Is the Facebook Ad Algorithm a Climate Discourse Influencer?

Anonymous Author(s)
Affiliation
Address
email

Abstract

Sponsored climate discourse, driven by both climate contrarians and advocates, influences public attitudes towards climate change. We present an experimental study suggesting that the Facebook advertisement algorithm also influences climate discourse. The algorithm preferentially delivers ads to Facebook audiences in certain locations and demographics, at least partially based upon the ad image. Further, the algorithm is biased in terms of how it delivers ads featuring images of non-renewable sources of energy, and does not always fulfill targeting intentions as requested. This may result in inadvertent manipulation of ad delivery with consequences for climate discourse and algorithmic fairness.

In 2022, the Intergovernmental Panel on Climate Change (IPCC) identified "rhetoric, misinformation, and politicization of science" as key barriers to climate action, further stating that "vested economic and political interests have organized and financed misinformation and contrarian climate change communications". Scholars across disciplines have documented the deceptive nature of contrarian climate communications [41, 42, 43, 15, 17, 28, 29, 2, 19, 20, 22, 24, 31, 1, 4, 41, 42, 43], but the roles of intermediary algorithms in delivering this content has been overlooked. These algorithms comprise deep image networks and automated recommendation systems, and their implicit bias has only recently begun to be studied by computer science researchers [8, 9, 7].
In 2022, Rolnick et al [36] presented a set of climate change issues that could be tackled by the machine learning community. A missing item in their list is climate change disinformation and the role of algorithmic bias in perpetuating it. To date, disinformation activities have been studied in print and broadcast media, rather than social media platforms [11, 49, 32, 33, 12, 10, 25, 13]. This blind spot is glaring given the volume of climate discourse on social media platforms. Analyzing this data and reducing bias in these systems, could improve the effectiveness of online climate action campaigns, support litigation[48] efforts, and advance inoculation and communication strategies against climate disinformation [14, 44, 48, 35, 21], [47, 5, 40].
Current studies have restricted the notion of climate discourse influencers to commercial enterprises (eg. fossil fuel companies), financial institutions (eg. banks), and lobby organizations (eg. think tanks). However, the algorithmic systems on social media, such as Facebook's ad algorithm, may also deserve recognition. Among strong reasons for this are that it is possible that in the course of providing maximum engagement for the cheapest cost, some inadvertent biases in other dimensions have crept into these very large, highly complex, data driven and continuously updating ad delivery systems [8, 9, 46, 39, 7].
In this study, we conduct an ad campaign, under controlled and ethical conditions, where we assume the role of a Facebook advertiser. We design simple and informative experiments to uncover the foundations of bias in the ad algorithm, based on ads that are primarily an image. Our experiment supports the existence of algorithmic bias though more extensive studies are needed to gauge the exact extent of such bias. We also discuss the implications of such algorithmic decision making on the climate discourse.

Submitted to Computational Sustainability Workshop at NeurIPS 2023. Do not distribute.

1 Experiment

To isolate the algorithm's influence, we do not use the targeting features provided to advertisers. Instead, we request the algorithm to impartially deliver ads to audiences in all U.S. states, and of all genders and ages. We launch the ads for a period of 24 hours, and further request Facebook to optimize delivery to reach audiences likely to click on them. This design is similar to the experiments found in prior research [8]. When a Facebook user clicks on an ad, they are taken to a website. At the end of the 24h experiment, ad delivery data is collected from Facebook and comparisons are made between the delivery information of contrarian and advocacy ads.

Ad delivery data Given that we cede full control over the ad targeting to Facebook, we use the Delivery Ratio (D) of advocacy and contrarian ads for each ad destination category c (U.S. state, gender category, age group category). D is a measure of delivery observed during the experiment and delivery expected, given Facebook's ad audience estimates. The expected estimates are provided by Facebook at ad creation time (and only change minimally over several weeks or months). To calculate D, first the 'Reach' information is collected for each launched ad. This is a total count of unique Facebook accounts that were shown an ad, broken up by U.S. state, gender, and age group. Second, Facebook's self-reported audience estimates are collected for each U.S. state, and age group. These provide a measure of the expected delivery count that is proportional to the audience size matching an ad destination. Facebook also advertises these population estimates as being the population sizes from which an ad audience sample will be drawn. The 'Delivery Ratio' (D) for an ad destination category c and and ad i is given by $D_i^c = \frac{O_i^c}{E_i^c}$. 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(U.S. state, gender, or age) c, and E_i^c is Facebook's estimated reach of the ad for the same category.

1.1 Experimental Groups

(a) Oil rigs (x65) (b) Solar cells (x65) (c) Controls (x65) (d) Oil rig + ExxonMobil logo (x65) (e) Solar cell + ExxonMobil logo (x65) (f) Oil rig + Greenpeace logo (x65) (g) Solar cell + Greenpeace logo (x65)

Figure 1: Examples of images without logo, and images with logos of ExxonMobil and Greenpeace representing contrarian and advocacy actors. Ads containing images without logos are launched twice to check whether Facebook delivery is consistent.

We create 650 ads containing images of solar cells and oil rigs, and the same neutral text. This sample size is chosen based on an *a priori* power analysis, indicating 65 images per group to uncover a small to moderate effect size of 0.25, with power=0.8. We design 3 experimental groups to investigate delivery dependence on ad images relevant to climate change.

1. **Images** - Solar cells (65 images) and oil rigs (65 images) unmodified. We add
   (a) **Controls** - Control images (65 images)
   (b) **Duplicates** - Duplicate ads using images from the *Images* and *Controls* group to check that delivery is consistent with ad content (65 images x 3).
2. **Images + Contrarian Logo** - Solar cells and oil rigs with the logo of ExxonMobil, an established contrarian organization[41, 42] on the top left (65 images x 2).
3. **Images + Advocacy Logo** - Solar cells and oil rigs with the logo of Greenpeace, an established advocacy organization [45, 16, 6] on the top left (65 images x 2).

We select images of oil rigs and solar cells in the experimental ads because these objects are found across the U.S., and featured in both contrarian and advocacy ads. While contrarians advertise oil rigs to highlight engineering capabilities and economic advantages, advocates use them to campaign against drilling. Similarly, contrarian ads use solar cells to highlight their contributions to climate action, and advocates use them to promote the use of renewable energy. 65 images featuring each of these objects are selected. We launch ads with and without logos to investigate if the algorithm decision making identifies and uses the advertiser information for delivery. A state-of-the-art image classifier[34] is able to distinguish our probes with high accuracy. We also sample 65 controls from the ImageNet-21K dataset, using a sampling process that excludes overlapping categories in the dataset. This ensures that controls are also able to be distinguished by a machine classifier. The algorithm to sample controls is provided in Appendix 6.3.

2 Research Questions

We pose the following research questions:

1. Can observed ad delivery be consistently attributed to the image?
2. Does ad delivery ratio, D, differ when logos are present on an ad image? Is the effect similar for ads containing images of solar cells and oil rigs with logos?
3. Does ad delivery ratio, D, differ based on the ad image?
4. Is observed ad delivery proportional to Facebook's population estimates across ad destinations?

3 Results

3.1 Main takeaways

1. **Consistency** Ad delivery is consistent with ad image: The audience sizes of 90% of the ads featuring the same image is consistent across U.S. states. The audience sizes of 100% and 99% of the ads are consistent across audiences of the same gender and age group respectively. See 6.4.1 for more details.
2. **Logo effects** The ad delivery ratio, D, is not significantly different for images featuring ExxonMobil or Greenpeace logos. D is also not influenced by the presence or absence of a logo in most ad destinations. We therefore combine images with and without logos when studying image effects for additional statistical power. See 6.8.4 for more details.
3. **Image effects** The ad delivery ratio, D, is significantly different based on whether the ad image features solar cells, oil rigs, or controls, in 46% of U.S. states, and in all gender and age destinations. In states where See 6.8.5 for more details.
4. **Promised vs Fulfilled Delivery**: Observed audience sizes are not always proportional to Facebook's audience estimates. Across U.S. states, audience sizes for *Controls* are more likely (64%) to be proportional to Facebook's population estimates than for images of solar cells and oil rigs (42.5%). Across males and females, audience sizes for solar cells and oil rigs are more likely (67%) than controls (22%) to be proportional to Facebook's population estimates. Across audience age intervals, neither images of solar cells nor oil rigs (0%) nor controls (0%), are proportional to Facebook's population estimates. See 6.8.6 for more details.

A detailed analysis of these results is provided in Appendix 6.4

4 Discussion

Algorithmic bias We show experimentally that climate ads featuring different content are consistently delivered differently. This suggests that climate advertising is vulnerable to algorithmic decision-making. We also find that delivery decisions made by Facebook's advertising algorithm are not proportional to Facebook's ad audience estimates for U.S. states, males/females and different age groups. This may also indicate that delivery skew is not arbitrary. While, we do not verify this in our experiment, we note that past research has determined that delivery decisions are largely driven by automated image classification on the algorithm's side, and not due to interactions of ad audiences with the ad. Startlingly, [8] showed that ads that appear invisible to a human (but visible to an automatic image classifier system) are delivered similarly to ads that are fully visible to humans.

Preferential Pricing Our ad experiments set a common budget of $1 on each ad. Our results indicating preferential delivery, therefore, also indicate preferential pricing. It is 'cheaper' to advertise images of oil rigs to males and older audiences and images of solar cells to females and younger audiences. Advocacy organizations are cash strapped, with one dataset discovering that 15% of advocacy ads request for donations or subscriptions [38]. Preferential pricing could therefore adversely impact the advertising strategy employed by advocacy organizations.

Audience impacts Our results highlight how advertising algorithms may impact the consumption of climate discourse by audiences with different psychological, cultural, and political responses to the climate crisis. The Six Americas Report[27] segments the U.S. population into six groups based on their response to climate action – Alarmed, Concerned, Cautious, Disengaged, Doubtful and Dismissive. Communication studies have noted that these groups require different persuasion strategies, and information channels, for climate engagement [37]. For example, audiences in the Doubtful and Dismissive category are best engaged by adopting non-confrontational approaches, and by framing messages in ways that are consistent with their values. Audiences in these groups are also more likely to be older individuals and male and located in the interior regions of the U.S. [26, 30] – demographics and regions where the algorithm preferentially delivers ads featuring pictures of an oil rig. Advocacy ads featuring images of oil rigs are more likely to use these images to dissuade audiences from fossil fuels and are more likely to be directed at male and older audiences who are more likely to be in segments that are Doubtful or Dismissive about the climate crisis. Such messaging from advocacy organizations may further alienate these groups from the climate cause. Similarly, contrarian ads featuring renewable energy sources such as solar cells may sometimes be used to promote an image of sustainable practice, a practice called greenwashing. Research is divided on whether these ads lead individuals to view actors as being more [3] or less sustainable [23]. Therefore the implications of the algorithm's recommendation of ads featuring solar cells to female and younger audiences where solar cell ads were delivered preferentially) is unclear. However, our experiment shows that these audiences may be more vulnerable to greenwashing ads from contrarian actors, and may therefore need to be inoculated more frequently against the practice.

