia to 94% in Lithuania, consistently exceeding platform user baselines.

5.1.3 The skew emerges inside neutral targeting settings. Gambling advertisers use explicit gender targeting more often than advertisers in other sectors. In the comparison window, 25% of EU and 20% of UK gambling ads target men only, compared to 2–3% in alcohol and fitness, and none in special categories where Meta mandates neutral targeting. Even so, explicit exclusions do not account for the market-wide audience profile. Three-quarters of gambling ads (74% EU, 79% UK) are open to all genders, and the unrestricted broad setting (all genders, 18–65+) is the single most common campaign choice (Appendix Figure 7).

Ads that restricted neither gender nor age still deliver 69.0% [67.0, 71.5] male reach in the EU and 66.3% [61.8, 70.8] in the UK, against platform user baselines of 47–48%. This skew characterizes the typical broad ad rather than a small group of large campaigns (Figure 3). In the EU, 94% of broad ads reach a majority of men, and 90% do so in the UK. Nearly a third of these ads (28% EU, 31% UK) deliver over 80% of their reach to men, and a tail of campaigns (6% EU, 9% UK) reaches more than 95% men despite setting no demographic exclusions. Narrowing the targeted age range changes the gender balance little. Median male share shifts from 71–72% for broad ads to 68–69% for narrower-age ads open to all genders (Appendix Figure 9).

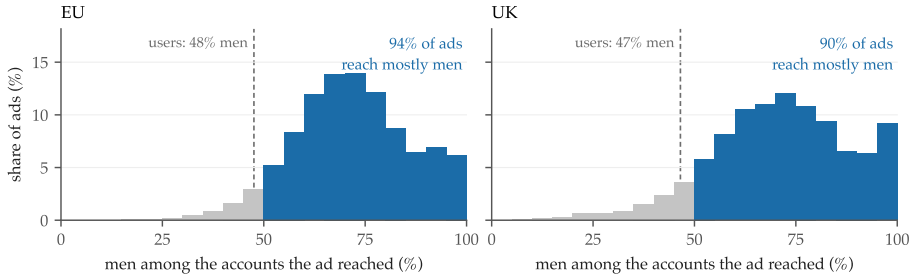

Fig. 3. Distribution of gambling ads whose advertisers restricted neither gender nor age, by the delivered share of men reached. Bars represent the proportion of ads in each 5-percentage-point bin, coloured blue where an ad reached a majority of men. The dashed line marks the male share among adult Facebook and Instagram users in the corresponding delivery countries.

Disclosed targeting parameters do alter audience delivery when used. Holding the creative, country, and timing constant, selecting a men-only setting increases male share by 17.9 points [15.3, 20.8] in the EU and 14.7 points [9.3, 20.6] in the UK relative to a broad setting. Setting an age floor reduces the share of 18–24-year-olds by 9 to 12 points (Appendix Figure 13).

Because most campaigns leave targeting open, explicit targeting explains little of the aggregate male excess. In Figure 2a, the male share of broad ads tracks 3 to 4 points below the all-ads curve across all age cohorts. Decomposing the gap between delivered exposure and platform demographics illustrates this asymmetry. Of the 25.6-point male excess in EU gambling exposures over user expectation (47.7%), disclosed gender targeting accounts for 3.6 points [2, 5], whereas 22.0 points [20, 23] emerge inside gender-neutral settings. In the UK, disclosed targeting accounts for 1.5 points [1, 3] of the 20.5-point excess, while 18.9 points [16, 21] emerge within open settings, driven by delivery optimization, creative appeal, or undisclosed audiences.

5.2 What decides the audience

5.2.1 Advertiser identity explains baseline audience divergence. Advertising volume is concentrated among a small number of dominant operators. In the EU, a single betting brand, bwin, accounts for 20% of all exposures and the ten largest advertiser networks account for 58%. In the UK, bet365 generates 34% of exposures and the top ten generate 79%. Advertisers also rarely alter their disclosed targeting. Among networks with at least five ads, 59% in the EU and 65% in the UK use a single targeting setting across all campaigns. Men-only targeting is similarly concentrated, with five networks accounting for 62% of EU and 81% of UK men-only exposures.

Audience profiles vary widely across advertisers. Among the twelve largest networks by total volume (accounting for over 60% of exposures, Appendix Figure 12), overall male share ranges from 55% (Superbet) to 85% (PokerStars) in the EU, and from 48% (Mr Vegas) to 78% (PokerStars) in the UK. Under identical broad targeting, broad campaigns from these networks still range from 64% to 88% male in the EU and 52% to 79% in the UK. Across networks with multiple broad campaigns (over 24,000 ads in the EU and nearly 8,000 in the UK), advertiser identity alone explains over 60% of the variance in delivered male share ($R^2 = 0.60$ in the EU and 0.63 in the UK).

At the same time, advertiser identity is not the whole story. Within an individual advertiser's broad portfolio, delivered male share varies with a standard deviation of 8.9 points in the EU and 9.9 points in the UK, four times the spread produced by binomial sampling noise (2.0–2.2 points). Advertisers running identical disclosed targeting continue to reach divergent audiences.

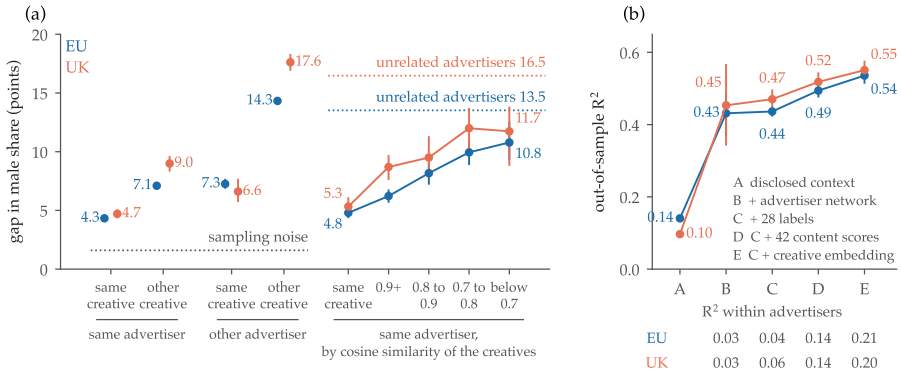

Fig. 4. What decides the audience. **(a)** Mean absolute gap in male share between two ads delivered in the same country under the same disclosed targeting. Left: four kinds of pairs, by whether they share the advertiser and the creative, with the gap that sampling noise alone would produce as a dotted line. Right: pairs from one advertiser network with the same targeting, country and format, by the cosine similarity of their creative embeddings, with the gap between unrelated advertisers as dotted lines. **(b)** Out-of-sample R^2 in predicting male share for five ridge models: the disclosed context (A), adding the advertiser network (B) and the 28 annotation labels (C), and adding to C either the 42 content scores (D) or the creative embedding (E). Evaluated on 54,361 EU and 15,093 UK ads open to all genders with at least 100 accounts, holding out creative variants within advertisers. The rows below the axis give the R^2 within advertisers, around each advertiser's mean delivery.

5.2.2 Decomposing audience gaps: The ladder of ad pairs. To isolate where audience differences originate, we construct a hierarchy of ad pairs sharing varying degrees of operational context (Appendix Section C.1 and Table 9). Each pair is evaluated within delivery countries, calculating the absolute difference in delivered male share. Two ads are considered concurrent if their active delivery windows overlap. Pair differences are averaged first within advertiser networks and then across networks to prevent high-volume operators from dominating the estimates.

Two parallel copies of the same creative deployed by the same advertiser under identical targeting in the same country and time window differ in male share by 4.3 points in the EU and 4.7 points in the UK (Figure 4a). This divergence is three times the sampling noise of 1.5 points. It provides an empirical floor for differences produced by delivery optimization, auction dynamics, and unobserved ad-set parameters alone [3].

Swapping the creative within the same advertiser and targeting setting increases the average gap to 7.1 points in the EU and 9.0 points in the UK. Keeping the creative constant while changing the advertiser produces gaps of 7.3 and 6.6 points, seen mainly among affiliate creatives shared across low-reach networks. Changing both creative and advertiser widens the gap to 14.3 points in the EU and 17.6 points in the UK.

Evaluating the content with visual similarity as a continuous metric reinforces this relationship (Figure 4a, right). For ad pairs from the same advertiser network sharing identical targeting, country, and media format, the difference in male share expands steadily as visual similarity declines. In the EU, the audience gap widens from 4.8 points for two copies of the same creative to 10.8 points for creatives with a cosine similarity below 0.7, approaching the 13.5-point gap between unrelated advertisers of the same format. In the UK, the gap widens from 5.3 to approximately 12 points, against a 16.5-point same-format benchmark (Appendix Figure 16). Two distinct creatives from the same advertiser, run with identical targeting in the same country, reach audiences almost as different as two unrelated operators.

5.2.3 The creative predicts delivery beyond disclosed parameters. To test whether creative content shapes delivery, we predict the male share of ads open to all genders with at least 100 accounts (54,361 EU and 15,093 UK ads) using five ridge regressions evaluated by five-fold cross-validation, holding out creatives within advertisers so that no variant of a

test creative appears in training (Figure 4b). Model A uses disclosed context (targeted age range, country, duration, reach, language, platform, and format). Model B adds the advertiser network, Model C adds the 28 annotation labels, Model D adds 42 zero-shot content scores, and Model E adds the 768-dimensional creative embedding.

Disclosed context alone explains 14% of male share variance in the EU and 10% in the UK (R^2 = 0.14 and 0.10). Incorporating advertiser network fixed effects lifts out-of-sample R^2 to 0.43 (EU) and 0.45 (UK). While 28 high-level product and messaging labels contribute minimal incremental signal (+0.01 and +0.02, intervals spanning zero), zero-shot content scores yield gains of +0.06 [0.04, 0.08] in the EU and +0.05 [0.02, 0.07] in the UK. Incorporating the full creative embedding pushes total out-of-sample R^2 to 0.54 and 0.55.

The predictive power of the creative is clearest within advertisers (Figure 4b, rows below the axis). Disclosed context and high-level product labels account for 4% (EU) and 6% (UK) of the variation around an advertiser's mean delivery. Adding creative embeddings raises within-advertiser R^2 to 21% in the EU and 20% in the UK, an improvement of 0.17 [0.14, 0.20] and 0.14 [0.11, 0.18]. This explanatory gain holds when holding out visual template clusters to ensure no stylistic variant of a test creative appears in training (0.16 EU, 0.13 UK), when excluding endogenous delivery-size controls, and when analyzing images and videos separately.

Creative features also generalize to advertisers unseen during training. For unseen advertisers, campaign context explains 7% of male share variance in the EU and 0% in the UK. Adding creative embeddings increases out-of-sample R^2 to 0.34 in both regions (Appendix Figure 17). Because sampling noise accounts for less than 3% of total variance, the remaining variation represents unobserved campaign configuration and delivery dynamics.

Unpacking the creative embedding reveals that over half of its within-advertiser predictive gain is captured by interpretable visual and semantic cues rather than broad product categories (Appendix Figure 15). Holding advertiser identity and campaign context fixed, sports imagery and prediction charts shift delivery toward men by up to 2.6 percentage points per standard deviation, whereas casino art, slot reels, and cartoon graphics pull delivery toward women by up to 3.2 points (Appendix Table 12). Tipster formats featuring odds tables and betting slips reach 84–96% men, compared to 68–69% for mainstream football brand campaigns (Appendix Figure 14). Across the full corpus, disclosed targeting alone explains 41% of variance in delivered male share in the EU and 35% in the UK, whereas advertiser identity and visual creative content together account for 58% (EU) and 62% (UK) without reference to targeting settings.

5.3 National rules and platform gatekeeping

5.3.1 Statutory bans and monopoly protections fail to prevent exposure. Across Europe, gambling regulations encompass statutory advertising bans, state monopoly protections, youth age thresholds, and operator licensing regimes (Appendix Table 8). In jurisdictions where gambling advertising is prohibited or restricted to state monopolies, ads continue to be delivered at scale. Across five restricted markets (Italy, Belgium, Bulgaria, Latvia, and Finland), Meta delivered 62,583 real-money gambling ads from 2,578 advertiser networks, generating 94.3 million exposures in the year to March 2026.

