

Facebook algorithm's active role in climate advertisement delivery

Received: 2 April 2024

Accepted: 26 March 2025

Published online: 13 June 2025

 Check for updates

Aruna Sankaranarayanan¹✉, Piotr Sapiezynski² & Una-May O'Reilly¹

Climate advertising on social media can shape attitudes towards climate change. Delivery algorithms, as key actors in the climate communication ecosystem, determine ad audience selection and might introduce demographic bias. Here, we present a two-part study—an observational analysis ($n = 253,125$) and a field experiment ($M = 650$)—to investigate algorithmic bias in Facebook's climate ad dissemination. Our findings provide preliminary evidence that the algorithm's selection of ad audiences can be explained by factors such as ad content, audience location (US states), gender and age group. Moreover, the cost-effectiveness of contrarian ads is linked with the conservative political alignment of a state, while the cost-effectiveness of advocacy ads correlates with liberal political alignment, higher population and per-capita gross domestic product; ad targeting strategies further modulate these effects. The skew in the distribution of climate ads across US states, age groups and genders reinforces existing climate attitudes.

Accurate climate communication is essential for large-scale climate action. However, despite a strong scientific consensus on the issue^{1–9}, lobbying the general public towards climate action has been challenging; decades of organic and paid climate contrarian communication has fuelled mistrust and cast doubt on the urgency of addressing climate change. In recognition of this issue, the Intergovernmental Panel on Climate Change identified 'rhetoric, misinformation and politicization of science' as key barriers to climate action. To counter and understand the influence of contrarian communications, scholars across disciplines have documented their deceptive nature^{10–24} and identified key players who disseminate and oppose them. However, these studies have primarily focused on print and broadcast media^{25–32} and traditional climate actors, such as fossil fuel and climate advocacy organizations, public relations firms and think tanks^{33–37}.

Social media platforms present challenges to accurate climate communication because algorithms for content and ad delivery, alongside contrarian and advocacy actors, heavily influence audience selection and pricing. These algorithms require investigation in the climate context because they have exhibited gender, racial^{38–41} and political biases⁴² in ad audience selection, even when advertisers fully relinquish control. These biases probably arise from algorithms optimizing for

ad 'relevance' in auction-based ecosystems, favouring users deemed most relevant rather than distributing ads randomly or purely based on bids⁴³. This can shape climate discourse—Goldberg et al.⁴⁴ found that climate advocacy ads delivered to Republicans improved their understanding of climate change by several percentage points.

The Meta Ad Library, a transparency effort by Facebook, allows researchers to investigate climate ads by asking questions like: Who funds climate content? How much are they spending? Where are they targeting campaigns? and To whom are ads delivered according to age and gender?—with greater precision than print or broadcast media. This dataset helps researchers understand the algorithm's role in disseminating climate contrarian and advocacy content. Using observational data from the Meta Ad Library and a field experiment, we examined whether Facebook's ad delivery system influences climate ad dissemination and if demographic biases emerge. We argue that advertising algorithms are key actors in the climate communication ecosystem.

The study was organized into two parts. In part 1, we analysed algorithmic bias in Facebook's advertising data from 2018 to 2023. Using a peer-reviewed list^{33–37} of climate contrarian and advocacy advertisers, we identified 63,542 contrarian and 139,012 advocacy ads after filtering for duplicates and relevance. For each ad, Facebook reports only

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA. ✉e-mail: arunas@mit.edu

Table 1 | US ads dataset summary

	Contrarian ads	Advocacy ads
Advertiser counts	245	453
Expenditure (US\$)	34–47 million	45–67 million
Impressions	1.9–2.4 billion	3.4–4.3 billion

delivery statistics, omitting the targeting criteria set by advertisers. This obscures advertiser intent. However, when ad delivery is isolated to a single state, gender or age group, we can infer an intention to target. Conversely, when ads reach all US states or demographics, we can assume no specific state-based or demographic-based targeting, although other interest-based or customized targeting methods may be in use. Under the constraints of these limitations, our analysis indicated evidence of algorithmic skew. For more details, see the section titled ‘Observational study of algorithmic bias in Facebook ad data’.

Detecting algorithmic bias, without such qualification, requires full transparency of both the advertiser’s targeting intent and the algorithm’s delivery decisions. To address this, in part 2, we investigated and found evidence for algorithmic bias in a field experiment. In this study, we tasked the platform with the impartial delivery of ads featuring images of oil drills and solar cells to all US states, genders and age groups using an ecologically valid experimental design that aligns with previous audits of the Facebook’s ad algorithm³⁸. More details can be found in the section titled ‘Field study of bias in the Facebook ad algorithm’. We tested, and partially rejected, the hypothesis that the delivery will be unbiased across gender, location (US states) and age groups.

Observational study of algorithmic bias in Facebook ad data

In our observational study, we investigated if the algorithm provided any advertising advantage to climate contrarian and advocacy ads using data from 63,542 contrarian and 139,012 advocacy ads collected between May 2018 and May 2023. We introduced a metric of ad impressions per dollar (AIPD), which refers to the total non-unique number of times an ad was shown on a Facebook user’s feed (including multiple displays to the same user) per dollar spent. A higher AIPD indicates a better ad price. While Facebook provides estimated audience sizes at the time of ad launch, the actual ad audience and pricing are dynamically determined by the algorithm. We compared AIPD values across gender, age and location groups to assess cost-effectiveness. Statistical comparisons of AIPD revealed whether it was cheaper to deliver contrarian or advocacy ads to specific US states, genders or age groups—the higher the AIPD for a demographic, the cheaper it was to deliver to that audience group. See Table 1 for the summary statistics from our dataset.

Among US states, we observed a ‘hollowing-out’ effect: contrarian ads were more cost-effective in interior states, while advocacy ads were more cost-effective on the East and West coasts (Fig. 1 and Supplementary Tables 1–3). Contrarian ads have a higher AIPD value across nearly all US states, indicating more views per dollar spent (Fig. 1a and Supplementary Table 1). However, advocacy ads perform better when targeted at a single state (Fig. 1b and Supplementary Table 2), with contrarian ads securing better pricing in only eight states.

For ads not targeted according to location or demographic (Fig. 1c and Supplementary Table 3), advocacy ads have higher AIPD on the East and West Coasts, while contrarian ads dominate the Midwestern, Southeastern and Southwestern USA. With targeted ads, the algorithm follows advertiser intent; however, in non-targeted ads, it has a larger role in decision-making (contrarian ads enjoy a price advantage in most US states).

Similarly, exposing males to contrarian ads and females to advocacy ads was cheaper. For males, contrarian ads had a higher mean AIPD (4.35) than advocacy ads (3.53) ($P < 0.001$). For females, advocacy

ads had a higher mean AIPD (5.4) than contrarian ads (3.51) ($P < 0.001$). See Fig. 2a and Supplementary Table 4 for more details.

For non-targeted ads, where algorithmic decision-making has a larger role, the mean AIPD for males was significantly higher for contrarian (4.85) than advocacy (3.11) ads ($P < 0.001$). For females, advocacy ads had a significantly higher mean AIPD (5.02) than contrarian ads (3.76) ($P < 0.001$) (see Fig. 2a and Supplementary Table 5 for more details). As advocacy ads are not exclusively targeted according to gender group, we did not compare gender-targeted ad delivery.