5 Conclusion

We find early evidence indicating that social media platforms are a set of new and constantly evolving climate discourse influencers. While this study uncovers algorithmic influence, we cannot be fully sure of the precise extent of the algorithm's influence over ad delivery. One limitation of our study is that we run experiments with a common budget of $1. Experiments with higher ad budgets may reveal if algorithmic decision making is conditional on ad budget. We urge social media and disinformation scholars to not just study the proliferation of information on social media platforms but to also account for information delivery patterns to engage diverse audiences towards climate action.

References

[1] The climate deception dossiers. https://www.ucsusa.org/sites/default/files/attach/2015/07/The-Climate-Deception-Dossiers.pdf. Accessed: 2023-06-12.

[2] Climate litigation rising: Hot spots to watch. https://www.americanbar.org/groups/environment_energy_resources/publications/trends/2021-2022/january-february-2022/climate-litigation-rising/. Accessed: 2021-12-22.

[3] How greenwashing fools us. https://www.nytimes.com/2022/08/23/climate/climate-greenwashing.html. Accessed: 2022-08-23.

[4] Smoke, mirrors & hot air: How exxonmobil uses big tobacco's tactics to manufacture uncertainty on climate science. https://www.ucsusa.org/sites/default/files/2019-09/exxon_report.pdf. Accessed: 2023-06-12.

[5] Laurenz Aisenpreis, Gustav Gyrst, and Vedran Sekara. How do u.s. congress members advertise climate change: An analysis of ads run on meta's platforms. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, pages 2–11, 2023.

[6] Meredian Alam, Pam Nilan, and Terry Leahy. Learning from greenpeace: Activist habitus in a local struggle. *Electronic Green Journal*, 1(42), 2019.

[7] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa Redmiles, and Piotr Sapiezynski. All things unequal: Measuring disparity of potentially harmful ads on facebook. In *6th Workshop on Technology and Consumer Protection*, 2022.

[8] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. Discrimination through optimization: How facebook's ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction*, 3(CSCW):1–30, 2019.

[9] Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. Ad delivery algorithms: The hidden arbiters of political messaging. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 13–21, 2021.

[10] Liisa Antilla. Climate of scepticism: Us newspaper coverage of the science of climate change. *Global environmental change*, 15(4):338–352, 2005.

[11] Toby Bolsen and Matthew A Shapiro. The us news media, polarization on climate change, and pathways to effective communication. *Environmental Communication*, 12(2):149–163, 2018.

[12] Maxwell T Boykoff and Jules M Boykoff. Balance as bias: Global warming and the us prestige press. *Global environmental change*, 14(2):125–136, 2004.

[13] Michael Brüggemann and Sven Engesser. Beyond false balance: How interpretive journalism shapes media coverage of climate change. *Global Environmental Change*, 42:58–67, 2017.

[14] Wändi Bruine de Bruin, Andrew M Parker, and Baruch Fischhoff. Individual differences in adult decision-making competence. *Journal of personality and social psychology*, 92(5):938, 2007.

[15] John Cook. Understanding and countering misinformation about climate change. *Research Anthology on Environmental and Societal Impacts of Climate Change*, pages 1633–1658, 2022.

[16] Julie Doyle. Picturing the clima (c) tic: Greenpeace and the representational politics of climate change communication. *Science as culture*, 16(2):129–150, 2007.

[17] Justin Farrell, Kathryn McConnell, and Robert Brulle. Evidence-based strategies to combat scientific misinformation. *Nature climate change*, 9(3):191–195, 2019.

[18] Christiane Fellbaum. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer, 2010.

[19] Benjamin Franta. Early oil industry disinformation on global warming. *Environmental Politics*, 30(4):663–668, 2021.

[20] Benjamin Franta. Weaponizing economics: Big oil, economic consultants, and climate policy delay. *Environmental Politics*, 31(4):555–575, 2022.

[21] Matthew H Goldberg, Abel Gustafson, Seth A Rosenthal, and Anthony Leiserowitz. Shifting republican views on climate change through targeted advertising. *Nature Climate Change*, 11(7):573–577, 2021.

[22] Marco Grasso. Oily politics: A critical assessment of the oil and gas industry's contribution to climate change. *Energy Research & Social Science*, 50:106–115, 2019.

[23] Menno D. T. De Jong, Karen M. Harkink, and Susanne Barth. Making green stuff? effects of corporate greenwashing on consumers. *Journal of Business and Technical Communication*, 32(1):77–112, 2018. PMID: 30034204.

[24] Clemens Kaupa. Smoke gets in your eyes: Misleading fossil fuel advertisement in the climate crisis. *Journal of European Consumer and Market Law*, 1, 2021.

[25] Derek J Koehler. Can journalistic "false balance" distort public perception of consensus in expert opinion? *Journal of Experimental Psychology: Applied*, 22(1):24, 2016.

[26] A Leiserowitz, E Maibach, S Rosenthal, J Kotcher, M Ballew, J Marlon, J Carman, M Verner, S Lee, T Myers, et al. Global warming's six americas, december 2022. *Yale University and George Mason University. New Haven, CT: Yale Program on Climate Change Communication*, 2023.

[27] Anthony Leiserowitz, Edward Maibach, Connie Roser-Renouf, and Nicholas Smith. Global warming's six americas, may 2011. *Yale University and George Mason University*, 2011.

[28] Stephan Lewandowsky. Climate change disinformation and how to combat it. *Annual Review of Public Health*, 42:1–21, 2021.

[29] Rakoen Maertens, Frederik Anseel, and Sander van der Linden. Combatting climate change misinformation: Evidence for longevity of inoculation and consensus messaging effects. *Journal of Environmental Psychology*, 70:101455, 2020.

[30] Jennifer Marlon, Peter Howe, Matto Mildenberger, Anthony Leiserowitz, and Xinran Wang. Yale climate opinion maps 2020. *Yale program on climate change communication*, 2, 2020.

[31] Aaron M McCright and Riley E Dunlap. Challenging global warming as a social problem: An analysis of the conservative movement's counter-claims. *Social problems*, 47(4):499–522, 2000.

[32] Aaron M McCright and Riley E Dunlap. The politicization of climate change and polarization in the american public's views of global warming, 2001–2010. *The Sociological Quarterly*, 52(2):155–194, 2011.

[33] Eric Merkley and Dominik A Stecula. Party elites or manufactured doubt? the informational context of climate change polarization. *Science Communication*, 40(2):258–274, 2018.

[34] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses, 2021.

[35] Justin Rolfe-Redding, Edward W Maibach, Lauren Feldman, and Anthony Leiserowitz. Republicans and climate change: An audience analysis of predictors for belief and policy preferences. *Available at SSRN 2026002*, 2011.

[36] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2):1–96, 2022.

[37] Connie Roser-Renouf, Neil Stenhouse, Justin Rolfe-Redding, Edward Maibach, and Anthony Leiserowitz. Engaging diverse audiences with climate change: Message strategies for global warming's six americas. In *The Routledge handbook of environment and communication*, pages 388–406. Routledge, 2015.

[38] Aruna Sankaranarayanan, Erik Hemberg, Piotr Sapiezynski, and Una-May O'Reilly. The facebook algorithm's active role in climate advertisement delivery. *arXiv preprint arXiv:2308.03191*, 2023.

[39] Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Aaron Rieke, and Alan Mislove. Algorithms that" don't see color" measuring biases in lookalike and special ad audiences. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 609–616, 2022.

[40] Dietram A Scheufele and Nicole M Krause. Science audiences, misinformation, and fake news. *Proceedings of the National Academy of Sciences*, 116(16):7662–7669, 2019.

[41] Geoffrey Supran and Naomi Oreskes. Assessing exxonmobil's climate change communications (1977–2014). *Environmental Research Letters*, 12(8):084019, 2017.

[42] Geoffrey Supran and Naomi Oreskes. Rhetoric and frame analysis of exxonmobil's climate change communications. *One Earth*, 4(5):696–719, 2021.

[43] Geoffrey Supran, S Rahmstorf, and N Oreskes. Assessing exxonmobil's global warming projections. *Science*, 379(6628):eabk0063, 2023.

[44] Kathie M d'I Treen, Hywel TP Williams, and Saffron J O'Neill. Online misinformation about climate change. *Wiley Interdisciplinary Reviews: Climate Change*, 11(5):e665, 2020.

[45] Haridimos Tsoukas. David and goliath in the risk society: Making sense of the conflict between shell and greenpeace in the north sea. *Organization*, 6(3):499–528, 1999.

[46] Giridhari Venkatadri, Elena Lucherini, Piotr Sapiezynski, and Alan Mislove. Investigating sources of pii used in facebook's targeted advertising. *Proc. Priv. Enhancing Technol.*, 2019(1):227–244, 2019.

[47] Iain Weaver, Ned Westwood, Travis Coan, Saffron O'Neill, and Hywel TP Williams. Sponsored messaging about climate change on facebook: Actors, content, frames. *arXiv preprint arXiv:2211.13965*, 2022.

[48] Jessica Wentz, Delta Merner, Benjamin Franta, Alessandra Lehmen, and Peter C Frumhoff. Research priorities for climate litigation. *Earth's Future*, page e2022EF002928, 2023.

[49] Jack Zhou. Boomerangs versus javelins: How polarization constrains communication on climate change. *Environmental Politics*, 25(5):788–811, 2016.

6 Appendix

6.1 Metadata

The metadata that is most relevant to our analysis and work are the following:

- `ad_reached_countries` - Facebook delivered the ads in these countries. We use this attribute to filter advertisements that were only shown in the United States.
- `delivery_by_region` - A state-wise breakdown of the ad delivery percentage.
- `demographic_distribution` - A gender and age wise breakdown of the ad delivery percentage.
- `impressions` - A range representing the minimum/maximum number of non-unique Facebook accounts that were shown an ad. The smallest bin represents ads that were shown to between 0 - 999 Facebook accounts and the largest bin contains ads that were shown to $>$ 1M Facebook accounts.
- `spend` - A range representing the minimum/maximum amount that was spent on an ad. The smallest bin represents ads whose expenditure was between \$0 - \$100 and the largest bin represents ads whose expenditure was $>$\$1M.

6.2 Ad Campaign Attributes

We briefly describe the attributes that were used for our ad campaigns.