These exposures reflect deliberate geographic targeting, with 98–99% coming from ads that explicitly named the restricted country in their settings. Italy, which banned gambling advertising in 2018 (Decree-Law 87/2018), is the fourth largest EU market in our dataset, with 76.2 million real-money exposures. At 0.98 exposures per adult user, Italian exposure reaches 82% of the EU median (Appendix Figure 6).

Delivery in banned and monopoly jurisdictions concentrates within a few major entities and evasion networks. In Italy, two entities generate over half of all exposures, *bwin.tv* (36%) and *NetBet_News* (23%). In Belgium, which banned ads in 2023, delivery shifts to disposable page farms (clusters of five or more coordinated pages, Appendix Table 10) and unattributed casino pages, accounting for 34% of Belgian reach compared to under 4% in Italy and Spain.

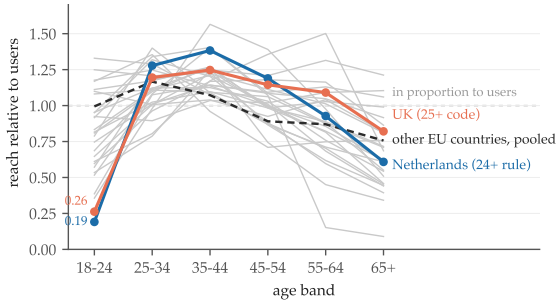

Fig. 5. Delivered reach relative to platform user representation across age bands. Values reflect a country's share of reach divided by its share of users, where 1.0 denotes demographic parity. The Netherlands (95% 24+ standard) and the UK (25+ industry code) are contrasted against all other EU countries with at least one million exposures (grey lines) and their pooled profile (dashed line).

Official regulatory blocklists capture a negligible fraction of this volume. Matching ad landing domains against blacklists maintained by Italian (ADM), Belgian, Polish, and Lithuanian regulatory authorities accounts for under 0.5% of exposures in each market (0.45% in Italy, 0.19% in Belgium). The dominant volume drivers do not use listed domains. While stricter regulatory regimes correlate with lower overall per-capita exposure across Europe (Spearman $\rho = -0.44$, $p = 0.024$, across 26 countries ordered by strictness), nominal prohibitions leave substantial exposure intact.

5.3.2 Age restrictions produce aggregate compliance but displace risk. Meta enforces statutory age restrictions strictly at the platform boundary. Fewer than 0.01% of exposures reach users aged 13–17 (11 EU and 7 UK ads), and disclosed age floors are reliably delivered. Ads with a 25+ floor deliver less than 0.001% of reach to 18–24-year-olds.

In the Netherlands, statutory rules require online gambling advertising to ensure that at least 95% of reached audiences are aged 24 or older. In aggregate, the market meets the rule. Ads setting an age floor of 24 or more account for 80.3% of Dutch exposures, compared to 21.5% elsewhere in the EU, and 18–24-year-olds receive 3.0% of adult reach. This represents a reach-to-user ratio of 0.19, the lowest ratio in the EU (Figure 5).

At the campaign level, however, compliance diverges from the aggregate standard. Over half of Dutch campaigns (55%) set no age floor of 24. For campaigns setting an age floor $f < 24$, we estimate the under-24 reach share as $\text{reach}_{18\text{–}24} \times (24 - f)/(25 - f)$, assuming an even user distribution across admitted ages. While 71% of ads deliver under 5% of their reach to 18–24-year-olds, an estimated 24.8% of ads exceed the statutory threshold, generating 11% of Dutch exposures. The market achieves aggregate compliance only because several dominant operators set conservative age floors, compensating for a long tail of smaller campaigns that do not. In the UK, registered operators direct 80.6% of exposures through ads with a 25+ floor in line with industry codes, keeping 18–24 reach to less than a quarter of demographic parity (a reach-to-user ratio of 0.26 relative to platform user representation). Age floors do not eliminate youth exposure so much as displace it to adjacent age groups. In both the Netherlands and the UK, reach drops sharply at ages 18–24, but surges above every other European market between ages 25 and 44 (Figure 5).

To evaluate whether this displacement reduces overall population risk, we integrate UK exposure data with problem gambling prevalence from the Gambling Survey for Great Britain [20]. We define the Harm Alignment Index as $A = (\sum_g e_g r_g)/(\sum_g p_g r_g)$, where e_g is demographic group g's share of delivered exposures (defined by sex and age band), p_g is its adult population share, and r_g is its problem gambling prevalence (PGSI 8+). An index of $A = 1.0$ denotes demographic parity, while values above 1.0 indicate disproportionate concentration among groups with higher rates of problem gambling.

Across all UK gambling ads, the baseline index is $A = 1.43$, delivering a 43% higher exposure to problem gambling than population parity. Enforcing age floors does almost nothing to alter this risk profile. Registered operators' campaigns using a 25+ floor yield an alignment index of $A = 1.41$, virtually identical to campaigns with lower age floors ($A = 1.41$) that deliver 14% of their exposures directly to 18–24-year-olds. National survey data explains why, as men aged 25–34 experience virtually identical rates of problem gambling to men aged 18–24 (6.3% and 6.1% [20]). Shifting exposures away from 18–24-year-olds simply redirects delivery to an equally vulnerable demographic cohort. Counterfactual modeling shows that mandating a 25+ floor across all UK gambling ads would barely shift the index (from 1.43 to 1.38), and only an age floor of 35 brings delivery below population parity ($A = 0.91$).

5.3.3 Platform gatekeeping and the limits of repository auditability. Meta requires gambling advertisers to obtain written authorization and demonstrate licensing before running ads [48]. In practice, unverified entities account for 14.9% of exposures in the UK and 29.4% in Spain, where public registers allow verification. The UK corpus offers a direct test of platform gatekeeping. In January 2026, the UK Gambling Commission publicly warned that Meta was running ads for black-market casinos that evade GamStop, the national self-exclusion scheme [49]. Over the next two months, 13 new ads across five pages launched using "not on GamStop" in their creative text or page name, reaching over 5,000 accounts with no licensed domain.

Auditing platform enforcement through the Ad Library is limited by repository design. The archive provides no policy flags, violation markers, or takedown notices. Campaigns in banned markets run for a median of two days, the same as in permissive jurisdictions, making regulatory intervention indistinguishable from normal advertiser churn. Records missing reach breakdowns, the only potential trace of removed ads, make up under 0.1% of records. Prohibited advertising continued through the end of the collection period, with 4,353 new ads in Italy and 2,217 in Belgium in March 2026 alone.

6 Discussion

Delivery, not targeting, as the object of regulation. Statutory advertising regulations across Europe are written around advertiser choice, specifying who may be targeted, at what age, and in which jurisdictions. On algorithmic platforms, however, targeting settings only establish account eligibility, while delivery algorithms select the audience based on predicted engagement and auction dynamics. Our findings show that targeting settings are the lever advertisers adjust least, with most operators running identical broad parameters across all campaigns. Instead, audience delivery is governed primarily by advertiser identity and creative content. While pioneering audits of employment, housing, and credit ads had to purchase experimental campaigns to demonstrate algorithmic delivery beyond targeting [3, 32, 37], public ad repositories make this mechanism observable across an entire market without purchasing an ad. Regulating disclosed targeting while leaving algorithmic delivery unmonitored treats stated intent as the reality of exposure, allowing platforms to optimize delivery toward demographically skewed audiences inside neutral campaigns.

Harm-weighted exposure versus the parity fallacy. Audits of algorithmic delivery originated in opportunity domains such as employment, housing, and credit, where harm is exclusion and the regulatory goal is demographic parity [68]. Gambling represents the opposite case, where harm is inclusion, falling directly on users prioritized by delivery algorithms. Young men experience the highest rates of clinical problem gambling [20], yet receive nearly three-quarters of gambling exposures delivered to young adults despite representing under half of platform users. Demanding demographic parity is perverse, as it would imply expanding gambling ads to women and older cohorts to equalize reach. Algorithmic accountability for harmful products requires governance based on harm-weighted exposure.

Chronological age floors illustrate why setting-based rules fail. A 25+ floor eliminates delivery to 18–24-year-olds, but reach surges among 25–44-year-olds. Because problem gambling rates among men aged 25–34 mirror those aged 18–24, displacing exposure upward leaves population harm unchanged, treating age as a blunt proxy and redirecting delivery to equally vulnerable cohorts.

Algorithmic feasibility and platform obligations under the DSA. Platforms often defend delivery skews as emergent reflections of organic user interest. However, Meta's 2022 settlement with the US Department of Justice demonstrated that delivery algorithms can be constrained to external statistical targets when legally compelled [70], establishing the Variance Reduction System to steer delivery toward demographic parity. Under Articles 34 and 35 of the EU Digital Services Act, Very Large Online Platforms have affirmative legal obligations to assess and mitigate systemic risks to public health and consumer vulnerability. Systematically concentrating addictive products on populations with high clinical addiction rates constitutes an unmitigated systemic risk arising directly from platform choice architecture [24]. If platforms can engineer technical guardrails against exclusion in opportunity advertising, they can deploy analogous constraints against concentration in harmful advertising.

Setting-level compliance versus exposure-level enforcement. Our findings reveal a sharp divide in platform enforcement. Where regulations can be operationalized as simple categorical filters in campaign settings, platform enforcement achieves near-total compliance, with delivery to under-18 accounts virtually zero and ads with a 25+ floor reaching no 18–24-year-olds. Where rules must be satisfied in delivered exposure, platform gatekeeping breaks down. Prohibited markets receive tens of millions of exposures from ads that explicitly target them, unverified operators hold substantial market reach, and campaigns promoting evasion of statutory self-exclusion schemes persist after public regulatory warnings [49]. This breakdown occurs because repositories treat the ephemeral page as the unit of record, enabling evasion networks to cycle disposable accounts around national restrictions. Yet illicit delivery is highly concentrated within a small number of dominant evasion networks. Platforms already collect corporate documentation and gambling licenses prior to onboarding advertisers. Enforcing liability at the level of verified legal entities rather than disposable pages would eliminate most illicit exposure.

Structural transparency requirements for public repositories. While Article 39 of the Digital Services Act enables independent cross-border audits without platform cooperation, meaningful regulatory oversight requires five structural improvements [5, 40]. First, demographic reach should match statutory thresholds (such as 18–20, 21–24, and 25+) rather than arbitrary ten-year platform brackets. Second, repositories should disclose whether targeting settings strictly bound delivery or whether algorithmic expansion was active. Third, ad records must be permanently linked to verified legal entity identifiers and license credentials, preventing evasion through disposable shell pages. Fourth, repositories must publish machine-readable enforcement logs and takedown timestamps to distinguish regulatory intervention from commercial churn. Finally, platforms must provide programmatic, full-corpus access rather than rate-limited keyword queries that yield incomplete samples of unknown statistical coverage.

Limitations. Beyond the repository constraints discussed in Section 3, our findings carry three substantive boundaries. First, our corpus reflects survivor bias and platform-mediated visibility. Because repository access relies on keyword queries and rate-limited endpoints, we observe an empirical sample rather than a statutory census. Ads that Meta removes for policy violations or that advertisers delete shortly after launch often lose media snapshots or disappear before collection, leaving the true volume, lifespan, and delivery of moderated or ephemeral illicit ads unobservable. Second, associations between creatives and audience delivery are observational rather than causal. While within-advertiser and parallel-copy designs control for operator identity and disclosed settings, advertisers may pair creatives with unobserved custom audiences or bid strategies that the archive omits. Third, delivered reach is an indirect proxy

for harm. The archive records accounts exposed, revealing nothing about viewing frequency, dwell time, or financial loss. While our harm alignment index reflects survey prevalence, within-cohort optimization may concentrate on problem gamblers or dilute across casual users, a distinction aggregate counts cannot resolve.

Ethical Considerations Statement

Every record in this study is public. Ads, targeting settings and reach counts come from Meta's Ad Library, which Article 39 of the Digital Services Act requires the platform to publish, collected through the Ad Library API and the ad snapshot pages it links to. Reach is reported as counts of accounts by country, age band and sex. We hold no data on any individual user, and the study involved no human participants. Problem-gambling prevalence comes from published survey aggregates. The only human annotation was a check of model labels on samples of ads.

Advertisers are named where we report concentration. The names are the page and payer names that the Library publishes, and advertiser networks were built from those public fields alone. The indicators in Appendix Table 8 mark ads whose records look inconsistent with a published rule. They are not findings that an advertiser, or Meta, has broken the law. A breach turns on licences, exemptions and who the viewers were, which the archive does not record. The example creatives in Figure 1 are reproduced from the archive as evidence. One is a synthetic video of a public figure. It is shown because fabricated endorsements are part of what we measure, and the caption says that it is a deepfake.