AIPD values according to age group showed a ‘hollowing-out’ effect like ad delivery trends across US states. Advocacy ads are cheaper for ages 13–34 and 65+, while contrarian ads cost less for ages 35–65 (Fig. 2b and Supplementary Table 6). Independent *t*-tests showed that advocacy ads have significantly higher mean AIPD values in the 13–34 and 65+ groups, while contrarian ads dominate in the 35–64 range ($P < 0.001$). This pattern holds for non-targeted ads, with advocacy ads being significantly cheaper for younger audiences (13–17) and contrarian ads more cost-effective for ages 35–64 (Fig. 2b and Supplementary Table 7). The *t*-tests showed that advocacy ads have significantly higher AIPD values in the 13–17 age group ($P < 0.001$), while contrarian ads dominate the 35–64 age range ($P < 0.001$). Differences in other age groups were not statistically significant. As advocacy ads are not strictly age-targeted, age-targeted comparisons are not made.

Regression analysis

We used robust linear regression to assess state-level cost differences based on population (millions), per-capita gross domestic product (GDP, in millions of dollars), oil production (millions of barrels) and 2020 Republican lean, with $\log(\text{AIPD})$ as the dependent variable. The model included fixed effects for the year and state categories. (See the section on ‘Regression analysis methods’ in the Methods for our regression methodology and Supplementary Tables 8 and 9 for results.) For non-targeted ads, we found that Republican-leaning states saw an 18.57% rise in AIPD for contrarian ads and a 58% drop for advocacy ads ($P < 0.001$). A 1-unit increase in per-capita GDP reduced AIPD by 42% for contrarian ads ($P < 0.001$) and 5.9% for advocacy ads ($P < 0.05$), while a 1-unit population increase lowered AIPD by 6.6% for contrarian ads ($P < 0.001$) and 0.55% for advocacy ads ($P < 0.05$). A 1-unit rise in oil production decreased contrarian AIPD by 96% ($P < 0.001$).

For explicitly targeted ads, advocacy ads face a higher cost disadvantage across all state factors except political alignment—Republican states raised advocacy AIPD by 87% ($P < 0.001$) and lowered contrarian AIPD by 20%. A 1-unit population rise cut AIPD by 9.4% for contrarian ads and 28% for advocacy ads ($P < 0.001$). A 1-unit GDP increase reduced advocacy AIPD by 95.1% ($P < 0.001$) but had no significant effect on contrarian ads. A 1-unit oil production rise increased contrarian AIPD by 31% ($P < 0.05$) and decreased advocacy AIPD by 71% ($P < 0.001$).

The observational study revealed skewing in the distribution of contrarian and advocacy ads across state, age and gender groups, particularly in non-targeted ads where algorithmic delivery varied according to ad type. In US states, this skew is partly explained by political alignment. However, these findings are limited by the proxies used for ad targeting.

Field study of bias in the Facebook ad algorithm

To directly detect algorithmic skewing, we conducted a field experiment controlling all targeting parameters. We tasked the algorithm with impartially delivering 650 ads featuring solar cells, oil drills or neutral controls (Fig. 3), with logos from either a contrarian organization (ExxonMobil), an advocacy group (Greenpeace) or no logo, to US audiences. We then compared ad delivery across states, genders and age groups.

We used the ad delivery ratio (*D*)—the ratio of observed ad views from our experiment to expected views based on Facebook’s estimated audience sizes across US states—to measure and compare




Fig. 1 | Comparison of AIPD values across US states. a, AIPD for the entire dataset ($n = 4,217,991$), including both targeted and non-targeted ads, showing that contrarian advertisers have a general advantage. **b**, AIPD for ads targeted ($n = 82,177$) at a single US state, where climate advocates hold an advantage. **c**, AIPD for non-targeted ads ($n = 2,895,737$), that is, ads whose delivery was influenced by algorithmic decision-making. In this case, contrarian advertisers hold an advantage. States where the mean AIPD favours contrarians (C) are

highlighted in yellow, while those favouring advocates (A) are highlighted in green. States where the AIPD distributions of contrarian and advocacy ads do not differ significantly are shown in grey. Statistical significance was assessed using a two-sided Welch's t -test for groups with at least 30 samples and a two-sided Mann–Whitney U -test otherwise. To control for the false discovery rate (FDR), P values were adjusted using the Benjamini–Hochberg correction.

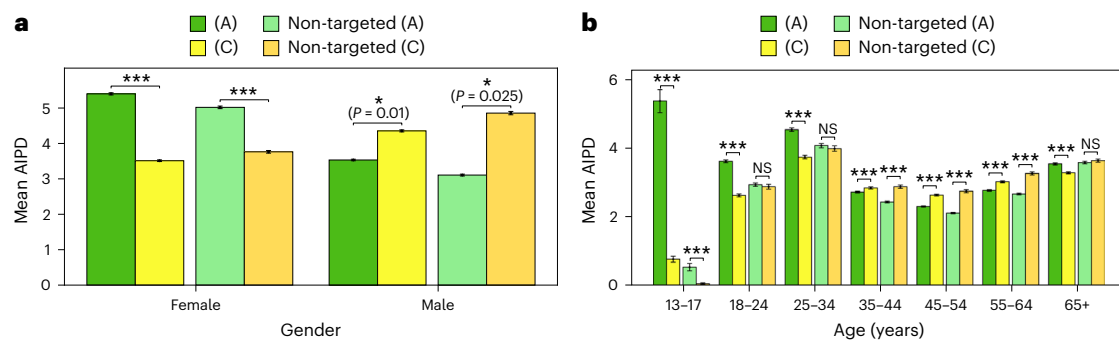


Fig. 2 | Comparisons of the mean AIPD of contrarian (C) and advocacy (A) ads according to audience gender, age and targeting strategies. a, AIPD comparisons according to gender ($n = 3,028,137$ overall; $n = 1,472,610$ for non-targeted ads; n varies according to gender subgroup). **b**, AIPD comparisons

according to age (same n values as in **a**, varying according to age subgroup). Data are presented as the mean \pm 95% confidence intervals. Statistical significance was determined using a two-sided Welch's t -test with Benjamini–Hochberg correction to control for the FDR. *** $P < 0.001$, * $P < 0.05$. NS, not significant.

delivery patterns. This allowed us to compare the observed AIPD from Facebook's algorithm with the AIPD from a uniform sampling baseline of Facebook's self-reported audience estimates, as well as report the over- or under-representation of a demographic in the ad delivery. For example, we found a significant difference in the delivery of oil drill versus solar cell ads to female audiences, with females being under-represented (median delivery ratio < 1) and males over-represented (median delivery ratio > 1) in the audience, as shown in Fig. 4b. The consistency of results between solar cell and control ads, and oil drill and control ads, in Fig. 4b,c serves as a sanity check to validate this metric. (See the 'Experimental design', 'Covariate analysis' and 'Field experiment' sections in Methods for more details on the initial covariate analysis, experimental design and methodology, respectively).