- *Duration* - The 652 ad campaigns were run in 3 batches, such that each batch was run for a period of 24h in order to reach all timezones of the U.S.[1]
- *Ad media* We use images of oil rigs, solar cells or controls (Fig. **??**). Each image was modified with the logo of a contrarian or advocacy organization, depending on the treatment group it was assigned to.
- *Ad text* For each ad, we included the text, "Use our website to tell us what you think about this picture."
- *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 micro-targeting features.
- *Ad placement* - We specified that the ads could only be shown on the Facebook platform, and could only be situated on a user account's Facebook feed.
- *Ad budget* - We specified a daily ad budget of \$1/day.
- *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.
- *Ad type* - We ran the ads under the 'Social issues, elections or political issues' category, in accordance with Facebook's advertising guidelines.

6.3 Experiment: Sampling control images

To sample control images, we utilize the ImageNet-21k dataset [34] and the WordNet [18] hierarchy. The ImageNet-21k dataset contains images grouped under 21,841 classes; WordNet is a large lexical database of English. In WordNet, nouns, verbs, adjectives and adverbs are grouped into sets of cognitive synonyms (synsets)[18], each expressing a distinct concept. The 21,841 labels in the ImageNet-21k dataset are a direct mapping to the noun synsets in Wordnet. We devise a methodology

[1] The 652 ads were run in 3 batches since Facebook has an upper limit of 250 concurrent ads that can be run by an advertiser whose advertising budget is less than \$1,000,000/month. Batch 1 (22 ads per campaign) was run from X to Y on Z. Batch 2 (22 ads per campaign) was run from X to Y on Z. Batch 3 (21 ads per campaign) was run from X to Y on Z. Since the ads are run simultaneously and run for a time period that spans all the timezones in the U.S, we minimize any market effects to the extent possible.

to randomly sample diverse ImageNet categories, such that a sampled category contains at least one image of width and height greater than 600px, which is a criteria required by Facebook's Ad Platform.

ImageNet Labels Tree We begin by constructing the WordNet tree for all the labels (synsets) in the ImageNet-21K dataset. The root of this tree is the synset, "entity"[18], level 1 of this tree contains nodes that are descendants of the "entity" node, level 2 contains descendants of nodes in level 1 and so on.

We then devise a methodology to randomly sample 300 different terminal nodes of this tree, such that these nodes are not related to each other, and the ImageNet category associated with the node contains images of width and height greater than 600px. We found, empirically, that it was necessary to sample roughly 4x the number of images we needed, in order to gather images that satisfied the Facebook Ad Platform's size criteria. To select 65 control ad images, we therefore sampled 300 categories.

In order to gather diverse images, we started at a tree level that has $>$ 300 nodes. Level 6 of the tree is the highest level to have > 300 nodes at 1188. We begin by sampling a random category on level 6 of the WordNet tree. For each category sampled on level 6, we sample a random sub category on the subsequent level, repeating this process until we sample a category that has no descendants. We repeat this process 300 times, to sample 300 unique categories from the 21,841 synsets. From each selected category, we sample a random image having at least 600px width and 600px height to satisfy Facebook criteria for ad images; only 103 categories satisfy this condition. We randomly sample 65 categories from this filtered set to get 65 control images.

We use the *random* library on python for all our sampling needs.

6.4 Detailed Analysis

6.4.1 RQ1: Can observed ad delivery be consistently attributed to the ad image?

We test if the observed ad delivery can be consistently attributed to the ad image, by duplicating the ads featuring images without logos. This includes images of solar cells, oil rigs, and controls without logos. Since statistical tests that compare distributions of categorical variables rely on count values, we compare the 'Reach' of an ad and its copy among various ad destinations (U.S. states, gender, age).

U.S states Audiences in many U.S. states receive 0 views of some ads, and several states receive < 5 views. To satisfy the assumptions of the Fisher's Test, we first group states based on Facebook's population estimates (See table 18). States that are expected to receive close to 0%, 1%, or 2% of the ad are grouped together (and their reach counts are summed), while states expected to receive greater than 2% of the ad are retained as is. This gives us 14 possible state destinations where an ad can be distributed. We use Fisher's Test to compare these observed delivery samples of an ad with its copy. In 89.7% of the ads, we find that the observed delivery sample of an ad is not significantly different from that of its duplicate (N = 195, $p < 0.05$). The exact values from our analysis are present in 4.

Gender Among gender-based destinations, we find that in 100% of the ads, the observed delivery sample of an ad and its duplicate are not significantly different (N = 195, $p < 0.05$). The exact values from our analysis are present in 5.

Age Among age-based destinations, we find that in 99% of the ads, the observed delivery sample of an ad and its duplicate are not significantly different (N = 195, $p < 0.05$). The exact values from our analysis are present in 6.

6.5 Additional Analyses: RQ1, Region axis

When we exclude the images containing logos from the groups considered above, the Kruskal-Wallis test finds that in 35 states there is a statistically significant difference ($p < 0.05$) between the D^R of ads featuring solar cells, oil fields, and controls. Within each state, we then investigate the pairwise differences between the 3 image groups using a Mann Whitney U Test with a Bonferroni correction, and find that in 17 states, there's a statistically significant difference ($p < 0.05$) between the delivery of images showing solar cells and oil fields. In 24 states, there's a statistically significant difference ($p < 0.05$) between the D^R of solar cell images and the control images and in 11 states there's a statistically significant difference ($p < 0.05$) between the D^R of oil field images and the control images.

When we exclude the control images from our omnibus test, and directly investigate if there's a difference between the solar cell images and the oil field images in different states in the U.S, we actually see that in 29 states, where there is a significant difference between the D^R of solar cells and oil fields.

6.6 Ad Delivery and Objects in an Ad Image

H-statistic and associated p-values from the Kruskal-Wallis test investigating if the population medians of ads featuring solar cells and oil rigs with no logo, logo of an advocacy organization and logo of a contrarian organization are significantly different. Results for U.S. State based ad destinations, gender based ad destinations and age based ad destinations are available in tables 1, 2, and 3.

6.7 Ad Delivery Attribution to Ad Image

p-values from Fisher's exact test, comparing the delivery of two ads featuring the same image and running at the same time, for U.S. State based ad destinations, gender based ad destinations and age based ad destinations are available in tables 4, 5 and 6

6.8 Ad Delivery vs Facebook's Population Estimates

6.8.1 Gender

χ^2 statistics and associated p-values for the 3 batches of ads in our experiment show whether observed ad delivery was proportional to Facebook's population estimates. We find that in a majority of cases, the values were not proportional, as shown in tables 7, 8 and 9

6.8.2 Age

χ^2 statistics and associated p-values for the 3 batches of ads in our experiment show whether observed ad delivery was proportional to Facebook's population estimates. We find that in a majority of cases, the values were not proportional, as shown in tables 10, 11 and 12

6.8.3 U.S. States

χ^2 statistics and associated p-values for the 3 batches of ads in our experiment show whether observed ad delivery was proportional to Facebook's population estimates. We find that in a majority of cases, the values were not proportional, as shown in tables 13, 14 and 15

6.8.4 RQ2: Does ad delivery ratio, D, differ when logos are present on an ad image?

We compare D samples of ads featuring images with a contrarian, advocacy, and no logos.

Gender	(H-statistic, p-value)
Female	(324.71, 0.0)
Male	(321.33, 0.0)
Unknown	(36.96, 0.0)

Table 1: H-statistic and p-values from the Kruskal-Wallis H-tests testing the null hypothesis that the population median of ads featuring solar cells, oil rigs, and controls are different.

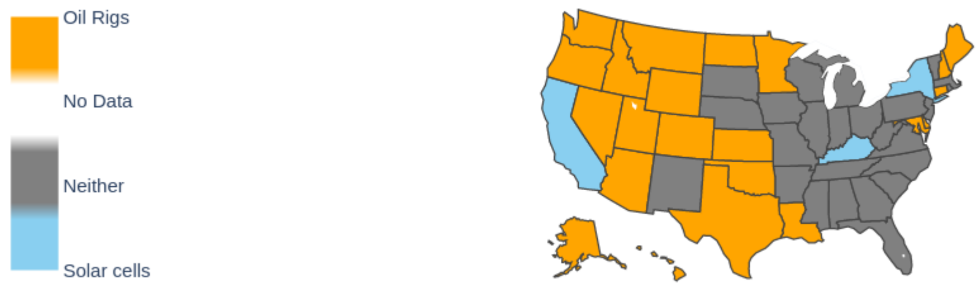

Figure 2: Regions where the D of ads featuring solar cells is significantly different from the D of ads featuring oil rigs. Regions in orange are regions where $D_{\texttt{oil rigs}} > D_{\texttt{solar cells}}$, and regions in blue are regions where $D_{\texttt{solar cells}} > D_{\texttt{oil rigs}}$. There are 25 states where there's a significant difference between the D of ads featuring oil rigs and solar cells. In 88% of these states, the $D_{\texttt{oil rigs}} > D_{\texttt{solar cells}}$

RQ2a: Does D differ based on logos present in an oil rig images? We find that D samples of ads featuring oil rigs with different types of logos (contrarian, advocacy, none) are not significantly different in audiences of different genders, ages or those belonging to different locations.

RQ2b: Does the ad delivery ratio, D differ based on logos present in a solar cell image? We find that D samples of ads featuring oil rigs with different types of logos (contrarian, advocacy, none) are not significantly different in audiences of different genders, ages or those belonging to different locations.

6.8.5 RQ3: Does ad delivery ratio, D, change based on the content of an ad image?

We compare and investigate differences in the D samples of ads featuring solar cells, oil rigs and controls.

U.S. states In 38 states, the D sample of at least one of the three groups (Solar cells, oil rigs and controls) is significantly different from the others (N=650, $p < 0.05$, k=3). Upon investigating the pairwise differences, we find that, in 25 states, there's a significant difference (N = 520, $p < 0.05$) between the D samples of solar cells and oil rigs. In 30 states, there's a significant difference (N=390, $p < 0.05$) between the D samples of solar cells and controls and in 18 states there's a significant difference (N=390, $p < 0.05$) between the D samples of oil rigs and controls. See 2 for a map visualizing these states, and table 3 for the Kruskall-Wallis H Statistic and p values.

Gender In both male and female audiences, the D samples of the three groups (solar cells, controls, or oil rigs) are significantly different(N=650, $p < 0.05$, k=2) as seen in Fig. **??**. Upon investigating the pairwise differences, we find that in both male and female audiences, the D samples of oil rigs

Age	(H-statistic, p-value)
18-24	(113.96, 0.0)
25-34	(73.85, 0.0)
35-44	(7.59, 0.02)
45-54	(6.29, 0.04)
55-64	(19.34, 0.0)
65+	(48.68, 0.0)

Table 2: H-statistic and p-values from the Kruskal-Wallis H-tests testing the null hypothesis that the population median of ads featuring solar cells, oil rigs, and controls are different.

and solar cells are significantly different (N=520, $p < 0.05$, k=2), and of oil rigs and controls (N=390, $p < 0.05$, k=2) are significantly different. Further, ads featuring oil rigs are preferentially delivered to males while ads featuring solar cells are preferentially delivered to females. The D samples further reveal that male populations are over represented ($D > 1$) while female populations are under represented in the ad audiences selected by Facebook ($D < 1$). See table 1 for the Kruskall-Wallis H Statistic and p values.