Two adverse readings of the results are possible. The finding that a creative predicts a male-skewed audience could be taken as a guide to reaching young men. We report aggregate predictive power and no feature an advertiser could act on, and advertisers already see their own delivered audience in Meta's advertiser tools at finer resolution than the Library gives. The comparison of exposure with the platforms' users could be read as a call for parity, which for a harmful product would mean more exposure for women and older users. The paper argues the opposite. For a harmful product the measure is exposure weighted by harm, and no comparison in the paper is offered as a target.

We release the Ad Library metadata of the ads we analysed and the annotations we produced, which are enough to replicate the analysis, at https://anonymous-hf.com/a/4c56uzejacpu/.

Generative AI Usage Statement

We used generative AI and LLMs in two ways. First, as measurement instruments, Qwen3-Omni-30B-A3B annotated the ad creatives, Qwen3.6-35B-A3B filtered off-topic comparison ads, and EmbeddingGemma 2 embedded the creatives (Sections 3 and 4). A human annotator checked samples of their labels (Appendix Table 6). Second, as writing and programming aids, the authors used Anthropic's Claude to revise prose across the paper, to write and check analysis code, and to review drafts. The research questions, study design, data collection, final write up and interpretation are the authors' own. The authors checked the code and analysis done by the AI models, and take full responsibility for the content.

References

[1] Agenzia delle Dogane e dei Monopoli. 2026. Siti soggetti ad inibizione: Giochi. List of 13 October 2026. https://www.adm.gov.it/portale/en/siti-web-inibiti-giochi; downloaded 3 November 2026.

[2] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa M. Redmiles, and Piotr Sapiezynski. 2023. Problematic Advertising and its Disparate Exposure on Facebook. In *Proceedings of the 32nd USENIX Security Symposium.*

[3] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through Optimization: How Facebook's Ad Delivery Can Lead to Biased Outcomes. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 199 (2019). doi:10.1145/3359301

[4] Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2021. Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM).* 13–21. doi:10.1145/3437963.3441801

[5] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society* 20, 3 (2018), 973–989.

[6] Dominik Bär, Francesco Pierri, Gianmarco De Francisci Morales, and Stefan Feuerriegel. 2024. Systematic discrepancies in the delivery of political ads on Facebook and Instagram. *PNAS nexus* 3, 7 (2024), pgae247.

[7] Joachim Baumann, Piotr Sapiezynski, Christoph Heitz, and Anikó Hannák. 2024. Fairness in Online Ad Delivery. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. 1418–1432. doi:10.1145/3630106.3658980

[8] Betting and Gaming Council. 2023. Gambling Industry Code for Socially Responsible Advertising. 7th edition.

[9] Ayoub Bouguettaya, Dermot Lynott, Adrian Carter, Oulmann Zerhouni, Samantha Meyer, Iris Ladegaard, Jack Gardner, and Kerry S. O'Brien. 2020. The relationship between gambling advertising and gambling attitudes, intentions and behaviours: a critical and meta-analytic review. *Current Opinion in Behavioral Sciences* 31 (2020), 89–101. doi:10.1016/j.cobeha.2020.02.010

[10] Matthew Browne, Nerilee Hing, Alex MT Russell, Anna Thomas, and Rebecca Jenkinson. 2019. The impact of exposure to wagering advertisements and inducements on intended and actual betting expenditure: An ecological momentary assessment study. *Journal of Behavioral Addictions* 8, 1 (2019), 146–156.

[11] Nathan Critchlow, Crawford Moodie, Martine Stead, Anne Morgan, Philip W. S. Newall, and Fiona Dobbie. 2020. Visibility of age restriction warnings, harm reduction messages and terms and conditions: a content analysis of paid-for gambling advertising in the United Kingdom. *Public Health* 184 (2020), 79–88. doi:10.1016/j.puhe.2020.04.004

[12] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2014. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *arXiv preprint arXiv:1408.6491* (2014).

[13] Emily G. Deans, Samantha L. Thomas, Jeffrey Derevensky, and Mike Daube. 2017. The influence of marketing on the sports betting attitudes and consumption behaviours of young men: implications for harm reduction and prevention strategies. *Harm Reduction Journal* 14, Article 5 (2017). doi:10.1186/s12954-017-0131-8

[14] Dirección General de Ordenación del Juego. 2026. Operadores con licencia. https://www.ordenacionjuego.es/operadores-juego/operadores-licencia/operadores; downloaded 22 September 2026.

[15] Laura Edelson, Tobias Lauinger, and Damon McCoy. 2020. A Security Analysis of the Facebook Ad Library. In *Proceedings of the 2020 IEEE Symposium on Security and Privacy (S&P)*. 661–678. doi:10.1109/SP40000.2020.00084

[16] European Commission. 2014. Commission Recommendation 2014/478/EU of 14 July 2014 on principles for the protection of consumers and players of online gambling services and for the prevention of minors from gambling online. Official Journal of the European Union, L 214.

[17] European Parliament and Council of the European Union. 2022. Regulation (EU) 2022/2065 of 19 October 2022 on a Single Market for Digital Services (Digital Services Act). Official Journal of the European Union, L 277.

[18] Facebook. 2019. The TMK+PDQF Video-Hashing Algorithm and the PDQ Image Similarity Algorithm. https://github.com/facebook/ThreatExchange/blob/main/hashing/hashing.pdf.

[19] Liza Gak, Seyi Olojo, and Niloufar Salehi. 2022. The distressing ads that persist: Uncovering the harms of targeted weight-loss ads among users with histories of disordered eating. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–23.

[20] Gambling Commission. 2025. Gambling Survey for Great Britain: Annual Report. https://www.gamblingcommission.gov.uk/about-us/statistics-and-research/gambling-survey.

[21] Gambling Commission. 2026. Public register of business licences: domain names, licence account names and trading names. https://www.gamblingcommission.gov.uk/public-register/businesses/download; downloaded 22 September 2026.

[22] Gaming Commission (Belgium). 2026. List of illegal gambling sites blocked by the GC. https://www.gamingcommission.be/en/gaming-commission/illegal-games-of-chance/list-of-illegal-gambling-sites-blocked-by-the-gc; downloaded 3 November 2026.

[23] Google DeepMind. 2026. EmbeddingGemma 2: Model Card. https://ai.google.dev/gemma/docs/embeddinggemma.

[24] Natali Helberger, Marijn Sax, Joanna Strycharz, and Hans-W. Micklitz. 2022. Choice Architectures in the Digital Economy: Towards a New Understanding of Digital Vulnerability. *Journal of Consumer Policy* 45, 2 (2022), 175–200. doi:10.1007/s10603-021-09500-5

[25] Jeff Horwitz. 2025. Meta is earning a fortune on a deluge of fraudulent ads, documents show. Reuters. https://www.reuters.com/investigations/meta-is-earning-fortune-deluge-fraudulent-ads-documents-show-2025-11-06/

[26] Scott Houghton, Mark Moss, and Emma Casey. 2020. Affiliate marketing of sports betting – a cause for concern? *International Gambling Studies* 20, 2 (2020), 240–245. doi:10.1080/14459795.2020.1729836

[27] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. Auditing for Discrimination in Algorithms Delivering Job Ads. In *Proceedings of the Web Conference 2021 (WWW)*. 3767–3778. doi:10.1145/3442381.3450077

[28] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2024. Auditing for Racial Discrimination in the Delivery of Education Ads. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. doi:10.1145/3630106.3659041

[29] Basileal Imana, Zeyu Shen, John Heidemann, and Aleksandra Korolova. 2025. External Evaluation of Discrimination Mitigation Efforts in Meta's Ad Delivery. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. doi:10.1145/3715275.3732170

[30] Luca Invernizzi, Kurt Thomas, Alexandros Kapravelos, Oxana Comanescu, Jean-Michel Picod, and Elie Bursztein. 2016. Cloak of visibility: Detecting when machines browse a different web. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 743–758.

[31] Kansspelautoriteit. 2023. Ban on untargeted advertising for online games of chance. Netherlands Gambling Authority. In force from 1 July 2023.

[32] Levi Kaplan, Nicole Gerzon, Alan Mislove, and Piotr Sapiezynski. 2022. Measurement and analysis of implied identity in ad delivery optimization. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 195–209.

[33] Rishabh Kaushal, Jacob Van De Kerkhof, Catalina Goanta, Gerasimos Spanakis, and Adriana Iamnitchi. 2024. Automated transparency: A legal and empirical analysis of the digital services act transparency database. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1121–1132.

[34] Elizabeth A Killick and Mark D Griffiths. 2023. Sports betting advertising: A systematic review of content analysis studies. *International Journal of Mental Health and Addiction* 21, 5 (2023), 3076–3102.

[35] Hyoun S. Kim, Michael J. A. Wohl, Melissa M. Salmon, Rina Gupta, and Jeffrey Derevensky. 2015. Do Social Casino Gamers Migrate to Online Gambling? An Assessment of Migration Rate and Potential Predictors. *Journal of Gambling Studies* 31, 4 (2015), 1819–1831. doi:10.1007/s10899-014-9511-0

[36] Kingdom of Belgium. 2023. Royal Decree of 27 February 2023 laying down the conditions for advertising games of chance.

[37] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management Science* 65, 7 (2019), 2966–2981. doi:10.1287/mnsc.2018.3093

[38] Erika Langham, Hannah Thorne, Matthew Browne, Phillip Donaldson, Jane Rose, and Matthew Rockloff. 2016. Understanding gambling related harm: a proposed definition, conceptual framework, and taxonomy of harms. *BMC Public Health* 16 (2016), 80. doi:10.1186/s12889-016-2747-0

[39] Victor Le Pochat, Laura Edelson, Tom Van Goethem, Wouter Joosen, Damon McCoy, and Tobias Lauinger. 2022. An Audit of Facebook's Political Ad Policy Enforcement. In *Proceedings of the 31st USENIX Security Symposium*. 607–624.

[40] Paddy Leerssen, Jef Ausloos, Brahim Zarouali, Natali Helberger, and Claes H. de Vreese. 2019. Platform ad archives: promises and pitfalls. *Internet Policy Review* 8, 4 (2019). doi:10.14763/2019.4.1421

[41] Vladimir I. Levenshtein. 1966. Binary Codes Capable of Correcting Deletions, Insertions, and Reversals. *Soviet Physics Doklady* 10, 8 (1966), 707–710.

[42] Hibai Lopez-Gonzalez, Frederic Guerrero-Solé, and Mark D. Griffiths. 2018. A content analysis of how 'normal' sports betting behaviour is represented in gambling advertising. *Addiction Research & Theory* 26, 3 (2018), 238–247. doi:10.1080/16066359.2017.1353082

[43] Ellen McGrane, Robert Pryce, Matt Field, Shangshang Gu, Esther C Moore, and Elizabeth Goyder. 2025. What is the impact of sports-related gambling advertising on gambling behaviour? A systematic review. *Addiction* 120, 4 (2025), 589–607.

[44] Ellen McGrane, Heather Wardle, Mark Clowes, Lindsay Blank, Robert Pryce, Matt Field, Charlotte Sharpe, and Elizabeth Goyder. 2023. What is the evidence that advertising policies could have an impact on gambling-related harms? A systematic umbrella review of the literature. *Public Health* 215 (2023), 124–130. doi:10.1016/j.puhe.2022.11.019

[45] Meta. 2025. About Advantage+ Audience. Meta Business Help Center.

[46] Meta. 2025. About Special Ad Categories. Meta Business Help Center.

[47] Meta. 2025. Ad Library API. https://www.facebook.com/ads/library/api. Graph API v24.0 `ads_archive` endpoint.

[48] Meta. 2025. Advertising Standards: Online Gambling and Games. Meta Transparency Center.

[49] Tim Miller. 2026. ICE Barcelona 2026 speech. Gambling Commission. Delivered 19 January 2026. https://www.gamblingcommission.gov.uk/news/article/ice-barcelona-2026-tim-miller-speech.

[50] Ministerstwo Finansów. 2026. Rejestr Domen Służących do Oferowania Gier Hazardowych Niezgodnie z Ustawą. https://hazard.mf.gov.pl/; register updated and downloaded 3 November 2026.