Demographic features influence ad delivery

We analysed ad delivery differences across US states, genders and age groups using a one-way Kruskal–Wallis test, followed by Dunn's test with Benjamini–Hochberg correction. For the ad delivery ratio D across US states, the Kruskal–Wallis test showed significant differences ($P < 0.001$) in oil drill, solar cell and control ads. Post-hoc analyses found significant oil drill versus solar cell differences in 25 US states (Fig. 4a and Supplementary Table 10), with the delivery ratio in 21 US states favouring oil drill ads. A regression analysis with the state category as a fixed effect (Supplementary Table 11) showed that a 1-unit increase in state oil production raised the delivery ratio of drill ads by 28.98% and solar ads by 4.89% ($P < 0.001$). A 1-unit population increase decreased both by 5% ($P < 0.001$). Republican-leaning states saw a 59% rise in the

delivery ratio of solar ads and 51% for drill ads ($P < 0.001$). A 1-unit GDP increase reduced the delivery ratio of solar ads by 81% ($P < 0.05$) with no significant effect on drill ads. For the ad delivery ratio D across genders, the Kruskal–Wallis test showed marginal significance ($P = 0.054$), indicating weak group-level differences. The post-hoc analyses found a significant difference between oil drill and solar cell ads ($P < 0.05$). We also found (Fig. 4b and Supplementary Table 12) a significant difference in the delivery ratio D between oil drills and solar cells for both males and females ($P < 0.001$). Solar cell ad delivery was higher among females and lower among males compared to oil drills. Additionally, as seen in prior work, females were under-represented, while males were over-represented in ad audiences⁴⁵. For the ad delivery ratio D across age groups, the Kruskal–Wallis test showed significant group-level differences ($P < 0.001$) in oil drill, solar cell and control ads. The post-hoc analysis found significant differences between oil drill and solar cell ads and between oil drill and control ads ($P < 0.05$). Further analysis revealed significant D differences between oil drills and solar cells across all age groups except those aged 45–54 years ($P < 0.05$) (Fig. 4c and Supplementary Table 13). The delivery ratio for solar cell ads was higher among younger audiences (18–44) and lower among older audiences (55+) compared to oil drill ads. Furthermore, younger individuals (18–54) were under-represented and older individuals (55+) were over-represented in our ad audiences.

Image features influence ad delivery

Analyses of data from our field experiment found a skew in ad delivery, similar to our observational study. We found that the ad delivery ratio (D) significantly ($P < 0.05$) varied depending on whether the ad image




Fig. 3 | Reference images resembling the ones used in the ads for our experimental study. Ads featuring the oil drill and solar image sets were also launched with the logos of ExxonMobil and Greenpeace, representing contrarian and advocacy actors, respectively (images not shown). Ads containing images without logos were launched twice to assess the consistency of Facebook's

delivery. These reference images were sourced from the Metropolitan Museum of Art (Ewer with Elephant-Headed Spout, <https://www.metmuseum.org/art/collection/search/7237>) and the US Geological Survey^{54–56}. **a**, Oil drills ($\times 65$). **b**, Solar cells ($\times 65$). **c**, Controls ($\times 65$).

featured solar cells, oil drills or controls in 25 US states, all genders and five of six age groups, indicating potential location, gender and age bias in the algorithm (Fig. 4a–c).

Ad delivery was consistent for a given ad image, as shown by duplicate ads featuring solar cells, oil drills or control images without a logo. A Fisher's test revealed consistent audience sizes for 90% of duplicate pairs across US states, 100% across genders and 99% across age groups ($P < 0.05$). This suggests that the algorithm's decision-making is considerably influenced by the ad image (see Supplementary Tables 14–16 for more details).

Effects of logo and audience sizes on ad delivery

A one-way Kruskal–Wallis test (with Dunn's test and Benjamini–Hochberg correction) showed that, contrary to our hypothesis, the ad delivery ratio D did not differ significantly ($P > 0.05$) for oil drills or solar cells featuring the logos of ExxonMobil, Greenpeace or no logo (see Supplementary Tables 17–22 for more details). We found significant effects only when considering pairwise post-hoc tests in the delivery of oil drills and solar cells with an ExxonMobil logo in Washington, Arizona, Montana, Nevada, Oklahoma and Colorado ($P < 0.05$). Therefore, we combined images with and without logos in our tests for additional statistical power.

Audience sizes in our experiment did not always align with Facebook's ad manager estimates (Supplementary Tables 23–25). A chi-squared goodness of fit test showed that control ads matched state-level estimates 64% of the time, compared to 42.5% for solar cell and oil drill ads ($P < 0.05$). For gender, solar cell and oil drill ads aligned 67% of the time, while control ads matched only 22% of the time ($P < 0.05$). Across age groups, no ad type (0%) matched the estimates ($P < 0.05$). This suggests that the algorithm weights audience characteristics differently, possibly influenced by factors such as the over-representation of older Facebook users making them cheaper. Our results also align with findings that males are cheaper to reach online than females⁴⁵ (see Supplementary Tables 26–34 for more details).

Discussion

We discovered a 'hollowing-out' effect: advocacy ads are cost-effective in West Coast and Northeastern states, while contrarian ads are cost-effective in interior states. Contrarian ads are cheaper for older audiences and males, while advocacy ads are cheaper for younger audiences and females. These patterns hold for ads not explicitly targeted according to location or demographics. When location-based targeting is applied, advocacy ads become cheaper across most US states,

indicating that climate advocacy groups benefit from targeted ads, while contrarian groups gain from algorithmic decision-making when selecting conservative audiences.

Previous research showed that males and older individuals are more likely to be dismissive or doubtful on climate change^{46–49}, while females and younger people tend to be more alarmed or cautious^{46–49}. Our analysis indicates that political leaning in US states explains the cost advantages for contrarian and advocacy ads, suggesting that ad patterns reinforce existing climate attitudes. We observed a reversal in cost patterns between non-targeted and targeted settings: contrarian ads are cheaper in conservative states when not targeted, but in liberal states when targeted. An important trend is the significant cost disadvantage to contrarian ads in oil-producing states in the 'non-targeted' setting, but this disadvantage also manifests for advocacy ads in the 'targeted' setting. Increasing state population or per-capita GDP favours a cost advantage and disadvantage for advocacy ads in the non-targeted and targeted settings. Our field study revealed that ad image content influences algorithmic delivery, with image type affecting demographic representation, while corporate logos do not significantly impact this effect. The algorithm does not sample uniformly from Facebook's audience estimates, resulting in skewed delivery based on audience and ad characteristics, with this skew being consistent rather than arbitrary. Past research suggested that delivery is largely driven by automated image classification rather than ad–audience interactions³⁸; delivery skew has been shown to occur in ads invisible to humans but detectable by classifiers³⁸.