Age The D samples of ads featuring solar cells, controls, or oil rigs (Fig **??**) are significantly different ($p < 0.05$, N=650, k=3) in audiences belonging to all age groups (18-24, 25-34, 35-44, 45-54, 55-64, 65+). Upon investigating the pairwise differences, we find that except for audiences in the ages of 45-54, the D samples of solar cells and oil rigs are significantly different in all age groups ($p < 0.05$, N=520, k=2). Ads featuring oil rigs are preferentially delivered to older audiences while ads featuring solar cells are preferentially delivered to younger audiences. See table 2 for the Kruskall-Wallis H Statistic and p values. Additional sub-analyses can be found in Appendix 6.5

6.8.6 RQ4: Is the observed ad delivery proportional to Facebook's population estimates within U.S. state, age, and gender-based ad destinations?

U.S. states In 47% of all ads, the observed delivery matches Facebook's population estimates(N=650, $p < 0.05$; See table 18 for population estimates by state.). The observed delivery of 64% of controls (N=130, $p < 0.05$), and 42.5% (N=520, $p < 0.05$) of non-control images (solar cells or oil rigs) matches Facebook's population estimates. The exact H-statistics and p-values are provided in tables 13, 14, and 15

Gender The observed delivery of 28% of non-control images (oil rigs or solar cells, N = 520, $p < 0.05$) and 54% of control images (N=130, $p < 0.05$) matches Facebook's population estimates (See table 16 for population estimates by gender). The exact H-statistics and p-values are provided in tables 7, 8, and 9

Age The observed delivery of none of the non-control images (N=520, $p < 0.05$) and none of the control images (N=130, $p < 0.05$) matches Facebook's population estimates (See table 17 for population estimates by age). The exact H-statistics and p-values are provided in tables 10, 11, and 12

6.9 Facebook Estimated Audience Size Estimates

Facebook's estimated audience size estimates are available in tables 16, 17, and 18.

U.S. State	H-statistic, p-value
Alabama	(10.38, 0.01)
Alaska	(20.63, 0.0)
Arizona	(41.19, 0.0)
Arkansas	(0.7, 0.71)
California	(13.54, 0.0)
Colorado	(41.2, 0.0)
Connecticut	(31.24, 0.0)
Delaware	(21.66, 0.0)
Florida	(3.23, 0.2)
Georgia	(1.01, 0.6)
Hawaii	(13.36, 0.0)
Idaho	(10.49, 0.01)
Illinois	(9.74, 0.01)
Indiana	(4.66, 0.1)
Iowa	(10.03, 0.01)
Kansas	(8.49, 0.01)
Kentucky	(14.16, 0.0)
Louisiana	(8.67, 0.01)
Maine	(27.79, 0.0)
Maryland	(31.25, 0.0)
Massachusetts	(43.94, 0.0)
Michigan	(4.73, 0.09)
Minnesota	(13.15, 0.0)
Mississippi	(4.1, 0.13)
Missouri	(1.42, 0.49)
Montana	(18.5, 0.0)
Nebraska	(7.93, 0.02)
Nevada	(29.7, 0.0)
New Hampshire	(29.08, 0.0)
New Jersey	(13.19, 0.0)
New Mexico	(13.23, 0.0)
New York	(20.47, 0.0)
North Carolina	(2.52, 0.28)
North Dakota	(26.96, 0.0)
Ohio	(12.55, 0.0)
Oklahoma	(32.05, 0.0)
Oregon	(22.74, 0.0)
Pennsylvania	(11.89, 0.0)
Rhode Island	(11.39, 0.0)
South Carolina	(2.54, 0.28)
South Dakota	(10.34, 0.01)
Tennessee	(5.81, 0.05)
Texas	(42.88, 0.0)
Utah	(20.56, 0.0)
Vermont	(21.09, 0.0)
Virginia	(5.12, 0.08)
Washington	(15.18, 0.0)
West Virginia	(1.08, 0.58)
Wisconsin	(7.83, 0.02)
Wyoming	(16.69, 0.0)

Table 3: H-statistic and p-values from the Kruskal-Wallis H-tests testing the null hypothesis that the population median of ads featuring solar cells, oil rigs, and controls are different.

	Solar Cells			Oil rigs			Controls		
Ad ID	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3
1	0.12	0.56	0.53	0.01	1.0	0.34	1.0	0.73	0.55
2	1.0	0.06	1.0	0.49	0.01	1.0	0.53	1.0	0.15
3	1.0	1.0	0.54	0.01	0.77	0.37	0.56	0.19	0.28
4	0.5	1.0	0.23	0.46	0.36	0.69	0.11	0.63	0.26
5	1.0	0.67	0.52	0.63	0.56	1.0	0.65	0.02	0.23
6	1.0	1.0	0.44	0.69	0.13	0.48	0.41	0.02	0.76
7	0.6	1.0	0.31	1.0	1.0	0.14	0.39	0.76	0.34
8	1.0	0.17	0.42	1.0	1.0	0.9	1.0	0.27	1.0
9	1.0	0.36	0.53	0.03	0.73	1.0	0.14	0.07	0.2
10	1.0	0.92	0.22	1.0	0.34	0.04	0.06	1.0	1.0
11	0.59	1.0	0.09	1.0	0.77	0.06	0.63	0.58	0.02
12	0.25	0.17	1.0	0.53	0.0	0.2	0.54	0.74	1.0
13	0.1	1.0	0.26	0.64	1.0	0.85	0.42	0.71	0.75
14	0.3	0.03	0.03	0.54	0.33	0.89	0.54	0.06	1.0
15	0.38	1.0	0.73	0.25	0.14	0.9	0.12	0.11	0.12
16	0.08	0.48	0.0	0.73	0.0	1.0	0.48	0.44	0.91
17	1.0	1.0	0.04	0.35	1.0	0.1	0.06	0.57	0.35
18	0.04	1.0	0.09	0.04	0.69	1.0	0.76	0.02	1.0
19	0.49	1.0	0.03	0.27	0.45	0.08	0.55	0.51	0.5
20	0.12	0.64	1.0	1.0	0.0	0.48	0.81	1.0	0.71
21	1.0	0.55	1.0	0.51	0.0	0.5	1.0	0.59	0.32
22	0.46	0.06		0.51	0.08		0.14	1.0	

Table 4: Table showing p-values for the Two Sided Fisher's Exact Test. The test measures if observed ad delivery in different U.S. state based ad destinations was consistent between 2 ads featuring the same image and run at the same time. p-values were calculated using the exact test, without using Monte-carlo simulations. Confidence intervals, and an estimate of the odds ratio are not available since this is a 14x2 dataset. Delivery of 65 associated solar cell, oil rig and control ad pairs split into 3 batches were compared.

	Solar Cells			Oil rigs			Controls		
Ad ID	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3
1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
2	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
3	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
4	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
6	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
7	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
10	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
12	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
13	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
14	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
15	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
16	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
17	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
18	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
19	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
20	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
21	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
22	1.0	1.0		1.0	1.0		1.0	1.0	

Table 5: Table showing p-values for the Two Sided Fisher's Exact Test. The test measures if observed ad delivery in different gender based ad destinations was consistent between 2 ads featuring the same image and run at the same time. p-values were calculated using the exact test, without using Monte-carlo simulations. Confidence intervals, and an estimate of the odds ratio are not available since this is a 3x2 dataset. Delivery of 65 associated solar cell, oil rig and control ad pairs split into 3 batches were compared.

	Solar Cells			Oil rigs			Controls		
Ad ID	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3	Batch 1	Batch 2	Batch 3
1	1.0	1.0	1.0	1.0	1.0	1.0	0.2	1.0	1.0
2	0.07	1.0	1.0	1.0	1.0	0.5	0.07	1.0	0.5
3	1.0	1.0	1.0	1.0	0.4	0.2	1.0	1.0	0.2
4	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.5	0.07
5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.07	0.05
6	1.0	0.2	1.0	0.2	1.0	1.0	1.0	1.0	0.13
7	0.2	1.0	1.0	1.0	1.0	1.0	1.0	0.4	1.0
8	1.0	1.0	1.0	1.0	1.0	0.5	1.0	0.07	1.0
9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
10	1.0	0.4	1.0	0.5	1.0	1.0	0.13	1.0	1.0
11	1.0	1.0	1.0	1.0	1.0	0.07	1.0	1.0	0.4
12	1.0	1.0	1.0	1.0	1.0	0.4	1.0	0.07	1.0
13	1.0	0.2	0.2	0.2	0.4	1.0	1.0	0.2	1.0
14	1.0	1.0	1.0	0.07	0.5	1.0	1.0	0.4	1.0
15	1.0	1.0	1.0	1.0	0.2	1.0	1.0	1.0	0.2
16	1.0	1.0	1.0	1.0	1.0	1.0	0.07	0.2	1.0
17	1.0	0.4	1.0	1.0	1.0	1.0	1.0	0.2	1.0
18	1.0	1.0	1.0	1.0	1.0	0.02	0.4	0.2	0.07
19	0.2	1.0	1.0	1.0	0.13	1.0	1.0	1.0	1.0
20	1.0	1.0	1.0	0.07	1.0	1.0	1.0	1.0	1.0
21	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
22	1.0	1.0		1.0	1.0		1.0	1.0	