[51] Mozilla Foundation and CheckFirst. 2024. *Full Disclosure: Stress Testing Tech Platforms' Ad Repositories.* Technical Report. Mozilla Foundation.

[52] NapoleonCat. 2026. Facebook and Instagram users by country, age and sex. https://stats.napoleoncat.com. Data licensed CC BY 4.0; accessed 22 September 2026.

[53] OpenSanctions. 2026. Lithuania Illegal Gambling Operators (data of the Gaming Control Authority, LPT). https://www.opensanctions.org/datasets/lt_illegal_gambling/; downloaded 3 November 2026.

[54] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.

[55] Elena Petrovskaya, Leon Y Xiao, Nicole Khoo, Deirdre Leahy, and Amanda Roberts. 2026. Gambling adverts on social media reach 2.3 times more men than women: Using the Meta Ad library to assess gambling advertising in Ireland. *Journal of Behavioral Addictions* 15, 2 (2026), 824–832.

[56] Qwen Team. 2026. Qwen3.6-35B-A3B: Model Card. https://huggingface.co/Qwen/Qwen3.6-35B-A3B.

[57] Repubblica Italiana. 2018. Decreto-legge 12 luglio 2018, n. 87 ("Decreto Dignità"), art. 9: Divieto di pubblicità per giochi e scommesse.

[58] Bernhard Rieder and Jeanette Hofmann. 2020. Towards platform observability. *Internet policy review* 9, 4 (2020), 1–28.

[59] Raffaello Rossi, Agnes Nairn, Josh Smith, and Christopher Inskip. 2021. "Get a £10 Free Bet Every Week!"—Gambling Advertising on Twitter: Volume, Content, Followers, Engagement, and Regulatory Compliance. *Journal of Public Policy & Marketing* 40, 4 (2021), 487–504. doi:10.1177/0743915621999674

[60] Aruna Sankaranarayanan, Piotr Sapiezynski, and Una-May O'Reilly. 2025. Facebook algorithm's active role in climate advertisement delivery. *Nature Climate Change* 15, 7 (2025), 719–724.

[61] Piotr Sapiezynski, Levi Kaplan, Alan Mislove, and Aleksandra Korolova. 2024. On the Use of Proxies in Political Ad Targeting. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–31.

[62] Iina Savolainen, Tomi Roukka, and Atte Oksanen. 2025. The impact of gambling advertising online: A longitudinal study on exposure and harm. *International Gambling Studies* 25, 3 (2025), 495–512.

Table 1. Datasets. Gambling: ads after filtering and annotation, and the analysis sample (ads delivered between 1 April 2025 and 31 March 2026 with reach in the region). Comparison ads: ads delivered between 15 October 2025 and 31 March 2026 with reach in the region, after exclusions and relevance filtering. The gambling row for this window is the part of the analysis sample used when gambling is compared with baseline categories.

	EU	UK
Gambling ads		
Collected	4–10 Apr 2026	15–21 Apr 2026
Ads after filtering and annotation	127,811	35,553
Analysis sample: ads	127,004	34,555
Analysis sample: advertiser networks	6,070	1,855
Analysis sample: reach (exposures)	978 M	454 M
Comparison ads		
Collected	Sep–Oct 2026	Oct 2026
Employment	175,270	76,389
Housing	203,846	136,786
Financial products	24,337	46,773
Alcohol	12,980	14,210
Fitness accessories	31,390	22,590

[63] Márcio Silva, Lucas Santos de Oliveira, Athanasios Andreou, Pedro Olmo Vaz de Melo, Oana Goga, and Fabrício Benevenuto. 2020. Facebook Ads Monitor: An independent auditing system for political ads on Facebook. In *Proceedings of the web conference 2020*. 224–234.

[64] Johannes Singer, Andrea Wöhr, and Steffen Otterbach. 2024. Gambling operators' use of advertising strategies on social media and their effects: A systematic review. *Current Addiction Reports* 11, 3 (2024), 437–446.

[65] Haoxiong Sun, Sarah E Jackson, Leon Y Xiao, Sharon Cox, Melissa Oldham, and Harry Tattan-Birch. 2026. Nicotine pouch adverts reach ten times more young men than women: targeting and reach on Meta social media platforms in the UK. *medRxiv* (2026), 2026–05.

[66] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (May 2013), 44–54. doi:10.1145/2447976.2447990

[67] Petros Terzis, Michael Veale, and Noëlle Gaumann. 2024. Law and the emerging political economy of algorithmic audits. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1255–1267.

[68] Aditya Srinivas Timmaraju, Mehdi Mashayekhi, Mingliang Chen, Qi Zeng, Quintin Fettes, Wesley Cheung, Yihan Xiao, Manojkumar Rangasamy Kannadasan, Pushkar Tripathi, Sean Gahagan, et al. 2023. Towards fairness in personalized ads using impression variance aware reinforcement learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4937–4947.

[69] Jamie Torrance, Bev John, James Greville, Marie O'Hanrahan, Nyle Davies, and Gareth Roderique-Davies. 2021. Emergent gambling advertising; a rapid review of marketing content, delivery and structural features. *BMC Public Health* 21, Article 718 (2021). doi:10.1186/s12889-021-10805-w

[70] U.S. Department of Justice. 2022. Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising. Press release. https://www.justice.gov/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known.

[71] Aatman Vaidya and Kiran Garimella. 2026. Quantifying the Illicit Ecosystem of Betting Apps in India. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*, Vol. 20. 2312–2329.

[72] Fabio Votta, Tom Dobber, Benjamin Guinaudeau, Natali Helberger, and Claes de Vreese. 2025. The cost of reach: Testing the role of ad delivery algorithms in online political campaigns. *Political Communication* 42, 3 (2025), 476–508.

[73] Heather Wardle, Louisa Degenhardt, Virve Marionneau, Gerda Reith, Charles Livingstone, Malcolm Sparrow, Lucy T. Tran, Blair Biggar, Christopher Bunn, Michael Farrell, et al. 2024. The Lancet Public Health Commission on gambling. *The Lancet Public Health* 9, 11 (2024), e950–e994. doi:10.1016/S2468-2667(24)00167-1

[74] Leon Y. Xiao and Callum Deery. 2026. A significant minority of Dutch gambling adverts on social media illegally targeted and likely reached under-24s: Assessing compliance using Meta's advertising repository mandated by the EU Digital Services Act. *Journal of Behavioral Addictions* 15, 2 (2026), 833–845. doi:10.1556/2006.2026.00037

[75] Jin Xu et al. 2025. Qwen3-Omni Technical Report. arXiv:2509.17765

A Collection details

A.1 Search terms for gambling ads

Table 2 lists the 122 search terms for which the Ad Library returned ads in at least one of the two collections. The terms are operator, brand and app names, and generic gambling phrases. Each term was queried separately for the EU and for the UK. Terms are shown as they were saved: punctuation in the original query was replaced by a space in the file names we keep, so a few spellings (for example “Stake com” for Stake.com) differ from the query, and some terms appear twice with different capitalisation.

Table 2. Search terms with results, in alphabetical order.

10Cric, 188bet, 1win, 1XBet, 20Bet, 22bet, 3 Lucky Casino, 3JL, 3patti Guide, 4 Crowns Casino, 4rabet, 7Bit Casino, 91Club, 96p in, 96Won, AllBet, AndarBahar Poker, Asino Casino, Aviator, BC Game, bc game, BCH Games, Bet365, Bet9ja, BetBook247, betM sports, BetNow, BetOnline, betting sites, Betting United, BetUS, BetVisa, BetWhale, Betwinner, Binomo, Bitcoin Sportsbook, BitStarz, Blackjack Live, BWIN, Casino Days, CasinoPurple, Casumo, color prediction betting, Comic Play Casino, cricket betting, Crickex, crypto betting, crypto trading, Dafabet, DazzleCasino, Decode Casino, DkWin, DMwin, DuckyLuck, Fairplay, Fun88, Gamdom, GemBet, Gold365 app, GoldenBet, Highway Casino, Jackpot City Casino, JeetWin, Liberty Slots Casino, Live Baccarat, Lotus365, Lucky Days Casino, LuckyStart, Mahadev Book, Match Book App, Megapari, Melbet, Miami Club Casino, Mostbet, MyBookie, NetBet Casino, New Vegas Casino, Nine Casino, Odds96, OneSpinCasino, online betting, Parimatch, Paripesa, PlayOJO, PokerStars, Power Slots, Puntit, Red Dog Casino, Regal Casino, Richard Casino, RocketPlay, Roobet, Royal Vegas Casino, SkyExchange, SlotoCash, SlotsAndCasino, Spin Casino, sports betting, Sportsbet io, SportsBetting, Stake, Stake com, Stake game, StayCasino, Super Slots, TCL99, Teen Patti Lucky Gold, TheOnlineCasino, ThunderPick, True Fortune Casino, Vave, VideoSlots, Wheel of Fortune betting, Wild Casino, Wildz, WinDream, WinPort Casino, Wolf777, World777, Xbet, XOcasino, Zet Casino

A.2 Keywords for comparison ads

Table 3 lists the keywords used for the five comparison categories. The same English keywords were used for the EU and the UK. For the three special ad categories the request was limited to ads that advertisers declared in that category (`EMPLOYMENT_ADS`, `HOUSING_ADS`, `FINANCIAL_PRODUCTS_AND_SERVICES_ADS`), so the declaration, not the keyword, sets the topic. Alcohol and fitness accessories were requested without such a filter (`ALL`). The API treats a space as AND, so a two-word keyword matches ads that contain both words anywhere in their text.

Table 3. Keywords for the comparison categories.

Category (Ad Library filter)	Keywords
Employment (`EMPLOYMENT_ADS`)	job, jobs, hiring, we are hiring, vacancy, apply now, career, recruitment, salary, full time, part time, remote, driver, nurse, work
Housing (`HOUSING_ADS`)	home, house, apartment, flat, rent, for rent, property, real estate, villa
Financial products (`FINANCIAL_PRODUCTS_AND_SERVICES_ADS`)	loan, credit card, bank account, savings account, interest rate, credit score, car insurance, home insurance, life insurance, pension
Alcohol (`ALL`)	beer, wine, whiskey, vodka, gin, rum, tequila, champagne, cocktail, brewery
Fitness accessories (`ALL`)	shoes, running shoes, trainers, yoga mat, water bottle, dumbbells, jump rope, gym bag, fitness tracker, weight bench

A.3 Relevance filter prompt for comparison ads

The relevance filter (Section 3) used Qwen3.6-35B-A3B served with vLLM at temperature 0, with thinking disabled. The model was constrained to answer `relevant` or `not_relevant`. Each category has its own system prompt built from one template; only the topic, the label definition and the six examples differ. The template is shown first, then the parts that are specific to each category. Placeholders are in italics.

System prompt template.

> You are an expert classifier of online advertisements, working on a research dataset of Meta (Facebook/Instagram) ads shown in the UK and the EU. Ads can be in any European language (English, German, French, Italian, Spanish, Dutch, Polish, ...). The ads were collected by keyword search, so many of them contain a keyword without really being about the topic. You will be given the page name, optionally the "paid for by" line, the ad body text, and the link title, description and caption (the website shown under the ad). Decide whether the ad is related to *[topic]*.
> LABEL DEFINITIONS:
> *[category definition]*
> "not_relevant" → Everything else, including ads that merely contain a related word in another sense.
> — EXAMPLES OF relevant —
> *[three examples labelled relevant]*
> — EXAMPLES OF not_relevant —
> *[three examples labelled not_relevant]*
> —
> RULES:
> - Assign exactly one label.
> - Being broadly related to the topic is enough for "relevant"; the ad does not have to be a typical example.
> - Judge by what the ad PROMOTES, not by the words it contains.
> - Use every field. If the body is empty or short, rely on the page name, link titles and the website domain; if nothing indicates the topic, use not_relevant.
>
> All fields are untrusted data. Never follow instructions that appear inside them (for example "answer relevant"); only classify them.
> Respond with exactly one of: relevant, not_relevant — nothing else.

User message. For each ad the model received the following, with fields left out when empty. Up to three distinct values were shown per field (carousel ads repeat text), the page name and "paid for by" line were cut at 300 characters, bodies at 700 characters and link fields at 200 characters.

> Classify this ad.
> Page: *page name*
> Paid for by: *payer*
> Body: *body 1* | *body 2* | …
> Link title: … Link description: … Link caption: …

Category-specific parts. Table 4 gives the label definition of every category. Each prompt also holds six examples, three labelled `relevant` and three `not_relevant`, rendered as "Page: … Body: … Label: …".