Since each ad in our experiments was allocated an identical budget of US\$1, the over-representation of certain demographics in delivery – as indicated by higher delivery ratios – reflects differential allocation by Facebook's delivery system. Our results also connect ad exposure to psychological, cultural and political influences on climate discourse. Exposure to climate information can raise perceived scientific consensus by 16 points, especially in conservative areas⁵⁰, underscoring concerns about algorithmic bias amplifying contrarian content. The Six Americas report segments the United States into six climate response groups, each requiring tailored messaging^{46,47}. Doubtful and dismissive groups – typically older males in the interior US states – respond best to non-confrontational messaging, yet algorithms frequently serve them oil drill ads. Advocacy ads featuring oil rigs may reinforce scepticism, while contrarian ads using greenwashing (for example, solar panels) reach younger females. The impact of greenwashing on perceived climate commitment remains debated^{51,52}, but misleading climate messaging weakens pro-environment beliefs and misinformation




Fig. 4 | Results of our field experiment showing differences in the delivery of ads featuring oil drills and solar cells. a, Delivery of ads across US states ($n = 22,024$). **b,** Delivery of ads across gender groups ($n = 1,300$), **c,** Delivery of ads across age groups ($n = 3,415$). The P value significance is shown in the figure:

*** $P < 0.001$, * $P < 0.05$. A median delivery ratio of less than 1 suggested that a group was under-represented while a median delivery ratio greater than 1 suggested that a group was over-represented in our ad audiences.

detection⁵³. Inoculation strategies are largely ineffective against these biases, leaving even climate-supportive groups vulnerable to algorithmic influence⁵³.

We present preliminary evidence that ad recommendation algorithms are emerging and evolving influencers of climate discourse, although their full impact on ad delivery remains unclear. Our US\$1 budget experiments focused on demographic representation, but larger budgets may reveal spending-based algorithmic shifts. Grouping solar and oil ads in the same Facebook ad set could improve ecological validity, while broader climate-related ads may uncover further bias. We hypothesize that Facebook's algorithms shape these patterns, warranting research into biases according to ethnicity and political affiliation. Our analysis uses pointwise representations over noisy and constrained data; further investigations into the structure and dynamics of ad delivery would help validate and expand these findings. We urge climate scholars to examine online advertising, audience targeting and algorithmic bias to counter contrarian content and amplify climate action messaging.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41558-025-02326-w>.

References

- Oreskes, N. *The Scientific Consensus on Climate Change: How Do We Know We're Not Wrong?* (Springer International Publishing, 2018).
- van der Linden, S. L., Leiserowitz, A. A., Feinberg, G. D. & Maibach, E. W. The scientific consensus on climate change as a gateway belief: experimental evidence. *PLoS ONE* **10**, e0118489 (2015).
- Oreskes, N. The scientific consensus on climate change. *Science* **306**, 1686 (2004).
- Bray, D. The scientific consensus of climate change revisited. *Environ. Sci. Policy* **13**, 340–350 (2010).
- Boehmer-Christiansen, S. Scientific consensus and climate change: the codification of a global research agenda. *Energy Environ.* **4**, 362–407 (1993).
- Stuart Carlton, J., Perry-Hill, R., Huber, M. & Prokopy, L. S. The climate change consensus extends beyond climate scientists. *Environ. Res. Lett.* **10**, 094025 (2015).
- Doran, P. T. & Zimmerman, M. K. Examining the scientific consensus on climate change. *EOS Trans. Am. Geophys. Union* **90**, 22–23 (2009).
- Cook, J. et al. Consensus on consensus: a synthesis of consensus estimates on human-caused global warming. *Environ. Res. Lett.* **11**, 048002 (2016).
- Cook, J. et al. Quantifying the consensus on anthropogenic global warming in the scientific literature. *Environ. Res. Lett.* **8**, 024024 (2013).
- Supran, G. & Oreskes, N. Assessing ExxonMobil's climate change communications (1977–2014). *Environ. Res. Lett.* **12**, 084019 (2017).
- Supran, G. & Oreskes, N. Rhetoric and frame analysis of ExxonMobil's climate change communications. *One Earth* **4**, 696–719 (2021).
- Supran, G., Rahmstorf, S. & Oreskes, N. Assessing ExxonMobil's global warming projections. *Science* **379**, eabk0063 (2023).
- Cook, J. in *Handbook of Research on Deception, Fake News, and Misinformation* (eds. Chiluba, I. & Samoilenko, S.) 281–306 (IGI Global, 2019).
- Farrell, J., McConnell, K. & Brulle, R. Evidence-based strategies to combat scientific misinformation. *Nat. Clim. Change* **9**, 191–195 (2019).
- Lewandowsky, S. Climate change disinformation and how to combat it. *Annu. Rev. Public Health* **42**, 1–21 (2021).
- Maertens, R., Anseel, F. & van der Linden, S. Combatting climate change misinformation: evidence for longevity of inoculation and consensus messaging effects. *J. Environ. Psychol.* **70**, 101455 (2020).
- Franta, B. *Climate Litigation Rising: Hot Spots to Watch* (American Bar Association, 2021); https://www.americanbar.org/groups/environment_energy_resources/publications/trends/2021-2022/january-february-2022/climate-litigation-rising/
- Franta, B. Early oil industry disinformation on global warming. *Environ. Politics* **30**, 663–668 (2021).
- Franta, B. Weaponizing economics: big oil, economic consultants, and climate policy delay. *Environ. Politics* **31**, 555–575 (2022).
- Grasso, M. Oily politics: a critical assessment of the oil and gas industry's contribution to climate change. *Energy Res. Soc. Sci.* **50**, 106–115 (2019).
- Kaup, C. Smoke gets in your eyes: misleading fossil fuel advertisement in the climate crisis. *J. Eur. Consumer Market Law* **10**, 21 (2021).
- McCright, A. M. & Dunlap, R. E. Challenging global warming as a social problem: an analysis of the conservative movement's counter-claims. *Soc. Probl.* **47**, 499–522 (2000).
- Mulvey, K. et al. *The Climate Deception Dossiers: Internal Fossil Fuel Industry Memos Reveal Decades of Corporate Disinformation* (Union of Concerned Scientist, 2015); <https://www.ucsusa.org/sites/default/files/attach/2015/07/The-Climate-Deception-Dossiers.pdf>