Table 6: Table showing p-values for the Two Sided Fisher's Exact Test. The test measures if observed ad delivery in different age based ad destinations was consistent between 2 ads featuring the same image and run at the same time. p-values were calculated using the exact test, without using Monte-carlo simulations. Confidence intervals, and an estimate of the odds ratio are not available since this is a 6x2 dataset. Delivery of 65 associated solar cell, oil rig and control ad pairs split into 3 batches were compared.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(16.12, 0.0)	(9.14, 0.01)	(25.14, 0.0)	(23.77, 0.0)	(50.83, 0.0)	(88.08, 0.0)	(40.5, 0.0)	(39.07, 0.0)	(75.92, 0.0)	(87.77, 0.0)	2.0
2	(16.26, 0.0)	(3.78, 0.15)	(23.31, 0.0)	(38.66, 0.0)	(51.56, 0.0)	(58.89, 0.0)	(36.93, 0.0)	(44.47, 0.0)	(41.48, 0.0)	(41.88, 0.0)	2.0
3	(52.66, 0.0)	(30.63, 0.0)	(22.8, 0.0)	(2.86, 0.24)	(47.2, 0.0)	(42.7, 0.0)	(45.58, 0.0)	(3.24, 0.2)	(65.37, 0.0)	(65.76, 0.0)	2.0
4	(10.44, 0.01)	(2.9, 0.23)	(30.92, 0.0)	(16.03, 0.0)	(50.83, 0.0)	(42.58, 0.0)	(45.95, 0.0)	(74.85, 0.0)	(53.74, 0.0)	(79.24, 0.0)	2.0
5	(1.19, 0.55)	(49.32, 0.0)	(42.23, 0.0)	(53.52, 0.0)	(32.22, 0.0)	(43.06, 0.0)	(65.23, 0.0)	(53.24, 0.0)	(30.53, 0.0)	(56.01, 0.0)	2.0
6	(4.16, 0.12)	(12.13, 0.0)	(28.86, 0.0)	(30.81, 0.0)	(37.3, 0.0)	(62.94, 0.0)	(40.67, 0.0)	(60.41, 0.0)	(1.53, 0.47)	(3.64, 0.16)	2.0
7	(10.62, 0.0)	(6.41, 0.04)	(0.72, 0.7)	(0.74, 0.69)	(54.86, 0.0)	(67.38, 0.0)	(41.88, 0.0)	(88.93, 0.0)	(7.36, 0.03)	(10.93, 0.0)	2.0
8	(7.4, 0.02)	(3.78, 0.15)	(18.41, 0.0)	(2.33, 0.31)	(49.35, 0.0)	(49.09, 0.0)	(33.42, 0.0)	(56.85, 0.0)	(3.99, 0.14)	(6.35, 0.04)	2.0
9	(15.09, 0.0)	(7.28, 0.03)	(39.72, 0.0)	(9.25, 0.01)	(39.34, 0.0)	(30.29, 0.0)	(52.75, 0.0)	(56.22, 0.0)	(15.46, 0.0)	(16.45, 0.0)	2.0
10	(1.11, 0.57)	(6.9, 0.03)	(14.66, 0.0)	(23.27, 0.0)	(45.81, 0.0)	(43.08, 0.0)	(64.21, 0.0)	(22.19, 0.0)	(4.92, 0.09)	(0.73, 0.69)	2.0
11	(2.87, 0.24)	(2.33, 0.31)	(44.44, 0.0)	(8.53, 0.01)	(76.58, 0.0)	(71.44, 0.0)	(56.23, 0.0)	(72.08, 0.0)	(0.87, 0.65)	(1.39, 0.5)	2.0
12	(28.08, 0.0)	(25.09, 0.0)	(19.99, 0.0)	(5.91, 0.05)	(55.93, 0.0)	(20.08, 0.0)	(36.31, 0.0)	(74.86, 0.0)	(82.86, 0.0)	(58.39, 0.0)	2.0
13	(14.94, 0.0)	(24.7, 0.0)	(26.21, 0.0)	(35.05, 0.0)	(46.47, 0.0)	(51.3, 0.0)	(77.6, 0.0)	(81.78, 0.0)	(2.65, 0.27)	(2.21, 0.33)	2.0
14	(22.26, 0.0)	(1.62, 0.45)	(26.61, 0.0)	(3.17, 0.21)	(16.75, 0.0)	(28.12, 0.0)	(25.21, 0.0)	(21.38, 0.0)	(6.75, 0.03)	(5.92, 0.05)	2.0
15	(8.83, 0.01)	(11.26, 0.0)	(11.46, 0.0)	(4.67, 0.1)	(35.61, 0.0)	(42.25, 0.0)	(29.25, 0.0)	(48.21, 0.0)	(2.74, 0.25)	(7.94, 0.02)	2.0
16	(15.65, 0.0)	(15.55, 0.0)	(22.97, 0.0)	(17.2, 0.0)	(21.59, 0.0)	(10.91, 0.0)	(28.34, 0.0)	(72.67, 0.0)	(8.19, 0.02)	(2.16, 0.34)	2.0
17	(22.58, 0.0)	(48.87, 0.0)	(32.43, 0.0)	(23.99, 0.0)	(53.51, 0.0)	(46.78, 0.0)	(59.3, 0.0)	(56.97, 0.0)	(3.21, 0.2)	(0.75, 0.69)	2.0
18	(12.19, 0.0)	(6.63, 0.04)	(21.24, 0.0)	(15.05, 0.0)	(39.62, 0.0)	(60.54, 0.0)	(40.06, 0.0)	(66.24, 0.0)	(32.08, 0.0)	(28.26, 0.0)	2.0
19	(2.34, 0.31)	(7.41, 0.02)	(1.5, 0.47)	(1.68, 0.43)	(50.19, 0.0)	(39.69, 0.0)	(46.12, 0.0)	(86.43, 0.0)	(11.26, 0.0)	(4.6, 0.1)	2.0
20	(5.01, 0.08)	(0.16, 0.92)	(15.55, 0.0)	(9.31, 0.01)	(71.32, 0.0)	(65.09, 0.0)	(58.39, 0.0)	(69.39, 0.0)	(70.2, 0.0)	(78.02, 0.0)	2.0
21	(29.55, 0.0)	(16.22, 0.0)	(35.49, 0.0)	(8.66, 0.01)	(89.27, 0.0)	(63.42, 0.0)	(31.57, 0.0)	(61.9, 0.0)	(0.75, 0.69)	(1.08, 0.58)	2.0
22	(1.76, 0.41)	(2.67, 0.26)	(6.02, 0.05)	(0.81, 0.67)	(53.2, 0.0)	(64.18, 0.0)	(63.98, 0.0)	(93.0, 0.0)	(23.03, 0.0)	(30.34, 0.0)	2.0

Table 7: Batch1: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different gender based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(13.15, 0.0)	(13.54, 0.0)	(19.51, 0.0)	(19.35, 0.0)	(43.08, 0.0)	(48.36, 0.0)	(36.02, 0.0)	(38.27, 0.0)	(23.12, 0.0)	(16.98, 0.0)	2.0
2	(6.36, 0.04)	(8.66, 0.01)	(13.02, 0.0)	(27.35, 0.0)	(41.96, 0.0)	(50.33, 0.0)	(20.44, 0.0)	(44.13, 0.0)	(43.32, 0.0)	(48.34, 0.0)	2.0
3	(0.1, 0.95)	(1.88, 0.39)	(12.98, 0.0)	(6.58, 0.04)	(28.94, 0.0)	(44.72, 0.0)	(34.86, 0.0)	(31.57, 0.0)	(4.15, 0.13)	(1.11, 0.58)	2.0
4	(4.15, 0.13)	(0.93, 0.63)	(7.71, 0.02)	(5.04, 0.08)	(35.49, 0.0)	(29.24, 0.0)	(39.93, 0.0)	(42.25, 0.0)	(3.67, 0.16)	(2.52, 0.28)	2.0
5	(1.8, 0.41)	(0.8, 0.67)	(5.29, 0.07)	(12.08, 0.0)	(45.68, 0.0)	(30.06, 0.0)	(29.74, 0.0)	(43.96, 0.0)	(13.22, 0.0)	(0.7, 0.71)	2.0
6	(5.02, 0.08)	(27.25, 0.0)	(15.24, 0.0)	(32.66, 0.0)	(46.2, 0.0)	(42.89, 0.0)	(29.24, 0.0)	(36.31, 0.0)	(15.63, 0.0)	(25.89, 0.0)	2.0
7	(0.62, 0.73)	(1.45, 0.48)	(1.24, 0.54)	(1.21, 0.55)	(46.9, 0.0)	(37.57, 0.0)	(32.17, 0.0)	(44.45, 0.0)	(5.46, 0.07)	(12.39, 0.0)	2.0
8	(11.68, 0.0)	(15.83, 0.0)	(21.01, 0.0)	(21.34, 0.0)	(19.35, 0.0)	(22.34, 0.0)	(29.8, 0.0)	(26.61, 0.0)	(21.14, 0.0)	(10.0, 0.01)	2.0
9	(7.75, 0.02)	(17.76, 0.0)	(16.68, 0.0)	(10.8, 0.0)	(32.24, 0.0)	(19.59, 0.0)	(42.23, 0.0)	(25.35, 0.0)	(0.62, 0.73)	(1.14, 0.57)	2.0
10	(8.57, 0.01)	(9.78, 0.01)	(6.82, 0.03)	(13.52, 0.0)	(44.89, 0.0)	(46.08, 0.0)	(45.04, 0.0)	(21.24, 0.0)	(22.64, 0.0)	(21.59, 0.0)	2.0
11	(23.07, 0.0)	(7.83, 0.02)	(25.72, 0.0)	(28.07, 0.0)	(32.81, 0.0)	(50.33, 0.0)	(31.83, 0.0)	(52.46, 0.0)	(42.25, 0.0)	(38.69, 0.0)	2.0
12	(0.8, 0.67)	(3.67, 0.16)	(4.64, 0.1)	(12.75, 0.0)	(68.25, 0.0)	(37.73, 0.0)	(61.95, 0.0)	(33.42, 0.0)	(1.07, 0.59)	(0.87, 0.65)	2.0
13	(23.55, 0.0)	(24.22, 0.0)	(19.3, 0.0)	(5.92, 0.05)	(52.15, 0.0)	(61.95, 0.0)	(51.31, 0.0)	(36.31, 0.0)	(0.71, 0.7)	(2.53, 0.28)	2.0
14	(20.27, 0.0)	(32.3, 0.0)	(30.63, 0.0)	(20.02, 0.0)	(40.1, 0.0)	(21.37, 0.0)	(6.7, 0.04)	(37.37, 0.0)	(12.74, 0.0)	(20.28, 0.0)	2.0
15	(3.43, 0.18)	(6.47, 0.04)	(4.18, 0.12)	(17.1, 0.0)	(45.04, 0.0)	(66.64, 0.0)	(55.17, 0.0)	(56.27, 0.0)	(62.28, 0.0)	(30.07, 0.0)	2.0
16	(6.52, 0.04)	(5.03, 0.08)	(8.97, 0.01)	(2.23, 0.33)	(53.74, 0.0)	(16.32, 0.0)	(35.9, 0.0)	(38.07, 0.0)	(11.8, 0.0)	(4.98, 0.08)	2.0
17	(15.63, 0.0)	(2.39, 0.3)	(20.27, 0.0)	(22.73, 0.0)	(69.74, 0.0)	(32.17, 0.0)	(33.42, 0.0)	(53.08, 0.0)	(2.95, 0.23)	(1.91, 0.39)	2.0
18	(20.76, 0.0)	(28.73, 0.0)	(1.18, 0.55)	(5.63, 0.06)	(51.64, 0.0)	(35.57, 0.0)	(46.63, 0.0)	(54.0, 0.0)	(17.65, 0.0)	(23.36, 0.0)	2.0
19	(15.91, 0.0)	(21.0, 0.0)	(15.65, 0.0)	(14.85, 0.0)	(18.43, 0.0)	(14.12, 0.0)	(20.72, 0.0)	(43.73, 0.0)	(5.02, 0.08)	(16.43, 0.0)	2.0
20	(7.69, 0.02)	(5.59, 0.06)	(4.6, 0.1)	(4.81, 0.09)	(27.4, 0.0)	(62.71, 0.0)	(22.96, 0.0)	(25.22, 0.0)	(37.73, 0.0)	(57.77, 0.0)	2.0
21	(5.0, 0.08)	(13.87, 0.0)	(2.66, 0.26)	(16.76, 0.0)	(72.81, 0.0)	(44.36, 0.0)	(50.15, 0.0)	(40.06, 0.0)	(10.89, 0.0)	(12.07, 0.0)	2.0
22	(13.27, 0.0)	(4.12, 0.13)	(2.05, 0.36)	(11.01, 0.0)	(31.11, 0.0)	(32.02, 0.0)	(17.05, 0.0)	(50.76, 0.0)	(25.42, 0.0)	(31.7, 0.0)	2.0