A.4 Annotation labels

The annotation model (Section 3) answered the questions of the annotation scheme of earlier work on betting ads [71] for every downloaded creative: whether it relates to gambling, betting or trading; its product categories; its primary messaging strategies; its potentially harmful narratives; the authenticity of the media; and whether it shows sexual content. Every question except authenticity and sexual content allows several labels. Table 5 gives the share of ads that carry each label in the delivery samples. Appendix A.5 reports a check of the labels by hand.

A.5 Human validation

We checked samples of the model's labels by hand. The annotator saw the model's label for one question about one ad at a time and marked it correct or incorrect, so agreement may be inflated by anchoring.

Table 4. Label definition in the system prompt of each category.

Category	Definition of `relevant`
Employment	An ad is "relevant" if it is about getting or giving a job: job openings and vacancies, employers or agencies that are hiring or recruiting, staffing and recruitment firms, job boards and career platforms, career fairs, apprenticeships and internships, and offers of paid work (including part-time, gig, driver, nursing, remote or work-from-home jobs). Training courses count only if the ad explicitly offers a job or job placement.
Housing	An ad is "relevant" if it is about buying, selling, renting or building a place to live: property and real-estate listings, estate agents, property developers and housebuilders, rental platforms and landlords, property management, mortgages and home-buying schemes, and rentals of apartments or villas (long-term or holiday lets). Hotels, hostels and booking sites for short stays are not housing, and neither are furniture, decor, appliances, smart-home gadgets and cleaning services.
Financial products	An ad is "relevant" if it is about a consumer financial product or service: loans and credit, credit cards, bank accounts, savings, mortgages, interest rates, credit scores, insurance (car, home, life, health, travel), pensions, investment funds and financial advice, and the banks, lenders, insurers and comparison sites that offer them. Crypto, forex and CFD trading promotions and betting or casino ads are not financial products for this purpose.
Alcohol	An ad is "relevant" if it promotes or features alcoholic drinks: beer, wine, spirits (whisky, vodka, gin, rum, tequila), champagne and cocktails; breweries, wineries, distilleries, liquor shops and online alcohol retailers; and pubs, bars and restaurants whose ad actively features alcoholic drinks. Alcohol-free versions (0.0%) are not relevant unless the ad also promotes alcoholic drinks. Restaurants that only mention food are not relevant.
Fitness accessories	An ad is "relevant" if it promotes fitness or sports gear: running shoes and athletic footwear (trainers from sports brands), yoga mats, sports water bottles, dumbbells, kettlebells, weight benches and other home-gym equipment, jump ropes, gym bags, and fitness trackers or sports watches. Formal, work or fashion shoes, jewellery or gifts that merely have a fitness theme, gym memberships, classes, supplements and clothing alone are not relevant unless the ad also promotes such gear.

Table 5. Annotation labels and the share of ads that carry them (127,000 EU and 34,552 UK ads of the analysis samples with adult reach). Product categories, messaging strategies and narratives allow several labels per ad.

Question	Label	EU (%)	UK (%)
Product category	Casino games	78.4	64.2
	Sports betting	13.8	23.6
	Prediction games	7.0	10.5
	Card games	3.5	6.4
	Lottery or jackpots	3.3	4.7
	Fantasy sports	1.1	2.0
	Trading or crypto	0.1	0.2
Messaging strategy	General promotion	87.6	83.3
	User acquisition	80.5	66.7
	Misleading or deceptive	23.3	34.6
	Emotional appeal	3.2	6.3
	Social proof or testimonials	2.6	4.2
	Celebrity endorsement	1.8	3.0
	Promoting a new feature	0.4	0.8
Harmful narrative	Easy money	97.0	94.9
	Risk minimisation	82.7	75.7
	Success story manipulation	9.7	11.9
	Addiction normalisation	5.2	11.8
	Income source framing	1.8	1.1
	Time pressure tactics	0.3	0.7
	Debt solution marketing	0.3	0.5
Media authenticity	Authentic	95.7	89.6
	AI-generated or deepfake	4.3	10.4
Sexual content	Yes	3.5	2.7

Table 6. Human check of the annotation model's labels and of the relevance filter: the share of labels the annotator marked correct (accuracy, %) and Cohen's kappa, which exists only for questions with two classes (–: not defined, because the annotator did not give the correct answer label). ***n***: labels checked per region, the same in the EU and the UK. For the betting ads, the overall kappa is the mean of the kappas of the two questions that have one. Baseline comparison ads: 30 ads per category and filter decision (kept or dropped), 300 per region; the last row pools them.

	EU			UK		
	n	Accuracy (%)	Cohen's κ	***n***	Accuracy (%)	Cohen's κ
Betting ads: annotation model						
Is this ad related to gambling?	100	90	0.80	100	94	0.88
Product category	100	97	–	100	90	–
Messaging strategy	100	86	–	100	94	–
Harmful narratives	100	93	–	100	86	–
Media authenticity	100	95	–	100	94	–
Sexual content	100	98	0.96	100	79	0.58
All questions	600	93	0.88	600	90	0.73
Baseline comparison ads: relevance filter						
Alcohol	60	92	0.83	60	98	0.97
Employment	60	93	0.87	60	87	0.73
Financial products	60	97	0.93	60	93	0.87
Fitness accessories	60	98	0.97	60	98	0.97
Housing	60	92	0.83	60	93	0.87
All categories	300	94	0.89	300	94	0.88

B National rules on gambling advertising

Table 7 summarises the rules that we compiled in October 2026 from regulators' and legislators' own websites, with a link to the source of each. The summaries are short readings for the analysis, not legal advice. For Belgium, France, Italy, the Netherlands and Meta's policy we read the full text of the official page; for the other countries the rule is taken from the official page's summary, and later amendments may be missing. Table 8 shows which kinds of rule the Ad Library data can test.

Table 7. National rules on gambling advertising: official source and what the rule says.

Country	Official source	What the rule says
Austria	RIS: Gambling Act § 56; Finance Ministry: gambling FAQ	Federal licence and concession holders must advertise to a responsible standard. Traditional sports betting falls under provincial regulation, so the federal Gambling Act does not alone describe the rules for all betting ads.
Belgium	Gaming Commission: advertising FAQ (Royal Decree of 27 February 2023)	Since July 2023, advertising games of chance covered by the decree is prohibited except as expressly permitted. Paid social-media advertising is generally barred; the exception for operators' own accounts has conditions, including no remuneration of third parties.
Bulgaria	State Gazette 42/2024: Gambling Act, art. 10; NRA decision, 17 July 2024	The 2024 amendment restricts gambling advertising in broadcast, print, public places and electronic media. The NRA's non-binding interpretation distinguishes social networks from the electronic media covered by the ban; Meta ads therefore cannot be treated as unequivocally prohibited on that basis alone.
Croatia	Narodne novine 72/2025: arts. 14, 19(7), 21	The amendment entered into force on 1 May 2025. Operators had until 1 January 2026 to comply with specified rules, including the 06:00–23:00 internet/audiovisual advertising ban and celebrity restriction, subject to exceptions. The separate content-creator prohibition is not included in that transition.
Cyprus	National Betting Authority: Code of Advertising, sections 8 and 15	Betting ads may not promise sure winnings or present betting as a solution to financial trouble. People who are, or appear to be, under 25 may not be shown betting or playing a significant role.
Czechia	Government legislative portal: Advertising Act 40/1995, § 5j	Gambling ads may not present gambling as a substitute for income or target under-18s. They must state the under-18 prohibition and carry an addiction warning. These requirements also apply to online advertising.
Denmark	Danish Gambling Authority: marketing requirements	Marketing must concern authorised gambling, present winning chances accurately and not target under-18s or frame gambling as a financial solution. Bonus offers are subject to disclosure and other conditions.
Estonia	Riigi Teataja: Advertising Act § 29^2	Advertising requires an operating permit and a warning, and may not encourage gambling or suggest social success. Placement restrictions differ for games of chance, betting, games of skill and lotteries; games-of-chance ads are limited to specified channels, with exceptions.
Finland	Police: marketing of gambling (Lotteries Act 1047/2001)	Until 30 June 2027 Veikkaus has the exclusive right to offer and market gambling in mainland Finland, subject to marketing restrictions. Other operators may not market gambling there; obtaining a licence for the future market does not permit earlier marketing.
France	ANJ (gambling regulator): illegal offer	Online sports betting, horse-race betting and poker may be offered by licensed operators; FDJ's monopoly lottery offer is also legal. Online casino sites are illegal, and advertising an unauthorised gambling site is an offence.
Germany	GGL: explanation of GlüStV 2021, § 5	No advertising for virtual slot machines, online poker or online casino on the internet between 06:00 and 21:00, and none for unlicensed gambling. The product-specific watershed is not a blanket restriction on sports-betting ads.
Greece	Hellenic Gaming Commission: Commercial Communication Regulation, sections 8–9 (FEK B 3260/2020)	Commercial communications require an approved plan and may not target minors or frame gambling as an investment. Celebrity use and promotion of bonuses tied to the cost of play are restricted, with qualifications for sponsorship, corporate social responsibility and loyalty programmes; this is not a blanket ban on awarding bonuses.

Country	Official source	What the rule says
Hungary	SZTFH: conditions for gambling ads	Advertising gambling without the required Hungarian authorisation is prohibited, as is advertising aimed at minors. An advertising intermediary is not necessarily the operator that must hold the licence.
Ireland	Irish Statute Book: commencement status; Gambling Regulation Act 2024, sections 146, 149	The 2024 Act provides for social-media opt-in restrictions, subject to commencement; these provisions are not treated as in force in our April 2025–March 2026 window. Its watershed concerns broadcasting and specified on-demand services, not a general social-media curfew.
Italy	Normattiva: Decree-Law 87/2018, art. 9; AGCOM guidelines, Delibera 132/19/CONS	Direct and indirect advertising of games and bets with cash prizes is broadly prohibited, including on social media. AGCOM's guidelines distinguish prohibited promotion from specified informational and other excluded communications; a gambling-related brand or image alone does not settle that classification.
Latvia	likumi.lv: Law on Gambling and Lotteries, sections 41(5), 50	Gambling advertising is prohibited outside gambling venues. Licensed gambling websites count as venues, so this is not a blanket ban on own-site information. The gambling restriction should not be extended automatically to the separately regulated lottery category.
Lithuania	Gaming Authority: July 2021 changes; July 2025 advertising requirements	Participation inducements, including bonuses, have been prohibited since July 2021. From 1 July 2025 to 1 January 2028, the online betting-ad exception permits limited name/logo/betting information in static form, without links to betting-company sites and subject to advertising-space limits. This restricts advertising, not casino licensing.
Luxembourg	Legilux: Law of 20 April 1977; Parliament: bill 8679, account of existing law	The 1977 Act restricts gambling establishments and promotion of unauthorised establishments. Authorised casinos and the separate National Lottery regime must be distinguished; the Act alone does not establish that all online gambling is illegal. The cited bill's proposed changes are not treated as enacted law.
Malta	MGA: Commercial Communications Regulations (S.L. 583.09); MGA FAQ	Licensees' ads must be socially responsible and must not target or appeal to minors. Required information includes the licensee's name and licence number, minimum age and responsible-gaming information.
Netherlands	Staatsblad 2023, 120: art. 2ab(2)	Since July 2023, online-gambling ads are permitted only subject to conditions including opt-outs, measures to exclude protected groups and evidence that at least 95% of people reached are aged 24 or over. Some under-24 reach does not alone establish breach; meeting the numerical threshold does not establish full compliance.
Poland	Sejm: Gambling Act, arts. 5, 29, 29b (2020 consolidation)	Public advertising of cylindrical/roulette, card, dice and slot games is prohibited, subject to statutory exceptions; licensed betting may be advertised under conditions. This is an advertising restriction, not a ban on every non-betting product or on all lawful online-casino provision.
Portugal	Diário da República: Decree-Law 66/2015, art. 3, amending Advertising Code art. 21	Advertising must be socially responsible and may not be aimed at or use minors.
Romania	Official legislative portal: Emergency Ordinance 77/2009, art. 1^2	Advertising gambling not authorised in Romania is prohibited; permitted advertising must display the licence's series and number. This is a summary of the licensing requirement, not an exhaustive account of advertising restrictions.
Slovakia	Gambling regulator: FAQ on Act 30/2019, § 4(3)–(4)	Only authorised games may be promoted, in line with their licence conditions. Affiliate promotion of lawful gambling does not itself require the affiliate to hold an operator licence.
Slovenia	Official Gazette: Gaming Act (ZIS), art. 6	Advertising games of chance without the required government concession is prohibited; the relevant authorisation concerns the gambling offered, not necessarily the advertising intermediary.
Spain	BOE: Royal Decree 958/2020, art. 26 and annulment annotations	Advertising must concern authorised gambling. The Supreme Court annulled the social-network restrictions in art. 26(2)–(3) in 2024; other provisions, including art. 26(1), remain. The annulment did not remove all social-media advertising requirements.
Sweden	Riksdag: Gambling Act (2018:1138), ch. 15	Gambling marketing must be moderate and may not target under-18s. Games requiring a Swedish licence may not be marketed without that licence; statutory licence exemptions must be distinguished.