24. *Smoke, Mirrors & Hot Air: How ExxonMobil Uses Big Tobacco's Tactics to Manufacture Uncertainty on Climate Science* (Union of Concerned Scientists, 2007); https://www.ucsusa.org/sites/default/files/2019-09/exxon_report.pdf
25. Bolsen, T. & Shapiro, M. A. The US news media, polarization on climate change, and pathways to effective communication. *Environ. Commun.* **12**, 149–163 (2018).
26. Zhou, J. Boomerangs versus javelins: how polarization constrains communication on climate change. *Environ. Politics* **25**, 788–811 (2016).
27. McCright, A. M. & Dunlap, R. E. The politicization of climate change and polarization in the American public's views of global warming, 2001–2010. *Sociol. Q.* **52**, 155–194 (2011).
28. Merkley, E. & Stecula, D. A. Party elites or manufactured doubt? The informational context of climate change polarization. *Sci. Commun.* **40**, 258–274 (2018).
29. Boykoff, M. T. & Boykoff, J. M. Balance as bias: global warming and the US prestige press. *Glob. Environ. Change* **14**, 125–136 (2004).
30. Antilla, L. Climate of scepticism: US newspaper coverage of the science of climate change. *Glob. Environ. Change* **15**, 338–352 (2005).
31. Koehler, D. J. Can journalistic “false balance” distort public perception of consensus in expert opinion?. *J. Exp. Psychol. Appl.* **22**, 24–38 (2016).
32. Brüggemann, M. & Engesser, S. Beyond false balance: how interpretive journalism shapes media coverage of climate change. *Glob. Environ. Change* **42**, 58–67 (2017).
33. Farrell, J. Corporate funding and ideological polarization about climate change. *Proc. Natl. Acad. Sci. USA* **113**, 92–97 (2016).
34. *Climate Disinformation Database* (DeSmog, accessed May 2022); <https://www.desmog.com/climate-disinformation-database/>
35. Brulle, R. J. Networks of opposition: a structural analysis of US climate change counter-movement coalitions 1989–2015. *Sociol. Inq.* **91**, 603–624 (2021).
36. Brulle, R. J., Hall, G., Loy, L. & Schell-Smith, K. Obstructing action: foundation funding and US climate change counter-movement organizations. *Clim. Change* **166**, 1–7 (2021).
37. Brulle, R. J. & Werthman, C. The role of public relations firms in climate change politics. *Clim. Change* **169**, 8 (2021).
38. Ali, M. et al. Discrimination through optimization: how Facebook's ad delivery can lead to biased outcomes. In *Proc. ACM on Human-Computer Interaction 3* (eds Lampinen, A., Gergle, D. & Shamma, D. A.) 199 (Association for Computing Machinery, 2019).
39. Imana, B., Korolova, A. & Heidemann, J. Auditing for racial discrimination in the delivery of education ads. In *Proc. 2024 ACM Conference on Fairness, Accountability, and Transparency* (eds Leskovec, J. et al.) 2348–2361 (Association for Computing Machinery, 2024).
40. Imana, B., Korolova, A. & Heidemann, J. Auditing for discrimination in algorithms delivering job ads. In *Proc. Web Conference 2021* (eds Leskovec, J. et al.) 3767–3778 (Association for Computing Machinery, 2021).
41. Ali, M., Goetzen, A., Mislove, A., Redmiles, E. M. & Sapiezynski, P. Problematic advertising and its disparate exposure on Facebook. In *Proc. 32nd USENIX Security Symposium (USENIX Security 23)* (eds Calandrino, J. & Troncoso, C.) 5665–5682 (USENIX, 2023).
42. Ali, M., Sapiezynski, P., Korolova, A., Mislove, A. & Rieke, A. Ad delivery algorithms: the hidden arbiters of political messaging. In *Proc. 14th ACM International Conference on Web Search and Data Mining* (eds Lewin-Eytan, L. et al.) 13–21 (Association for Computing Machinery, 2021).
43. *About Ad Auctions* (Meta, 2024); <https://www.facebook.com/business/help/430291176997542?id=561906377587030>
44. Goldberg, M. H., Gustafson, A., Rosenthal, S. A. & Leiserowitz, A. Shifting Republican views on climate change through targeted advertising. *Nat. Clim. Change* **11**, 573–577 (2021).
45. Lambrecht, A. & Tucker, C. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Manage. Sci.* **65**, 2966–2981 (2019).
46. Leiserowitz, A., Maibach, E., Roser-Renouf, C. & Smith, N. *Global Warming's Six Americas* (Yale Univ. & George Mason Univ., 2011).
47. Roser-Renouf, C., Stenhouse, N., Rolfe-Redding, J., Maibach, E. & Leiserowitz, A. in *Routledge Handbook of Environment and Communication* (eds Hansen, A. & Cox, R.) 388–406 (Routledge, 2015).
48. Marlon, J., Howe, P., Mildenberger, M., Leiserowitz, A. & Wang, X. *Yale Climate Opinion Maps 2020* (Yale Program on Climate Change Communication, 2020).
49. Leiserowitz, A. et al. *Global Warming's Six Americas* (Yale Univ. & George Mason Univ., 2023).
50. Zhang, B. et al. Experimental effects of climate messages vary geographically. *Nat. Clim. Change* **8**, 370–374 (2018).
51. How greenwashing fools US. *The New York Times* (23 August 2022); <https://www.nytimes.com/2022/08/23/climate/climate-greenwashing.html>
52. De Jong, M. D. T., Harkink, K. M. & Barth, S. Making green stuff? Effects of corporate greenwashing on consumers. *J. Bus. Tech. Commun.* **32**, 77–112 (2018).
53. Spampatti, T., Hahnel, U. J. J., Trutnevyte, E. & Brosch, T. Psychological inoculation strategies to fight climate disinformation across 12 countries. *Nat. Hum. Behav.* **8**, 380–398 (2024).
54. *Oil Well Pad in the Bakken Formation in North Dakota* (US Geological Survey, 2009); <https://www.usgs.gov/media/images/oil-well-pad-bakken-formation-north-dakota>
55. Avera Hunter, B. *Solar Panel of Photovoltaic Cells* (US Geological Survey, 2010); <https://www.usgs.gov/media/images/solar-panel-photovoltaic-cells-1>
56. Doyle, J. Picturing the clima(c)tic: Greenpeace and the representational politics of climate change communication. *Sci. Cult.* **16**, 129–150 (2007).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2025

Methods

Ethics

This research, exempt from Massachusetts Institute of Technology Committee on the Use of Humans as Experimental Subjects (ID: E-4191), involved ad experiments using images of objects featured in ads by both contrarian and advocacy actors, whose logos were included. Ads complied fully with Facebook's Ad Review policies and featured climate-related images without using climate-related text to avoid spreading disinformation.

Observational study

We conducted an observational study of 203K US climate ads on Facebook. We identified contrarian and advocacy actors, compiled a dataset of past ads and analysed differences in delivery patterns.

Dataset

We assemble our actors dataset using a peer-reviewed list of contrarian and advocacy actors^{33–37}. We focused on fossil fuel corporations and their affiliates as contrarians, and environmental groups and renewable energy providers as advocates. Any ad from these groups is labelled as a 'climate' ad. From May 2018 to May 2023, we extracted 81,248 contrarian ads and 171,877 advocacy ads, filtering the dataset to 63,542 contrarian and 139,012 advocacy ads after removing duplicates and non-US ads. The ads were viewed for 1.36 million days, shown 5.4–6.4 billion times and cost 79–133 millions. To simplify the analysis and anonymize the data, each row in our dataset corresponded to delivery in a single US state or to one gender or age group. We released this decomposed data but withheld the full list of advertisers, following precautions from previous studies^{35–37}. The scrubbing and decomposing methodology are described in our code in 'Dataset availability'. The full de-anonymized dataset of advertisers, and ads, may be obtained by emailing arunas@mit.edu. (See the section on 'Dataset availability' for a link to the anonymized dataset.)

Pre-analysis

We analysed the collected data. Ignoring ad content, we focused our attention on ad delivery. We now describe possible ad destinations, and corresponding delivery attributes and proxies to surmise targeting intent.

Ad destinations and delivery information. Each ad in the dataset contained attributes describing delivery information for three possible ad destinations: US state, gender and age.

For US states, the `delivery_by_region` attribute of each ad contained delivery information for US state destinations. This attribute is a list of n tuples, $n \in [1, 52]$. The i th tuple is $(region_i, delivery_percentage_i)$, where $region_i$ can be one of 52 locations, including 50 US states, Washington DC and an 'unknown' category, and $delivery_percentage_i$ is a fraction such that $delivery_percentage_i \in [0, 1]$ and $\sum_{i=1}^{52} delivery_percentage_i = 1$.