Table 8: Batch2: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different gender based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(15.11, 0.0)	(8.34, 0.02)	(5.68, 0.06)	(4.1, 0.13)	(29.38, 0.0)	(41.0, 0.0)	(42.93, 0.0)	(18.44, 0.0)	(2.41, 0.3)	(0.92, 0.63)	2.0
2	(2.01, 0.37)	(2.02, 0.36)	(21.01, 0.0)	(18.57, 0.0)	(26.88, 0.0)	(15.92, 0.0)	(27.84, 0.0)	(50.68, 0.0)	(26.17, 0.0)	(18.83, 0.0)	2.0
3	(5.55, 0.06)	(12.94, 0.0)	(13.95, 0.0)	(7.4, 0.02)	(20.01, 0.0)	(18.17, 0.0)	(52.4, 0.0)	(39.28, 0.0)	(2.21, 0.33)	(8.07, 0.02)	2.0
4	(15.85, 0.0)	(18.57, 0.0)	(0.65, 0.72)	(10.15, 0.01)	(38.34, 0.0)	(12.01, 0.0)	(20.09, 0.0)	(47.73, 0.0)	(2.43, 0.3)	(0.65, 0.72)	2.0
5	(17.18, 0.0)	(13.51, 0.0)	(10.62, 0.0)	(2.85, 0.24)	(39.81, 0.0)	(37.08, 0.0)	(34.45, 0.0)	(27.29, 0.0)	(4.12, 0.13)	(10.49, 0.01)	2.0
6	(11.92, 0.0)	(5.46, 0.07)	(1.18, 0.55)	(2.02, 0.36)	(23.58, 0.0)	(19.89, 0.0)	(37.07, 0.0)	(47.22, 0.0)	(0.57, 0.75)	(0.56, 0.76)	2.0
7	(23.16, 0.0)	(13.4, 0.0)	(10.06, 0.01)	(15.85, 0.0)	(47.08, 0.0)	(40.41, 0.0)	(37.42, 0.0)	(29.76, 0.0)	(33.42, 0.0)	(76.46, 0.0)	2.0
8	(22.49, 0.0)	(18.74, 0.0)	(16.61, 0.0)	(9.19, 0.01)	(20.42, 0.0)	(27.88, 0.0)	(32.22, 0.0)	(28.91, 0.0)	(19.93, 0.0)	(7.18, 0.03)	2.0
9	(16.6, 0.0)	(25.17, 0.0)	(0.47, 0.79)	(2.94, 0.23)	(19.87, 0.0)	(53.24, 0.0)	(26.0, 0.0)	(25.1, 0.0)	(96.59, 0.0)	(81.96, 0.0)	2.0
10	(13.96, 0.0)	(3.0, 0.22)	(20.16, 0.0)	(7.28, 0.03)	(40.41, 0.0)	(37.58, 0.0)	(37.48, 0.0)	(42.85, 0.0)	(32.39, 0.0)	(31.47, 0.0)	2.0
11	(18.34, 0.0)	(0.11, 0.95)	(15.45, 0.0)	(12.82, 0.0)	(12.98, 0.0)	(47.29, 0.0)	(24.14, 0.0)	(33.38, 0.0)	(2.13, 0.34)	(0.53, 0.77)	2.0
12	(12.82, 0.0)	(17.18, 0.0)	(3.22, 0.2)	(8.05, 0.02)	(13.96, 0.0)	(17.61, 0.0)	(31.9, 0.0)	(36.31, 0.0)	(8.49, 0.01)	(3.65, 0.16)	2.0
13	(0.24, 0.89)	(1.25, 0.53)	(8.18, 0.02)	(2.96, 0.23)	(32.48, 0.0)	(38.63, 0.0)	(52.15, 0.0)	(42.56, 0.0)	(2.66, 0.26)	(1.07, 0.59)	2.0
14	(7.8, 0.02)	(6.87, 0.03)	(1.63, 0.44)	(3.64, 0.16)	(41.26, 0.0)	(36.21, 0.0)	(35.72, 0.0)	(20.82, 0.0)	(2.05, 0.36)	(0.73, 0.69)	2.0
15	(10.7, 0.0)	(5.47, 0.06)	(28.86, 0.0)	(6.97, 0.03)	(23.88, 0.0)	(23.17, 0.0)	(39.34, 0.0)	(15.25, 0.0)	(8.76, 0.01)	(16.17, 0.0)	2.0
16	(10.62, 0.0)	(11.17, 0.0)	(24.14, 0.0)	(15.55, 0.0)	(49.94, 0.0)	(30.71, 0.0)	(26.25, 0.0)	(23.72, 0.0)	(1.79, 0.41)	(2.11, 0.35)	2.0
17	(3.77, 0.15)	(3.02, 0.22)	(18.22, 0.0)	(2.45, 0.29)	(29.88, 0.0)	(27.83, 0.0)	(37.54, 0.0)	(30.83, 0.0)	(3.54, 0.17)	(9.17, 0.01)	2.0
18	(11.64, 0.0)	(7.98, 0.02)	(18.26, 0.0)	(4.55, 0.1)	(31.57, 0.0)	(45.95, 0.0)	(28.75, 0.0)	(17.71, 0.0)	(42.25, 0.0)	(37.21, 0.0)	2.0
19	(12.48, 0.0)	(8.18, 0.02)	(1.8, 0.41)	(11.73, 0.0)	(30.6, 0.0)	(43.29, 0.0)	(39.0, 0.0)	(18.53, 0.0)	(1.63, 0.44)	(1.08, 0.58)	2.0
20	(3.42, 0.18)	(5.09, 0.08)	(3.25, 0.2)	(8.02, 0.02)	(59.64, 0.0)	(40.24, 0.0)	(51.62, 0.0)	(32.48, 0.0)	(42.12, 0.0)	(52.65, 0.0)	2.0
21	(7.76, 0.02)	(15.83, 0.0)	(8.25, 0.02)	(20.84, 0.0)	(26.87, 0.0)	(42.25, 0.0)	(33.14, 0.0)	(62.08, 0.0)	(32.65, 0.0)	(44.65, 0.0)	2.0

Table 9: Batch3: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different gender based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(302.51, 0.0)	(385.41, 0.0)	(369.35, 0.0)	(294.82, 0.0)	(205.49, 0.0)	(233.23, 0.0)	(234.73, 0.0)	(289.0, 0.0)	(463.34, 0.0)	(417.05, 0.0)	5.0
2	(301.61, 0.0)	(282.14, 0.0)	(277.93, 0.0)	(381.29, 0.0)	(339.91, 0.0)	(341.34, 0.0)	(421.03, 0.0)	(362.77, 0.0)	(187.0, 0.0)	(312.64, 0.0)	5.0
3	(166.25, 0.0)	(255.77, 0.0)	(296.22, 0.0)	(296.22, 0.0)	(264.97, 0.0)	(225.35, 0.0)	(270.68, 0.0)	(353.26, 0.0)	(355.09, 0.0)	(381.51, 0.0)	5.0
4	(364.75, 0.0)	(426.58, 0.0)	(394.32, 0.0)	(321.63, 0.0)	(412.74, 0.0)	(290.78, 0.0)	(357.23, 0.0)	(388.61, 0.0)	(120.63, 0.0)	(135.6, 0.0)	5.0
5	(425.98, 0.0)	(243.8, 0.0)	(161.96, 0.0)	(332.67, 0.0)	(375.14, 0.0)	(403.02, 0.0)	(398.8, 0.0)	(558.47, 0.0)	(318.63, 0.0)	(307.82, 0.0)	5.0
6	(304.33, 0.0)	(246.75, 0.0)	(205.39, 0.0)	(286.74, 0.0)	(288.23, 0.0)	(432.48, 0.0)	(431.28, 0.0)	(362.94, 0.0)	(333.75, 0.0)	(322.45, 0.0)	5.0
7	(442.47, 0.0)	(378.23, 0.0)	(333.58, 0.0)	(599.85, 0.0)	(206.75, 0.0)	(265.03, 0.0)	(396.46, 0.0)	(227.15, 0.0)	(350.31, 0.0)	(369.44, 0.0)	5.0
8	(430.26, 0.0)	(360.6, 0.0)	(421.0, 0.0)	(379.51, 0.0)	(319.55, 0.0)	(287.94, 0.0)	(212.48, 0.0)	(247.41, 0.0)	(385.89, 0.0)	(404.44, 0.0)	5.0
9	(245.63, 0.0)	(267.54, 0.0)	(212.99, 0.0)	(449.16, 0.0)	(284.62, 0.0)	(343.23, 0.0)	(370.28, 0.0)	(336.25, 0.0)	(231.4, 0.0)	(276.03, 0.0)	5.0
10	(363.57, 0.0)	(357.38, 0.0)	(279.11, 0.0)	(357.89, 0.0)	(463.32, 0.0)	(482.42, 0.0)	(336.6, 0.0)	(489.62, 0.0)	(293.96, 0.0)	(359.55, 0.0)	5.0
11	(341.33, 0.0)	(315.76, 0.0)	(236.36, 0.0)	(375.15, 0.0)	(475.94, 0.0)	(335.32, 0.0)	(404.95, 0.0)	(475.71, 0.0)	(201.3, 0.0)	(240.13, 0.0)	5.0
12	(144.08, 0.0)	(147.04, 0.0)	(55.71, 0.0)	(268.9, 0.0)	(335.24, 0.0)	(302.92, 0.0)	(231.59, 0.0)	(419.68, 0.0)	(277.81, 0.0)	(207.62, 0.0)	5.0
13	(241.98, 0.0)	(282.95, 0.0)	(201.61, 0.0)	(478.99, 0.0)	(345.43, 0.0)	(392.46, 0.0)	(242.41, 0.0)	(355.85, 0.0)	(318.22, 0.0)	(310.22, 0.0)	5.0
14	(354.13, 0.0)	(399.64, 0.0)	(278.52, 0.0)	(453.76, 0.0)	(280.75, 0.0)	(302.37, 0.0)	(330.82, 0.0)	(328.58, 0.0)	(355.09, 0.0)	(286.91, 0.0)	5.0
15	(404.76, 0.0)	(474.57, 0.0)	(275.91, 0.0)	(460.24, 0.0)	(348.76, 0.0)	(349.85, 0.0)	(387.33, 0.0)	(275.88, 0.0)	(329.08, 0.0)	(308.53, 0.0)	5.0
16	(370.94, 0.0)	(282.51, 0.0)	(269.82, 0.0)	(452.03, 0.0)	(374.04, 0.0)	(339.74, 0.0)	(321.55, 0.0)	(289.84, 0.0)	(420.82, 0.0)	(326.29, 0.0)	5.0
17	(412.24, 0.0)	(292.65, 0.0)	(227.68, 0.0)	(283.6, 0.0)	(363.34, 0.0)	(365.32, 0.0)	(425.82, 0.0)	(418.92, 0.0)	(159.56, 0.0)	(133.34, 0.0)	5.0
18	(357.67, 0.0)	(251.11, 0.0)	(181.95, 0.0)	(288.47, 0.0)	(189.9, 0.0)	(281.19, 0.0)	(161.1, 0.0)	(294.71, 0.0)	(417.72, 0.0)	(488.35, 0.0)	5.0
19	(519.54, 0.0)	(506.76, 0.0)	(298.38, 0.0)	(314.85, 0.0)	(223.94, 0.0)	(243.3, 0.0)	(322.27, 0.0)	(417.26, 0.0)	(163.83, 0.0)	(182.85, 0.0)	5.0
20	(282.22, 0.0)	(332.25, 0.0)	(339.77, 0.0)	(334.1, 0.0)	(294.91, 0.0)	(232.2, 0.0)	(333.32, 0.0)	(409.62, 0.0)	(308.75, 0.0)	(304.12, 0.0)	5.0
21	(330.17, 0.0)	(394.24, 0.0)	(139.97, 0.0)	(348.03, 0.0)	(349.87, 0.0)	(217.6, 0.0)	(331.03, 0.0)	(181.07, 0.0)	(254.78, 0.0)	(295.49, 0.0)	5.0
22	(348.37, 0.0)	(191.77, 0.0)	(226.59, 0.0)	(307.12, 0.0)	(353.62, 0.0)	(258.88, 0.0)	(312.89, 0.0)	(385.93, 0.0)	(105.25, 0.0)	(86.11, 0.0)	5.0