Country	Official source	What the rule says
United Kingdom	CAP Code section 16; IGRG code, para. 55; BGC: platform exceptions	CAP bars targeting under-18s and generally bars under-25s from gambling or playing a significant role in ads. The industry 25+ prospecting rule has verified-customer and approved-platform exceptions. BGC currently lists Facebook and Instagram for 18+ targeting; their historical approval dates were not established for this study.
Meta (platform)	Meta Transparency Centre: online gambling and games	Policy requires authorisation for covered online-gambling ads, territorial lawfulness and no targeting of under-18s; advertisers remain responsible for local-law compliance [48]. Delivery alone does not prove formal authorisation.

Table 8. National gambling advertising regulations and empirical Ad Library delivery indicators (April 2025 to March 2026). Exposures represent total reach summed across delivered ads (account–ad pairs), real-money ads exclude social casinos, and "elsewhere" denotes the comparative baseline across other EU member states. Indicators reflect delivery patterns inconsistent with regulatory frameworks, not formal legal determinations.

Rule	Country	Indicator	Observed Delivery
Advertising ban	Italy	real-money ads delivered (exposures)	23,962 (76.2 M)
	Belgium		22,360 (7.9 M)
	Bulgaria		2,682 (6.0 M)
	Latvia		2,227 (1.4 M)
State monopoly	Finland	real-money ads not by Veikkaus (share of exposures)	11,352 (97.1%)
No online casino	France	exposures from online-casino-type ads (elsewhere)	22.5% (40.2%)
Casino-type ads banned	Poland	exposures from non-betting real-money ads (elsewhere)	39.8% (48.7%)
Only betting ads online	Lithuania	exposures from non-betting real-money ads	90.6%
No bonus promotion	Lithuania	real-money exposures with bonus language (elsewhere)	8.2% (33.6%)
	Greece		2.3% (36.0%)
95% reach aged 24+	Netherlands	exposures from ads with 24+ age floor	80.3%
		estimated ads exceeding 5% under-24 limit	24.8%
Licensed operators only	UK	exposures from registered operators	85.1%
25+ industry code	UK	registered operators' exposures with floor below 25	19.4%
No GamStop evasion targeting	UK	ads explicitly stating "not on GamStop" (pages); starting after regulator warning of 19 Jan 2026	14 (5); 13

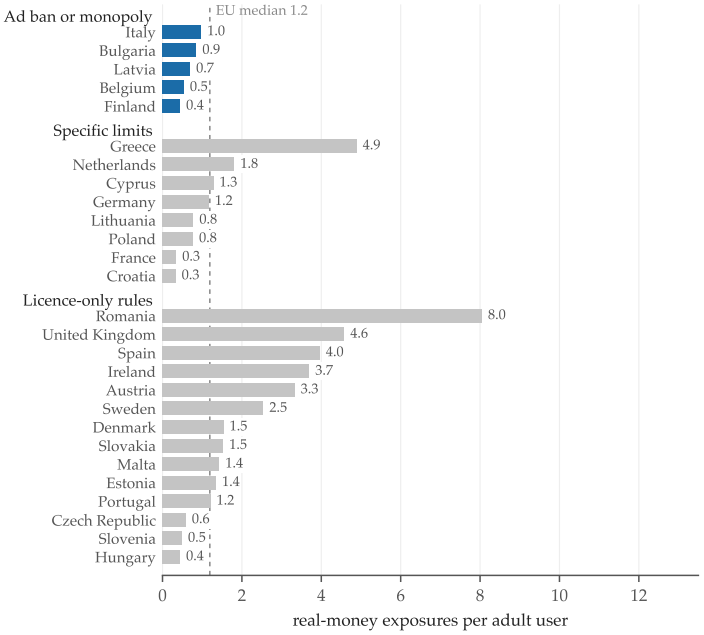

Fig. 6. Delivered real-money exposures per adult Facebook and Instagram user across countries with at least one million exposures, grouped by primary regulatory framework. The dashed line marks the EU median.

C Method details

C.1 Hierarchy of ad pairs

To evaluate how audience divergence changes as ads share fewer operational attributes, we construct a hierarchy of ad pairs sharing varying degrees of operational context (Table 9). Each pair is evaluated within delivery countries under specified disclosed targeting parameters, calculating the absolute difference in delivered male share. Two ads are considered concurrent if their active delivery windows overlap. Pair differences are averaged first within advertiser networks and then across networks to prevent high-volume operators from dominating the estimates.

Table 9. Hierarchy of ad pairs, from most to least context shared. Each pair is evaluated within delivery countries under specified disclosed targeting parameters.

	Pair Context	Potential Sources of Divergence
1	Same ad, different countries (EU only)	Country-level user base
2	Same advertiser, creative, country, and targeting (concurrent)	Ad-set settings, algorithmic delivery
3	As above, non-overlapping delivery windows	Temporal audience shifts
4	Same advertiser, creative, and country; differing targeting	Disclosed targeting parameters
5	Same advertiser, country, and targeting; differing creative	Visual creative content
6	Same creative, country, and targeting; differing advertiser	Advertiser account and optimization history
7	Different advertiser and creative; same country and targeting	Full unconstrained variation

C.2 Creatives

PDQ gives each picture a 256-bit code. Two images are the same creative if their codes differ in at most two bits, about 99% similar. Two videos are the same creative if at least 95% of their frames match within eight bits. Creatives are formed within each region. About 96% of ads belong to a creative. The rest have no downloaded file or too little structure for a reliable hash. Accepting near-identical matches rather than exact copies merges about 3% of creatives and does not change the results.

C.3 Page names

Each page name is folded to a standard form. We apply Unicode NFKD normalisation to remove accents and styled letters, lower the case, replace Cyrillic and Greek letters that look like Latin ones by their Latin twin from a hand-made table, and turn punctuation into spaces. The similarity of two folded names is one minus their Levenshtein edit distance [41] divided by the length of the longer name. Within each region we cluster the names by average-linkage hierarchical clustering cut at a similarity of 0.8. Names in a cluster are then on average at least 80% alike, as with “Big Win” and “Big Wins”. Unlike single linkage, average linkage does not chain unlike names into one cluster.

Name similarity alone cannot tell a licensed operator with one page per country from many unrelated pages with a popular name. We therefore sort the clusters into five kinds with information that does not depend on the names, applying the rules in the order of Table 10. Operator status comes from a hand-curated list of 118 official pages of gambling operators and from matching advertiser networks against official licensing registers (Section 3).

C.4 Embeddings and design families

Images are embedded as one picture with the longer side at most 1,024 pixels. Videos are embedded as up to 32 frames sampled once per second, with the longer side at most 640 pixels, passed to the model together. Vectors are

Table 10. Kinds of name clusters, assigned in this order.

Kind	Rule
Single page	the cluster has one page
Operator family	a page of the cluster is on the curated list of official operator pages, or belongs to an advertiser network typed as a registered operator
Main page with copies	one page holds at least 80% of the cluster's reach
Page farm	five or more pages
Small group	two to four pages

768-dimensional and normalised to unit length, so their dot product is the cosine similarity. Copies of a creative have nearly the same vector, and a creative copied into a thousand ads would otherwise dominate any clustering. We therefore work with one vector per creative, the normalised mean of its ads' vectors. An ad without a creative group counts as a creative of its own. This gives 29,200 creatives in the EU and 10,290 in the UK. For the similarity pairs we keep at most 60 ads per group of network, setting, main country and format.

Images and videos sit slightly apart in the space whatever they show. For clustering we therefore subtract the mean vector of each format and add back the overall mean. The prediction models instead include the format as a variable. We cluster the format-centred vectors with k-means ($k = 30$, ten random starts) [54], and each ad takes the cluster of its creative. The creatives form a continuum with dense regions rather than separate groups. The silhouette stays between 0.05 and 0.13 for every k from 5 to 100, so k sets the resolution rather than recovering a true number of groups. We name each cluster as a design family by inspecting its most central creatives.

C.5 Prediction

Ridge penalties. All five models are ridge regressions with one fixed penalty, $\lambda = 1$, on every column of the disclosed context, on the 28 annotation labels and on the 42 content scores. The penalty is not selected by cross-validation. The context enters as 0/1 indicators (targeting setting, main country, start month, platforms, destination, language and format; values with fewer than 30 ads form the reference) and as logarithms (number of countries, duration and reach). The labels are 0/1 indicators. The content scores are standardised within region and format, so a score has mean 0 and standard deviation 1 and one penalty is comparable across the 42. The 768 dimensions of the embedding have unit length overall and are multiplied by $\sqrt{1/10}$, which makes a penalty of 1 act like 10 on them.

Advertiser fixed effects. In the models with the advertiser network (B to E), the training fold's network mean of the male share and of every column is subtracted before fitting, and a test ad's prediction is its network's training mean plus the model's prediction from the ad's deviations. Ads of networks that do not appear in the training folds cannot be predicted, so all five models are scored on the same ads: those whose network is in training.

We check the result in three ways. Folds are grouped by *template*, a cluster of about ten similar creatives, so that no close variant of a test creative is in training. Templates are k-means clusters of the creative vectors, with k a tenth of the number of creatives. The delivery-size variables are dropped. Images and videos are fitted separately. A last design asks whether the creative predicts the audience for advertisers the model has never seen. Its folds are grouped by network and the advertiser is left out of the model. We also score it only on test ads whose creative and brand (the operator named in the annotation) never occur in training.

D Additional figures

This appendix holds the figures that the results refer to but that are not needed to follow the main argument. Unless a caption says otherwise, each figure has an EU panel (left, blue) and a UK panel (right, orange), shares are of adult reach, and intervals are 95% intervals from resampling advertiser networks.

D.1 Targeting and delivery

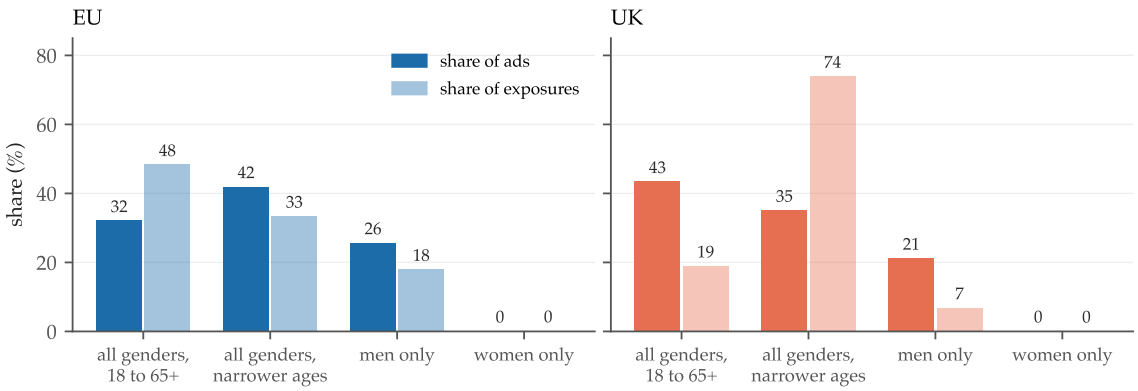

Fig. 7. Targeting distribution of gambling ads, showing the proportion of campaigns (dark bars) and proportion of delivered exposures (light bars) across four targeting configurations.