For age and gender destinations, the `demographic_distribution` attribute of each ad contained delivery information for age and gender destinations. This attribute is a list of n tuples, $n \in [1, 24]$. The i th tuple is $gender_i, age_i, delivery_percentage_i$, where $age_i \in [18–24, 25–34, 35–44, 45–54, 55–64, 65+, unknown]$ and $gender_i \in [male, female, unknown]$. $delivery_percentage_i$ is a fraction such that $delivery_percentage_i \in [0, 1]$ and $\sum_{i=1}^{52} delivery_percentage_i = 1$.

The delivery data show the audiences, locations and algorithm's role in delivering contrarian and advocacy ads. Ideally, both targeting and delivery data would separate the roles of advertisers and algorithms, but Facebook only provides delivery information according to US state, age and gender. Thus, we used proxies in the delivery data to infer targeting intent.

Targeting proxies. We used two proxies to infer targeting according to location and demographics. Ads delivered to a single US state, gender

or age group were considered 'targeted', while ads reaching at least two US states (Supplementary Fig. 1), all genders and all age groups were considered 'non-targeted' ads. As a robustness check, varying the number of US states between 2 and 50 yielded similar results. We hypothesize that algorithmic decision-making is evident in these ads, as suggested by previous research^{38,40,42}.

Comparisons of t -tests

We calculated AIPD to assess ad effectiveness across ad destinations. Impressions and spend are presented in range format in the dataset. We approximated them into point values by calculating the midpoint of the range. If $AIPD_{advocacy\ ad}$ was greater than $AIPD_{contrarian\ ad}$ for an audience, we considered advocacy ads cheaper for that audience. This was repeated for targeted and non-targeted ads. We assumed that higher spend would lead to a larger audience, although this is hard to confirm because of the range format of the data.

For each possible location, gender or age category, we verified if the distribution of advocacy and contrarian ads consisted of at least 30 samples, to satisfy assumptions of normality. If the distribution contained more than 30 samples, we verified if the $AIPD_{advocacy\ ad}$ and $AIPD_{contrarian\ ad}$ were homoscedastic using Levene's test. A two-sample Welch's t -test with Bonferroni correction was then used to investigate if the statistical distribution of $AIPD_{advocacy\ ad}$ was significantly different from that of $AIPD_{contrarian\ ad}$ for each location, gender and age category. For example, this analysis would be conducted for each US state, and reported at the state level, when we investigated audiences in each US state. If the distribution contained fewer than 30 samples, a two-sample Mann–Whitney U -test was used to investigate if the statistical distribution of $AIPD_{advocacy\ ad}$ was significantly different from that of $AIPD_{contrarian\ ad}$ for each location, gender and age category, after establishing the homoscedasticity assumption of the sample distributions using a Levene's test.

If the Welch's t -test or the Mann–Whitney U -test revealed that the statistical distribution of $AIPD_{advocacy\ ad}$ was significantly different from that of $AIPD_{contrarian\ ad}$ ($P < 0.05$), the sample means or medians of the sample distributions (depending on whether the data were normal or not) were used to further determine if contrarian or advocacy ads were more effectively targeted in the location, age or gender category being considered. We also applied the Benjamini–Hochberg procedure to control the FDR in multiple comparisons.

For ads delivered to a single location, or gender or age type, normalized ad counts were compared for contrarian and advocacy ads.

Regression analysis

For the regression analysis, data were preprocessed into rows with the ad ID, and the delivery start and end times, US state category, AIPD value, state oil production (millions of barrels), population (millions), per-capita GDP (millions) and a binary value of 1 if a state leaned Republican in the 2020 elections, or 0 if a state leaned Democrat in the 2020 elections. This dataset was used for a fixed-effects regression analysis, initially specified as follows:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + (\text{state}) + (\text{year}) + \epsilon$$

where Y is the dependent variable and X_1, X_2, X_3 and X_4 are the covariates representing the annual state population, annual oil production, Republican leaning and annual per-capita GDP in the state, respectively. Per-capita GDP was considered instead of state-level GDP so that it was not perfectly correlated with population. ϵ represents the error term. States and years in which the ad was run were added as fixed effects. Ordinary least squares (OLS) regression showed an exponential distribution for the residuals, so the original model was modified to be:

$$\log(Y) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4 + (\text{state}) + (\text{year}) + \epsilon$$

However, initial diagnostic checks, including Q–Q plots, showed significant deviations from normality and heavy-tailed residuals, with excess kurtosis greater than 1. These findings indicated that some data points were unduly influencing the model, leading to potential bias if OLS regression was used. We applied transformations (square root, inverse, Box–Cox, Yeo–Johnson) and winsorized AIPD at the 0.25 and 0.75 percentiles; however, heavy-tailed residuals persisted. To address this, we used robust linear regression with Tukey’s biweight M-estimator, reducing the influence of outliers by down-weighting large residuals. This approach was applied to all regressions in the study.

Field experiment

To isolate the algorithm’s influence, we launched ads of oil drills and solar cells on Facebook without enforcing constraints on audience selection. We selected ad images of oil drills and solar cells and affixed logos on them, as described in the section ‘Ad image selection’. We launched 650 ads with an equal distribution of ads featuring oil drills, solar cells, and oil drills and solar cells with the logos of ExxonMobil and Greenpeace (see the section on ‘Experimental groups’). The ads were launched on Facebook without targeting, allowing delivery to all US states, genders and age groups for 24 h. Facebook was instructed to optimize delivery to likely clickers, according to a design similar to prior research³⁸. When a user clicked on an ad, they were taken to a website, but this was irrelevant to our investigation and simply created a realistic ad experience for both the algorithm and user³⁸. After the 24-h experiment, ad delivery data were collected from Facebook to compare the delivery of contrarian and advocacy ads.

Experimental design

We aimed to isolate the algorithm’s influence on climate ad delivery. However, using real ads from ExxonMobil or Greenpeace risks promoting fossil fuels or inadvertently influencing the climate information ecosystem. Inspired by previous research³⁸, we ran a field experiment that primarily varied the images used in an ad as a workaround. The ad text was neutral and unrelated to climate change. Note that images of oil drills or solar cells are not inherently contrarian or advocacy images as ads from both ExxonMobil and Greenpeace feature these two objects; while Greenpeace may request audiences to stop using fossil fuels with the image of an oil drill, ExxonMobil may talk about the importance of the oil economy with the same image. Therefore, the messaging in our ads was implicit. Once the images were selected (see the section on ‘Ad image selection’), ads were launched in three batches. This is because Facebook has an upper limit of 250 concurrent ads for advertisers whose budget is less than US\$1,000,000 per month. Batch 1 (22 ads per campaign) ran from 9:01 Pacific Standard Time (PST) on 17 May 2023 to 9:01 PST, on 18 May 2023. Batch 2 (22 ads per campaign) ran from 16:01 PST on 18 May 2023 to 16:01 PST on 19 May 2023. Batch 3 (21 ads per campaign) ran from 4:01 PST on 20 May 2023 to 4:01 PST on 21 May 2023. Ads ran simultaneously across all time zones in the United States, controlling for market effects related to user traffic, ad auctions and bid pricing to the extent possible. Each ad image of a solar cell or oil rig was accompanied by the text ‘Use our website to tell us what you think about this picture’, along with a neutral website (<https://imagesweseepage.wordpress.com/>) created solely for the purpose of this experiment. All ads were allocated a budget of US\$1 per day and optimized to maximize audience traffic to the website—a design choice that was again inspired by past research³⁸. In accordance with Facebook guidelines, all ads were launched under the ‘social issues, elections or political issues’ category. Lastly, and most importantly, ads were scheduled to be delivered to anyone in the United States who matched Facebook’s default age criteria, irrespective of gender or location. We did not use features that allowed for interest-based targeting, which would allow us to promote the ads to users interested in climate change. This allowed us to control for the demographic characteristics preferred by Facebook. All design choices were accessible as features

on the Facebook Ad Manager (see the section on ‘Full analysis’). The groups used in our experiment to identify group-level differences are described below.