Table 10: Batch1: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different age based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(240.66, 0.0)	(264.51, 0.0)	(350.49, 0.0)	(323.42, 0.0)	(337.91, 0.0)	(555.35, 0.0)	(461.74, 0.0)	(261.57, 0.0)	(360.69, 0.0)	(360.18, 0.0)	5.0
2	(415.28, 0.0)	(400.2, 0.0)	(436.58, 0.0)	(420.91, 0.0)	(369.26, 0.0)	(409.56, 0.0)	(384.11, 0.0)	(400.43, 0.0)	(294.99, 0.0)	(308.44, 0.0)	5.0
3	(404.84, 0.0)	(315.4, 0.0)	(232.59, 0.0)	(306.24, 0.0)	(426.16, 0.0)	(469.0, 0.0)	(319.51, 0.0)	(434.67, 0.0)	(273.88, 0.0)	(323.4, 0.0)	5.0
4	(414.86, 0.0)	(547.21, 0.0)	(328.6, 0.0)	(506.07, 0.0)	(356.52, 0.0)	(356.09, 0.0)	(353.92, 0.0)	(372.42, 0.0)	(374.36, 0.0)	(359.82, 0.0)	5.0
5	(365.52, 0.0)	(419.78, 0.0)	(406.02, 0.0)	(408.53, 0.0)	(488.13, 0.0)	(431.04, 0.0)	(458.54, 0.0)	(403.56, 0.0)	(330.86, 0.0)	(350.85, 0.0)	5.0
6	(489.75, 0.0)	(332.99, 0.0)	(356.26, 0.0)	(324.06, 0.0)	(470.71, 0.0)	(365.23, 0.0)	(422.13, 0.0)	(385.0, 0.0)	(393.81, 0.0)	(285.07, 0.0)	5.0
7	(432.63, 0.0)	(384.93, 0.0)	(315.52, 0.0)	(438.9, 0.0)	(260.67, 0.0)	(412.63, 0.0)	(314.84, 0.0)	(474.65, 0.0)	(329.19, 0.0)	(495.3, 0.0)	5.0
8	(211.03, 0.0)	(332.04, 0.0)	(339.74, 0.0)	(415.89, 0.0)	(354.98, 0.0)	(403.07, 0.0)	(343.7, 0.0)	(487.06, 0.0)	(391.13, 0.0)	(407.68, 0.0)	5.0
9	(314.9, 0.0)	(201.05, 0.0)	(379.61, 0.0)	(412.39, 0.0)	(366.69, 0.0)	(360.67, 0.0)	(312.42, 0.0)	(451.65, 0.0)	(318.41, 0.0)	(399.02, 0.0)	5.0
10	(396.9, 0.0)	(480.14, 0.0)	(552.1, 0.0)	(365.92, 0.0)	(418.11, 0.0)	(407.0, 0.0)	(396.09, 0.0)	(445.16, 0.0)	(364.76, 0.0)	(395.97, 0.0)	5.0
11	(368.89, 0.0)	(471.84, 0.0)	(409.81, 0.0)	(402.1, 0.0)	(377.42, 0.0)	(465.99, 0.0)	(427.14, 0.0)	(442.5, 0.0)	(376.81, 0.0)	(398.4, 0.0)	5.0
12	(416.23, 0.0)	(404.96, 0.0)	(437.94, 0.0)	(519.32, 0.0)	(406.97, 0.0)	(378.15, 0.0)	(434.89, 0.0)	(330.76, 0.0)	(414.12, 0.0)	(557.19, 0.0)	5.0
13	(336.55, 0.0)	(322.36, 0.0)	(412.29, 0.0)	(415.35, 0.0)	(423.83, 0.0)	(415.47, 0.0)	(285.06, 0.0)	(436.22, 0.0)	(346.21, 0.0)	(430.66, 0.0)	5.0
14	(359.34, 0.0)	(324.33, 0.0)	(329.79, 0.0)	(346.03, 0.0)	(340.7, 0.0)	(322.75, 0.0)	(282.52, 0.0)	(334.25, 0.0)	(338.0, 0.0)	(420.29, 0.0)	5.0
15	(295.33, 0.0)	(387.94, 0.0)	(504.13, 0.0)	(394.17, 0.0)	(364.65, 0.0)	(435.43, 0.0)	(309.52, 0.0)	(379.58, 0.0)	(387.64, 0.0)	(390.2, 0.0)	5.0
16	(327.19, 0.0)	(359.88, 0.0)	(425.26, 0.0)	(433.55, 0.0)	(425.44, 0.0)	(471.97, 0.0)	(427.99, 0.0)	(460.84, 0.0)	(358.2, 0.0)	(359.79, 0.0)	5.0
17	(374.68, 0.0)	(369.65, 0.0)	(398.02, 0.0)	(501.06, 0.0)	(609.47, 0.0)	(400.75, 0.0)	(494.88, 0.0)	(440.04, 0.0)	(435.27, 0.0)	(353.39, 0.0)	5.0
18	(319.14, 0.0)	(309.92, 0.0)	(359.35, 0.0)	(519.02, 0.0)	(299.78, 0.0)	(382.94, 0.0)	(391.96, 0.0)	(481.99, 0.0)	(398.34, 0.0)	(384.75, 0.0)	5.0
19	(359.61, 0.0)	(467.86, 0.0)	(346.51, 0.0)	(397.74, 0.0)	(388.51, 0.0)	(501.11, 0.0)	(525.51, 0.0)	(563.1, 0.0)	(324.68, 0.0)	(328.59, 0.0)	5.0
20	(476.38, 0.0)	(397.95, 0.0)	(378.06, 0.0)	(400.8, 0.0)	(441.96, 0.0)	(463.14, 0.0)	(489.42, 0.0)	(496.54, 0.0)	(347.52, 0.0)	(357.86, 0.0)	5.0
21	(340.65, 0.0)	(419.51, 0.0)	(478.39, 0.0)	(444.12, 0.0)	(511.78, 0.0)	(400.6, 0.0)	(341.76, 0.0)	(434.09, 0.0)	(363.68, 0.0)	(314.35, 0.0)	5.0
22	(463.49, 0.0)	(297.57, 0.0)	(407.21, 0.0)	(442.11, 0.0)	(339.03, 0.0)	(432.22, 0.0)	(359.29, 0.0)	(533.67, 0.0)	(342.47, 0.0)	(467.64, 0.0)	5.0