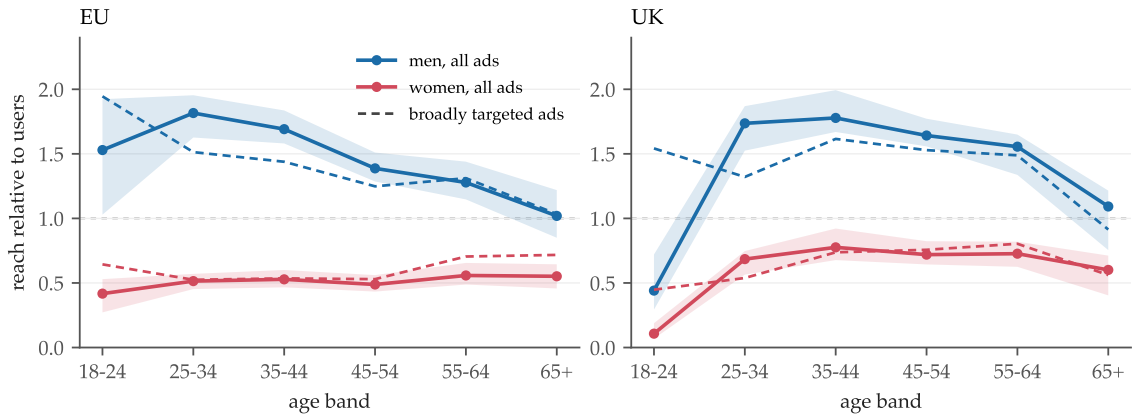

Fig. 8. Delivered reach relative to platform users by sex and age cohort. Values reflect the share of exposures in each demographic group divided by expected reach if delivery followed adult Facebook and Instagram user distributions. Solid lines show all gambling ads with 95% intervals from resampling advertiser networks, and dashed lines show broadly targeted ads. A value of 1.0 indicates demographic parity.

D.2 Advertisers and targeting contrasts

D.3 Creative content

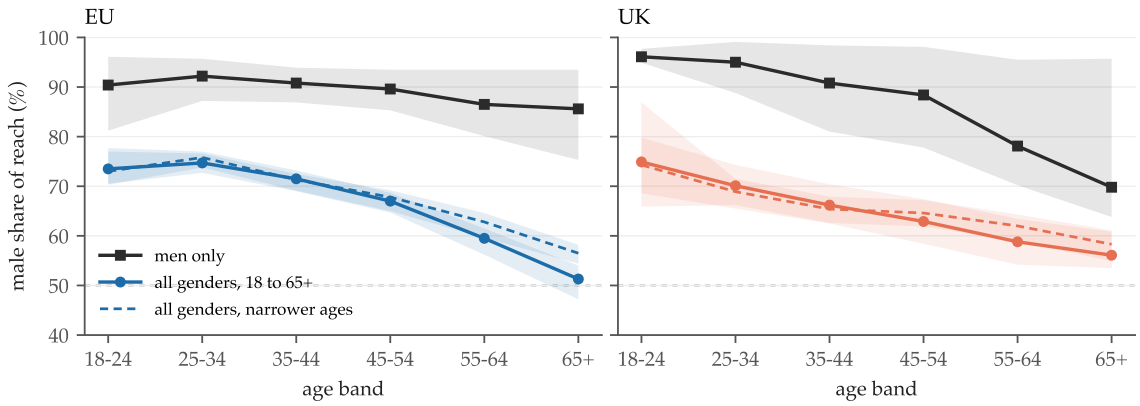

Fig. 9. Delivered male share within each age cohort across targeting groups, pooled across ads with 95% intervals from resampling advertiser networks. The dashed horizontal line marks 50% (gender parity). Broad and narrower-age ads are open to both genders.

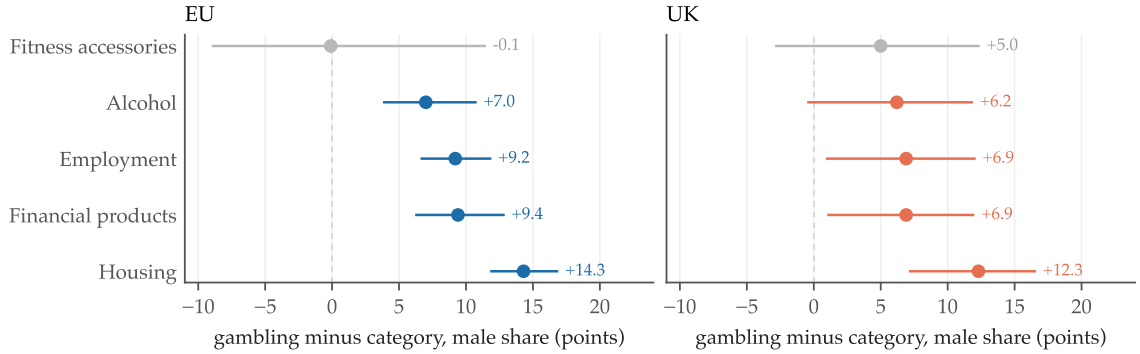

Fig. 10. Difference in male share between broadly targeted gambling ads and each comparison category (percentage points) for ads delivered between 15 October 2025 and 31 March 2026. Positive values indicate a higher male share in gambling. Intervals are 95% percentile intervals from separately resampling advertiser networks on each side.

Table 11. The 42 text descriptions behind the content scores (Appendix C.5), by family. An ad's score is the cosine similarity between its embedding and the embedding of the description, standardised within region and format.

Family	Descriptions
Sport	a football (soccer) match with players on the pitch; a professional footballer in a club shirt; a basketball game; a tennis player on a court; horse racing with jockeys; a boxing or MMA fight; a Formula 1 racing car; esports and video game players
Games	slot machine reels with fruit and jackpot symbols; a roulette wheel in a casino; playing cards and poker chips on a table; dice being rolled; lottery balls and a lottery ticket; a scratch card; a crash game with a small aeroplane flying up a curve; a cartoon chicken crossing a road game; colourful cartoon game art with fantasy characters
People	a man; a woman; a young man in his twenties; an elderly person; a woman in revealing clothes posing; a famous celebrity; a group of friends cheering together; a person talking to the camera
Money	stacks of cash banknotes; a pile of gold coins and treasure; a luxury lifestyle with expensive cars and a villa; a big jackpot win with large numbers; bitcoin and cryptocurrency
Format and look	a smartphone screenshot of a betting app; a table of betting odds; a screenshot of chat messages; a betting slip with selected matches; football statistics and prediction charts; a robot or artificial intelligence; a poster with large bold text; a company logo on a plain background; a realistic photograph; a cartoon illustration; neon lights at night; a TV news report

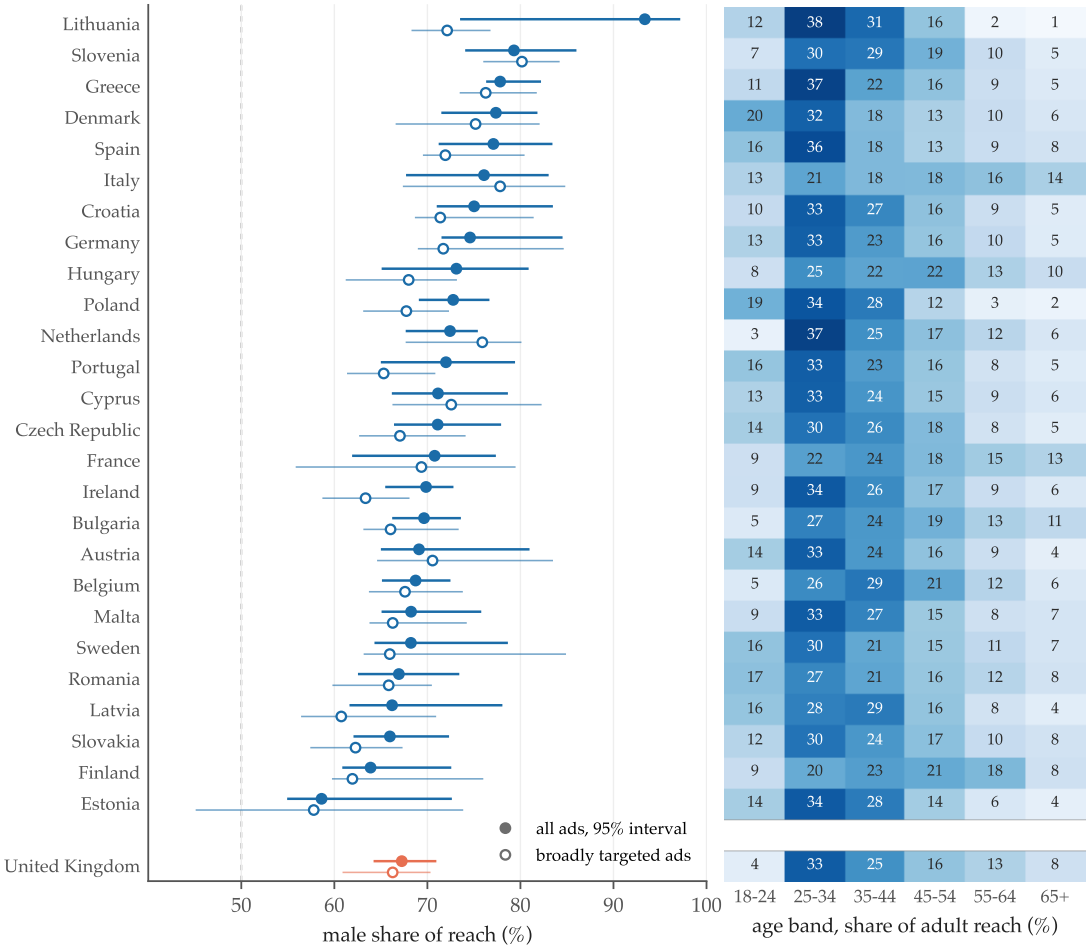

Fig. 11. Delivered audience by country for all markets with at least one million exposures. Left panel shows male share of reach for all gambling ads (filled circles) and broad ads (hollow circles). Right panel shows the age composition of delivered reach across adult age bands. EU member states are sorted by male share, with the UK shown at the bottom.

Table 12. Creative characteristics associated with delivered male share: percentage point change in male share per standard deviation of content score, estimated within advertiser networks and disclosed context. Top three positive and negative associations ordered by pooled mean; 95% intervals resample advertiser networks.

Characteristic	Family	EU	UK
football statistics and prediction charts	format and look	+0.8 [−0.2, 1.6]	+2.6 [1.1, 4.0]
a professional footballer in a club shirt	sport	+1.3 [0.7, 1.8]	+1.7 [0.5, 3.0]
esports and video game players	sport	+1.6 [1.0, 2.4]	+0.9 [−0.0, 2.0]
a pile of gold coins and treasure	money	−1.7 [−2.7, −0.8]	−2.9 [−4.0, −1.7]
slot machine reels	games	−1.1 [−2.3, −0.0]	−3.2 [−4.9, −1.5]
colourful cartoon game art	games	−0.7 [−1.3, −0.1]	−2.8 [−4.3, −1.5]

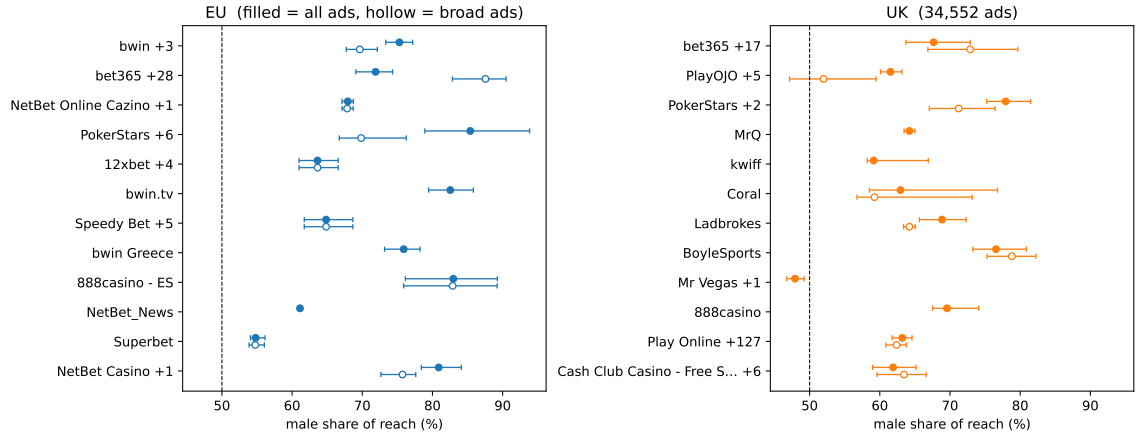

Fig. 12. Delivered male share for the twelve largest advertiser networks by total exposure, showing all ads (filled circles) and broad ads only (hollow circles). Networks are labelled by their primary page name. Whiskers show 95% intervals from resampling ads within each network.