Experimental groups

The 650 ads that were launched contained images of solar cells, oil rigs and controls, with the same neutral text. We used three experimental groups to investigate delivery dependence on ad images relevant to climate change: (1) images—solar cells (65 images) and oil drills (65 images) unmodified. We added (i) controls—control images (65 images) and (ii) duplicates—duplicate ads using images from the images and control group to check that delivery was consistent with ad content (65 images × 3); (2) images + contrarian logo—solar cells and oil drills with the logo of ExxonMobil, an established contrarian organization^{10,11} on the top left (65 images × 2); (3) images + advocacy logo × solar cells and oil drills with the logo of Greenpeace, an established advocacy organization^{57,58}, on the top left (65 images × 2).

Ad image selection

In our field study, ad images were the principal focus. We selected images classified as oil drills or solar cells with at least 70% probability using a state-of-the-art image classifier⁵⁹, and control images that were unrelated to these classes. Solar cell and oil rig images were manually collected from internet searches, while control images were sampled from the ImageNet-21K dataset⁵⁹ using the WordNet hierarchy⁶⁰. ImageNet-21K includes 21,841 classes mapped to WordNet noun synsets. To ensure diversity, we randomly sampled categories that were distant in the WordNet tree, with a requirement that sampled images should meet the Facebook Ad Platform size criteria (≥600 px width/height). We constructed the WordNet tree for all ImageNet-21K labels, with ‘entity’ as the root. Starting at level 6, which contained over 300 nodes, we randomly sampled 300 terminal nodes, ensuring they were unrelated and their categories included images meeting the size criteria. Of these, only 103 categories satisfied the size requirement; from this set, we randomly selected 65 control images. All sampling was performed using Python’s random library.

Ad campaign attributes

After sampling the images, we launch ads via Facebook Ad Manager, selecting the following attributes to form the core elements of our experimental design: (1) duration—652 ad campaigns were run in three batches, such that each batch was run for a period of 24 h to reach all time zones of the United States; (2) Ad media—we used images of oil rigs, solar cells or controls (Fig. 3). Each image was modified with the logo of a contrarian or advocacy organization, depending on the treatment group it was assigned to; (3) Ad text—for each ad, we included the text ‘Use our website to tell us what you think about this picture’; (4) desired audience attributes—the ads were scheduled to be delivered to anyone in the United States who belonged to the default age criteria on Facebook, irrespective of gender and location. We did not use any additional microtargeting features; (5) ad placement—we specified that the ads could only be shown on the Facebook platform and could only be situated on a user’s account Facebook feed; (6) ad budget—we specified a daily ad budget of US\$1 per day; (7) campaign objective—we specified that the ads’ objective was to maximize audience traffic to the website. This website collected opinions about the ad images when shared by a visitor. It did not contain content that revealed the intentions of our experiment, or a stance on climate action or climate change; (8) ad type—we ran the ads under the ‘social issues, elections or political issues’ category, in accordance with Facebook’s advertising guidelines.

Covariate analysis

We conducted a covariate analysis to identify image features that significantly affected the delivery ratio D . Covariates included object

coverage percentage (oil drills or solar cells), background, logo type (ExxonMobil, Greenpeace or none), image type and classification probability. The outcome variable, D , was modelled as a linear combination of these covariates. Because of the right-skewed distribution of D across ad destination types (US state, gender, age), OLS regression resulted in non-normal residuals. Therefore, we used a generalized linear model with a gamma distribution and a log link function to account for the multiplicative relationship between predictors and D . A Wald test was applied to assess the significance of individual coefficients and their influence on D . The covariate analysis for US states (Supplementary Table 35) showed that while the logo type and its interactions with the ad image did not significantly affect the delivery ratio ($P > 0.05$), image type, US state, their interaction and the three-way interaction with the logo type all had a significant impact ($P < 0.001$). The percentage of image coverage and classifier probability also influenced the delivery ratio ($P = 0.002$, $P = 0.046$). In gender-based analysis (Supplementary Table 36), image type, audience gender and their interactions, including the three-way interaction with the logo type, significantly affected the delivery ratio ($P < 0.001$), while object coverage and classifier probability did not ($P > 0.05$). For the age-based analysis (Supplementary Table 37), all covariates and interactions, except for oil drill probability, significantly influenced the delivery ratio ($P < 0.001$).

Full analysis

As we ceded control over ad delivery to Facebook, we compared the delivery ratio D of advocacy and contrarian ads across ad destination categories (US state, gender, age group). D is the ratio of observed to expected delivery, based on Facebook's ad audience estimates, which are provided at ad creation and appear to be stable over short time periods. This approach allowed us to analyse the algorithm's influence by comparing its observed decision-making to a uniform sampling baseline. To calculate D , we collected the 'reach' data—unique accounts shown an ad, segmented according to state, gender and age—and Facebook's audience estimates, which represent the expected delivery proportional to audience size for each category. The D for an ad destination category c and ad i is given as $D_i^c = O_i^c/E_i^c$. Here, $D \in \mathbb{R}^+$, and $O_i^c, E_i^c \in \mathbb{N}$. O_i^c is the unique number of times an ad i was shown in an ad destination (US state, gender or age) c , and E_i^c is Facebook's estimated reach of the ad for the same category if audiences are uniformly sampled from the corresponding population categories.

Most samples collected from Facebook were non-normal and heteroscedastic. To address research questions 1 and 2, we applied a one-way Kruskal–Wallis test to detect significant differences across D samples from different groups. Pairwise differences were analysed with Dunn's test, adjusting P values using the Benjamini–Hochberg correction; we report medians. For research question 3, a Fisher's test was used to assess consistency in audience sizes when the same image was delivered twice. For research question 4, we used a chi-squared goodness of fit test with Bonferroni correction to check if audience sizes aligned with Facebook's reported estimates.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The full observational and field experiment datasets are available on Zenodo at <https://doi.org/10.5281/zenodo.14914203> (ref. 60).

Code availability

The code used in this study is available on Zenodo at <https://doi.org/10.5281/zenodo.14914203> (ref. 60).

References

- Alam, M., Nilan, P. & Leahy, T. Learning from Greenpeace: activist habitus in a local struggle. *Electron. Green J.* **1**, 1–18 (2019).
- Ridnik, T., Ben-Baruch, E., Noy, A. & Zelnik-Manor, F. ImageNet-21K pretraining for the masses. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)* (2021).
- Fellbaum, C. in *Theory and Applications of Ontology: Computer Applications* (eds Poli, R., Healy, M., & Kameas, A.) 231–243 (Springer, 2010).
- Sankaranarayanan, A., Hemberg, E., Sapiezynski, P. & O'Reilly, U.-M. The Facebook algorithm's active role in climate advertisement delivery dataset and code. *Zenodo* <https://doi.org/10.5281/zenodo.14914203> (2025).

Acknowledgements

We thank B. Lipkin, M. Groh, H. Okabe, E. Hemberg and especially K. Arquilla, for their guidance on statistical inference. We thank S. Srikanth and E. Hemberg for valuable ideas on robust experimental design. We thank G. Supran and C. Hickey for aggregating and sharing the list of contrarian and advocate actors from peer-reviewed research, from which we retrieved the set of actors considered in this work. Finally, we thank the Climate Disinformation Research Coalition, including N. Oreskes, G. Supran, C. Hickey and U.-M. O'Reilly, and Influence Map (F. Holder and H. Groh for their feedback on the initial versions of this draft). We also thank S. Satish, A. Ranka, S. Satish and S. Narreddi at OwnPath, and R. Kumar Sah, M. Ahmad and P. Verma at Idea2MVP (<https://ramailo.tech/>) for helping us build an interface to explore the climate discourse on Facebook.

Author contributions

A.S. and U.-M.O. conceived the research question. A.S., P.S. and U.-M.O. designed the experiment. A.S. developed the experimental stimuli and conducted the experiments. A.S., P.S. and U.-M.O. analysed and interpreted the results. A.S., P.S. and U.-M.O. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41558-025-02326-w>.

Correspondence and requests for materials should be addressed to Aruna Sankaranarayanan.

Peer review information *Nature Climate Change* thanks Jiaojiao Ji and the other, anonymous, reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

Climate contrarian and advocacy advertisers were identified based on prior peer-reviewed research. Advertisements from these advertisers were retrieved from the Facebook Ad Library, covering the period from May 2018 to May 2023. An initial dataset of 81,248 contrarian ads and 171,877 advocacy ads was collected, which was subsequently refined by removing duplicates and non-U.S. advertisements, resulting in 63,542 contrarian ads and 139,012 advocacy ads. Each advertisement may have been disseminated to audiences in one or multiple U.S. states, as well as to specific demographic groups based on gender or age. To facilitate analysis while preserving anonymity, each row in the dataset represents ad delivery within a single U.S. state or to a specific gender or age group. The decomposed dataset is publicly released; however, consistent with ethical precautions established in prior studies, the complete list of advertisers is withheld.

Data analysis

Data was analyzed using Python notebooks provided on Zenodo: <https://zenodo.org/records/14942030>

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Data is available on <https://zenodo.org/records/14942030>. Imagenet-21K dataset is available on <https://www.image-net.org/>

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	Sex and gender classifications were used in accordance with the terminology employed by Facebook to ensure consistency. Facebook categorizes gender into three groups: male, female, and unknown. For this analysis, we excluded the "unknown" category and focused exclusively on data pertaining to the "male" and "female" categories.
Reporting on race, ethnicity, or other socially relevant groupings	N/A
Population characteristics	N/A
Recruitment	Participants were not recruited.
Ethics oversight	MIT COUHES, proposal E-4191

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

- Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	We conduct both an observational and an experimental study to isolate the impact of Facebook's algorithm on the delivery of climate-related advertisements—specifically, ads published by climate contrarian and climate advocacy groups. The experimental study further isolates the role of the algorithm by ensuring that advertising parameters were fully controlled by the platform. All data collection and analyses are quantitative. Detailed descriptions of both studies are provided in the Methods section.
Research sample	The observational study examines advertisements from climate advocacy and contrarian organizations collected between May 2018 and May 2023. These organizations were identified based on prior peer-reviewed research. The final research sample for the observational study comprises 63,542 contrarian ads and 139,012 advocacy ads. To facilitate analysis while preserving anonymity, each row in the dataset corresponds to ad delivery in a single U.S. state or to a specific gender or age group. The decomposed dataset, which contains 4,217,991 rows, is publicly released, while the full list of advertisers is withheld in accordance with ethical precautions established in prior studies. Similarly, the experimental study involved running 650 advertisements with a budget of \$1 per ad. The resulting dataset was decomposed using the same methodology, yielding 22,024 observations when analyzing audience effects by U.S. state, 1,300 observations when analyzing by gender, and 3,415 observations when analyzing by both gender and state. The sample is not representative, as ads were run for a single day with a minimal budget. As such, the findings provide preliminary evidence of algorithmic bias rather than definitive conclusions.
Sampling strategy	Sample sizes for the observational study were determined by the maximum number of available samples at the time of data collection, as the Facebook Ad Library provides data from May 2018 onward. For the experimental study, sample sizes were informed by a power analysis assuming a medium to large effect size. Images used in the experimental study were selected based on Facebook Ads Manager's requirements for image size and resolution. Additionally, they were classified as either solar cells or oil drills with at least 70% accuracy. To enhance statistical power in our primary analysis (Fig. 4), we included both images with and without logos, as no significant differences in ad delivery were observed between these categories. The final experimental dataset consists of 260 oil drill and solar cell ads and 130 control ads.

While power analyses provide useful guidance, they are inherently based on assumptions (<https://datacolada.org/4>). This study does not evaluate all possible pricing and targeting strategies but offers initial evidence on potential algorithmic bias in Facebook's ad delivery system.

Data collection

Observational data was collected from May 2018 - 2023 from the Facebook ads library. The experimental study was launched after results were aggregated for the observational study to clearly control for ad targeting parameters to more clearly identify bias, if any, in the algorithm. At the time of collecting observational data, <https://dl.acm.org/doi/pdf/10.1145/3359301> had already been released, so hypotheses were somewhat known before collection of data. At the time of collecting experimental data too, the hypotheses were somewhat known since the experimental study was launched to investigate the results from the observational study when control is more clearly assigned to the algorithm. Experimental data was collected directly from the Facebook ads interface as a CSV file after running the ads. Observational data was collected directly from the Facebook ads manager using their Ads API. Both observational and experimental data were collected by the corresponding author.

Timing

Observational data was collected between May 2018-May 2023. Experimental data was collected in 3 batches on 3 consecutive days -- 17th, 18th, 19th May 2023 to comply with Facebook's advertising policies. More details in the Methods section of the paper.

Data exclusions

Duplicate data and ads delivered to non U.S. locations were excluded from the observational study. No data was excluded from the experimental study.

Non-participation

No human participants were recruited or involved in the study.

Randomization

Ad images were randomly allocated to different experimental batches when running the experimental study.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

- | n/a | Involvement in the study |
|-------------------------------------|--|
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Antibodies |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Eukaryotic cell lines |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Palaeontology and archaeology |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Animals and other organisms |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Clinical data |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Dual use research of concern |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Plants |

Methods

- | n/a | Involvement in the study |
|-------------------------------------|---|
| <input checked="" type="checkbox"/> | <input type="checkbox"/> ChIP-seq |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Flow cytometry |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> MRI-based neuroimaging |

Plants

Seed stocks

Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.

Novel plant genotypes

Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.

Authentication

Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.