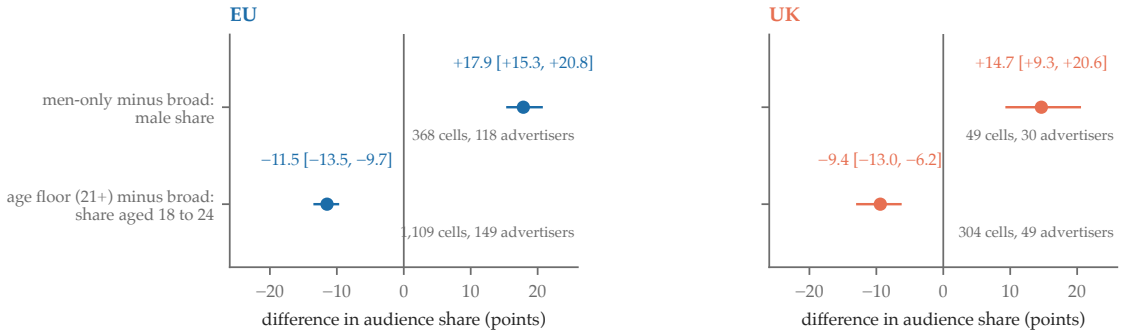

Fig. 13. Within-advertiser targeting contrasts holding creative content and delivery country constant. Left panel shows the difference in male share between men-only ads and broad ads. Right panel shows the difference in the share of 18–24-year-olds between ads setting an age floor of 21 or higher and broad ads. Values are averaged across advertisers, with 95% intervals from resampling advertisers.

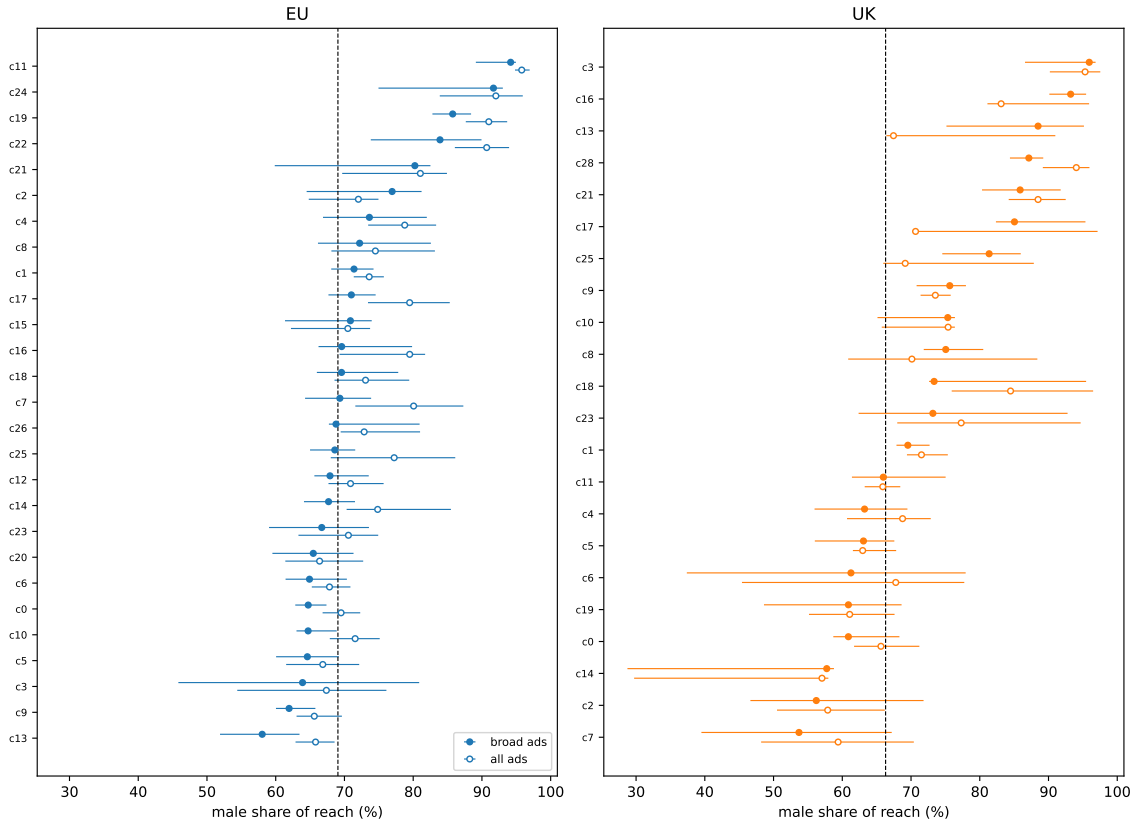

Fig. 14. Male share of reach across creative design clusters for broad ads (filled circles) and all ads (hollow circles). The sample includes clusters with at least 30 broad ads from five or more advertiser networks. The dashed horizontal line indicates the male share across all broad ads.

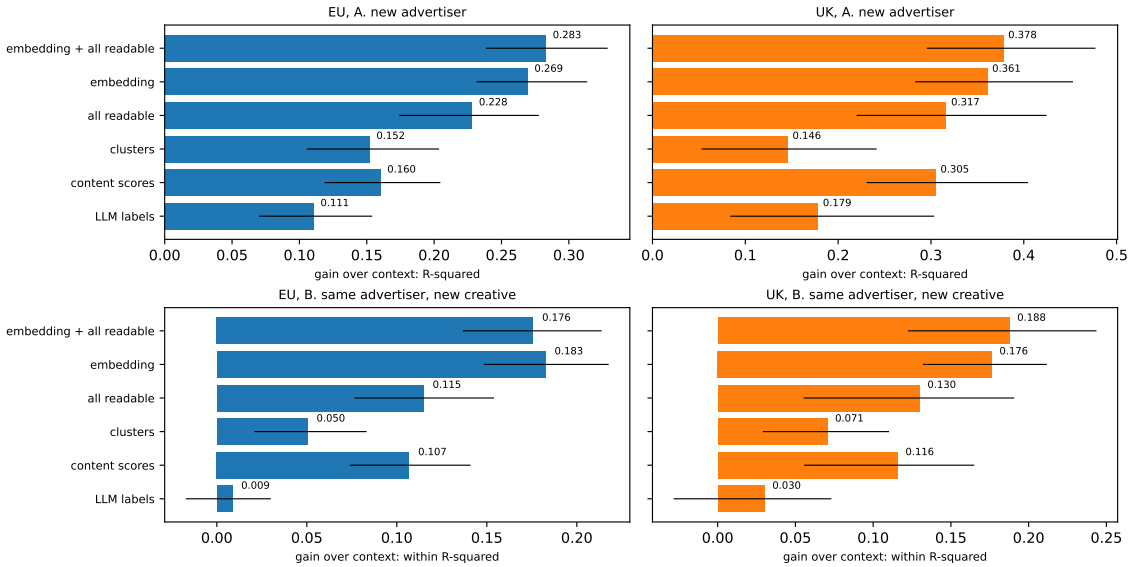

Fig. 15. Incremental gain in out-of-sample R^2 over observed campaign context when describing creatives through multimodal embeddings or interpretable feature sets (28 annotation labels, 42 zero-shot content scores, and 30 design clusters). Top panels show generalization to unseen advertisers, and bottom panels show within-advertiser prediction holding out creative variants.

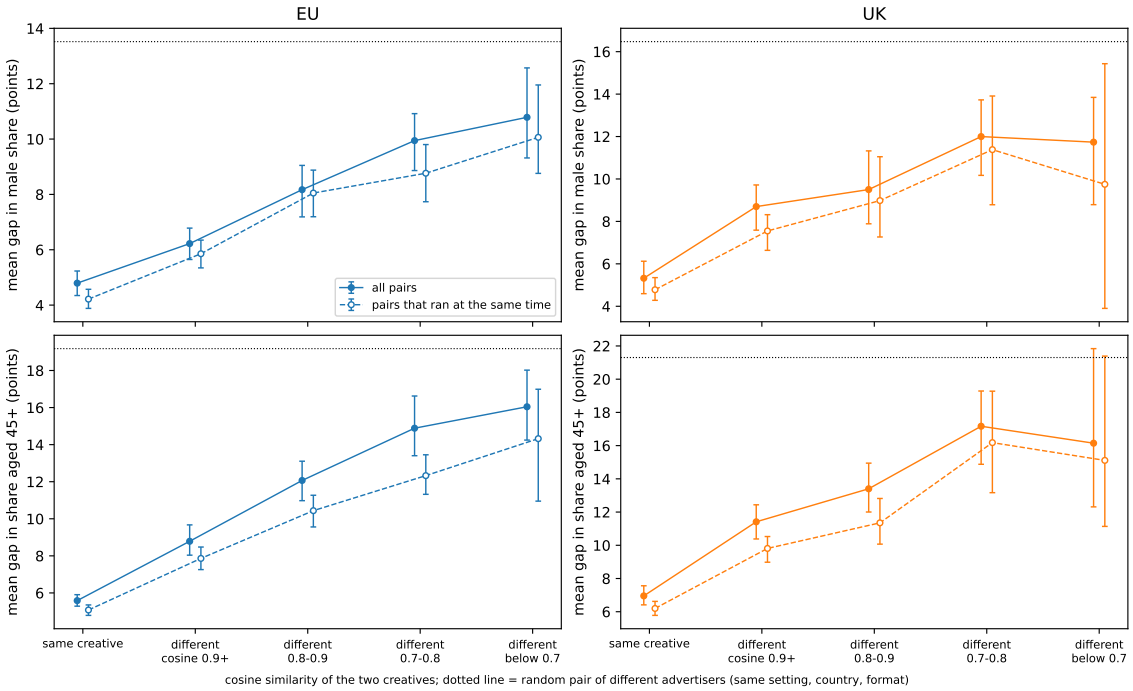

Fig. 16. Audience divergence between pairs of ads from the same advertiser network with identical targeting, country, and media format, plotted against the cosine similarity of their multimodal creative embeddings. Top panels show the mean absolute difference in male share, and bottom panels show the difference in the share aged 45 and over. Filled points show all pairs, and hollow points show pairs with overlapping delivery periods. The dotted line marks the reference divergence between random pairs of unrelated advertisers. Intervals are 95% intervals from resampling advertiser networks.

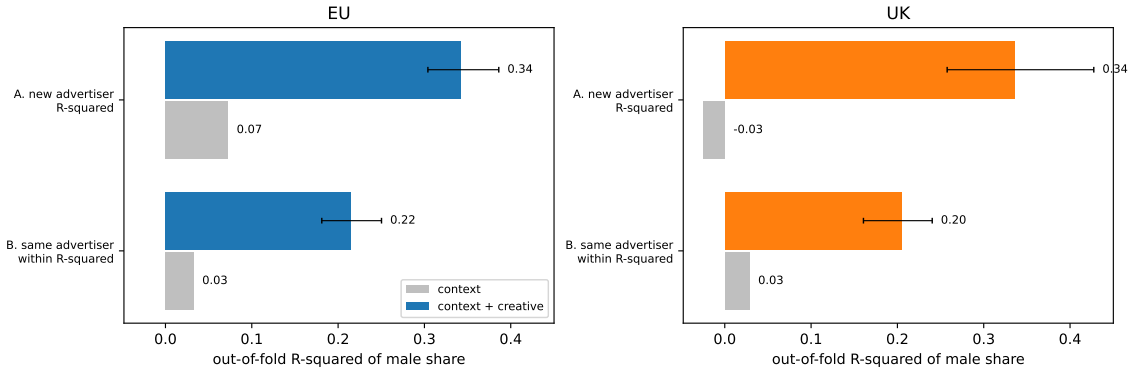

Fig. 17. Out-of-sample R^2 in predicting male share from observed campaign context alone (grey bars) and with multimodal creative embeddings added (coloured bars). Panel A evaluates generalization to unseen advertiser networks by grouping cross-validation folds by network. Panel B evaluates within-advertiser prediction using folds grouped by creative with advertiser fixed effects. Evaluated on ads open to all genders with at least 100 accounts. Whiskers show 95% intervals for the step gain from resampling advertiser networks.