Table 11: Batch2: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different age based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(253.28, 0.0)	(237.93, 0.0)	(257.05, 0.0)	(410.96, 0.0)	(401.91, 0.0)	(360.31, 0.0)	(296.15, 0.0)	(328.25, 0.0)	(372.85, 0.0)	(392.32, 0.0)	5.0
2	(441.44, 0.0)	(375.16, 0.0)	(232.82, 0.0)	(314.09, 0.0)	(412.43, 0.0)	(348.91, 0.0)	(330.51, 0.0)	(441.48, 0.0)	(340.34, 0.0)	(390.5, 0.0)	5.0
3	(251.24, 0.0)	(268.49, 0.0)	(269.96, 0.0)	(298.99, 0.0)	(197.26, 0.0)	(363.22, 0.0)	(239.99, 0.0)	(329.79, 0.0)	(352.6, 0.0)	(379.0, 0.0)	5.0
4	(259.29, 0.0)	(271.41, 0.0)	(252.32, 0.0)	(342.74, 0.0)	(353.05, 0.0)	(330.0, 0.0)	(308.8, 0.0)	(349.5, 0.0)	(413.49, 0.0)	(406.07, 0.0)	5.0
5	(386.55, 0.0)	(380.63, 0.0)	(339.54, 0.0)	(395.81, 0.0)	(350.5, 0.0)	(338.25, 0.0)	(357.79, 0.0)	(334.68, 0.0)	(434.17, 0.0)	(365.28, 0.0)	5.0
6	(203.4, 0.0)	(262.52, 0.0)	(262.56, 0.0)	(273.75, 0.0)	(346.04, 0.0)	(287.88, 0.0)	(369.62, 0.0)	(378.96, 0.0)	(310.31, 0.0)	(339.41, 0.0)	5.0
7	(304.04, 0.0)	(222.57, 0.0)	(215.38, 0.0)	(278.87, 0.0)	(343.1, 0.0)	(390.82, 0.0)	(240.37, 0.0)	(341.34, 0.0)	(204.08, 0.0)	(252.15, 0.0)	5.0
8	(252.87, 0.0)	(312.8, 0.0)	(276.86, 0.0)	(369.99, 0.0)	(348.62, 0.0)	(376.37, 0.0)	(233.06, 0.0)	(346.53, 0.0)	(297.64, 0.0)	(304.85, 0.0)	5.0
9	(277.78, 0.0)	(310.07, 0.0)	(360.29, 0.0)	(362.94, 0.0)	(323.56, 0.0)	(290.05, 0.0)	(291.7, 0.0)	(308.02, 0.0)	(219.66, 0.0)	(247.69, 0.0)	5.0
10	(313.84, 0.0)	(489.51, 0.0)	(247.21, 0.0)	(383.54, 0.0)	(210.4, 0.0)	(321.68, 0.0)	(201.07, 0.0)	(353.41, 0.0)	(250.04, 0.0)	(259.75, 0.0)	5.0
11	(295.57, 0.0)	(367.21, 0.0)	(301.12, 0.0)	(346.91, 0.0)	(318.6, 0.0)	(343.32, 0.0)	(194.6, 0.0)	(293.48, 0.0)	(361.54, 0.0)	(366.6, 0.0)	5.0
12	(282.94, 0.0)	(309.55, 0.0)	(313.07, 0.0)	(401.81, 0.0)	(342.97, 0.0)	(411.61, 0.0)	(252.58, 0.0)	(371.43, 0.0)	(328.99, 0.0)	(294.08, 0.0)	5.0
13	(387.32, 0.0)	(335.41, 0.0)	(347.62, 0.0)	(332.63, 0.0)	(359.76, 0.0)	(323.31, 0.0)	(339.42, 0.0)	(304.9, 0.0)	(389.55, 0.0)	(267.69, 0.0)	5.0
14	(285.8, 0.0)	(361.84, 0.0)	(289.23, 0.0)	(373.78, 0.0)	(299.52, 0.0)	(201.56, 0.0)	(387.61, 0.0)	(254.25, 0.0)	(314.79, 0.0)	(300.0, 0.0)	5.0
15	(343.33, 0.0)	(219.31, 0.0)	(211.65, 0.0)	(333.45, 0.0)	(229.25, 0.0)	(216.21, 0.0)	(221.94, 0.0)	(314.36, 0.0)	(360.07, 0.0)	(263.1, 0.0)	5.0
16	(238.9, 0.0)	(261.31, 0.0)	(220.22, 0.0)	(232.16, 0.0)	(224.09, 0.0)	(313.06, 0.0)	(285.35, 0.0)	(241.99, 0.0)	(312.41, 0.0)	(244.19, 0.0)	5.0
17	(363.2, 0.0)	(302.95, 0.0)	(246.09, 0.0)	(316.5, 0.0)	(277.05, 0.0)	(248.77, 0.0)	(230.11, 0.0)	(308.77, 0.0)	(259.78, 0.0)	(411.23, 0.0)	5.0
18	(452.7, 0.0)	(445.95, 0.0)	(411.74, 0.0)	(397.19, 0.0)	(352.02, 0.0)	(407.03, 0.0)	(298.23, 0.0)	(285.33, 0.0)	(286.37, 0.0)	(299.62, 0.0)	5.0
19	(261.09, 0.0)	(270.92, 0.0)	(288.69, 0.0)	(262.53, 0.0)	(348.25, 0.0)	(354.32, 0.0)	(310.8, 0.0)	(224.69, 0.0)	(362.8, 0.0)	(334.55, 0.0)	5.0
20	(348.0, 0.0)	(472.12, 0.0)	(285.83, 0.0)	(288.63, 0.0)	(380.67, 0.0)	(219.06, 0.0)	(220.82, 0.0)	(345.33, 0.0)	(270.62, 0.0)	(384.15, 0.0)	5.0
21	(352.55, 0.0)	(287.36, 0.0)	(349.64, 0.0)	(350.32, 0.0)	(227.33, 0.0)	(324.19, 0.0)	(336.45, 0.0)	(337.68, 0.0)	(471.93, 0.0)	(318.36, 0.0)	5.0

Table 12: Batch3: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different age based ad destinations is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(104.95, 0.0)	(106.24, 0.0)	(116.97, 0.0)	(76.82, 0.01)	(89.09, 0.0)	(62.19, 0.12)	(93.37, 0.0)	(118.92, 0.0)	(118.72, 0.0)	(77.22, 0.01)	50.0
2	(109.51, 0.0)	(86.17, 0.0)	(78.15, 0.01)	(79.53, 0.0)	(76.79, 0.01)	(90.42, 0.0)	(108.25, 0.0)	(92.04, 0.0)	(66.72, 0.06)	(73.0, 0.02)	50.0
3	(88.01, 0.0)	(112.27, 0.0)	(82.11, 0.0)	(108.85, 0.0)	(75.29, 0.01)	(128.11, 0.0)	(72.49, 0.02)	(94.48, 0.0)	(109.75, 0.0)	(76.67, 0.01)	50.0
4	(163.81, 0.0)	(208.42, 0.0)	(157.19, 0.0)	(91.36, 0.0)	(120.11, 0.0)	(59.43, 0.17)	(59.77, 0.16)	(88.37, 0.0)	(44.52, 0.69)	(50.24, 0.46)	50.0
5	(143.41, 0.0)	(52.59, 0.37)	(51.76, 0.41)	(223.41, 0.0)	(78.73, 0.01)	(111.96, 0.0)	(96.15, 0.0)	(97.94, 0.0)	(112.85, 0.0)	(48.15, 0.55)	50.0
6	(73.44, 0.02)	(106.57, 0.0)	(78.25, 0.01)	(69.79, 0.03)	(56.11, 0.26)	(117.85, 0.0)	(98.64, 0.0)	(63.06, 0.1)	(63.02, 0.1)	(83.72, 0.0)	50.0
7	(127.21, 0.0)	(78.81, 0.01)	(91.82, 0.0)	(126.46, 0.0)	(70.61, 0.03)	(129.3, 0.0)	(63.54, 0.09)	(51.12, 0.43)	(88.53, 0.0)	(69.61, 0.03)	50.0
8	(59.21, 0.17)	(102.69, 0.0)	(92.73, 0.0)	(85.69, 0.0)	(72.79, 0.02)	(129.47, 0.0)	(37.93, 0.89)	(92.84, 0.0)	(100.01, 0.0)	(49.32, 0.5)	50.0
9	(51.3, 0.42)	(95.26, 0.0)	(140.48, 0.0)	(39.67, 0.85)	(80.14, 0.0)	(77.47, 0.01)	(74.27, 0.01)	(69.22, 0.04)	(55.8, 0.27)	(65.28, 0.07)	50.0
10	(76.54, 0.01)	(71.17, 0.03)	(84.64, 0.0)	(76.53, 0.01)	(84.02, 0.0)	(80.62, 0.0)	(64.81, 0.08)	(153.8, 0.0)	(64.28, 0.08)	(66.76, 0.06)	50.0
11	(110.5, 0.0)	(76.3, 0.01)	(77.94, 0.01)	(85.99, 0.0)	(91.35, 0.0)	(75.54, 0.01)	(88.27, 0.0)	(78.17, 0.01)	(88.79, 0.0)	(67.95, 0.05)	50.0
12	(60.0, 0.16)	(91.5, 0.0)	(48.01, 0.55)	(58.38, 0.19)	(59.82, 0.16)	(66.35, 0.06)	(111.38, 0.0)	(100.16, 0.0)	(52.23, 0.39)	(57.04, 0.23)	50.0
13	(54.32, 0.31)	(91.8, 0.0)	(103.74, 0.0)	(111.11, 0.0)	(73.33, 0.02)	(104.07, 0.0)	(62.15, 0.12)	(80.67, 0.0)	(79.55, 0.0)	(126.3, 0.0)	50.0
14	(56.35, 0.25)	(86.02, 0.0)	(113.53, 0.0)	(71.24, 0.03)	(79.7, 0.0)	(106.38, 0.0)	(59.57, 0.17)	(67.45, 0.05)	(55.78, 0.27)	(76.21, 0.01)	50.0
15	(108.33, 0.0)	(157.93, 0.0)	(89.92, 0.0)	(131.43, 0.0)	(87.01, 0.0)	(75.94, 0.01)	(54.7, 0.3)	(50.23, 0.46)	(69.54, 0.04)	(51.05, 0.43)	50.0
16	(91.1, 0.0)	(63.91, 0.09)	(69.22, 0.04)	(144.7, 0.0)	(95.62, 0.0)	(98.78, 0.0)	(98.1, 0.0)	(55.75, 0.27)	(58.98, 0.18)	(64.39, 0.08)	50.0
17	(88.82, 0.0)	(73.08, 0.02)	(107.32, 0.0)	(88.96, 0.0)	(99.38, 0.0)	(100.44, 0.0)	(96.55, 0.0)	(68.06, 0.05)	(70.22, 0.03)	(83.82, 0.0)	50.0
18	(84.16, 0.0)	(57.33, 0.22)	(59.56, 0.17)	(86.1, 0.0)	(63.16, 0.1)	(66.81, 0.06)	(84.33, 0.0)	(71.63, 0.02)	(119.09, 0.0)	(91.66, 0.0)	50.0
19	(128.59, 0.0)	(99.51, 0.0)	(87.01, 0.0)	(86.74, 0.0)	(82.97, 0.0)	(86.93, 0.0)	(94.82, 0.0)	(264.14, 0.0)	(38.93, 0.87)	(69.34, 0.04)	50.0
20	(106.34, 0.0)	(93.06, 0.0)	(76.36, 0.01)	(105.89, 0.0)	(71.71, 0.02)	(56.53, 0.24)	(54.42, 0.31)	(81.16, 0.0)	(59.0, 0.18)	(51.02, 0.43)	50.0
21	(95.32, 0.0)	(83.27, 0.0)	(71.39, 0.03)	(150.18, 0.0)	(117.22, 0.0)	(92.15, 0.0)	(146.39, 0.0)	(63.03, 0.1)	(71.51, 0.02)	(54.35, 0.31)	50.0
22	(188.79, 0.0)	(117.99, 0.0)	(68.71, 0.04)	(64.27, 0.08)	(80.6, 0.0)	(126.53, 0.0)	(100.78, 0.0)	(128.06, 0.0)	(48.36, 0.54)	(77.07, 0.01)	50.0

Table 13: Batch1: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different U.S. states is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook. The Chi Square test was not able to be performed for the highlighted values.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(98.56, 0.0)	(74.01, 0.02)	(50.34, 0.46)	(55.08, 0.29)	(76.01, 0.01)	(172.15, 0.0)	(85.09, 0.0)	(78.18, 0.01)	(71.18, 0.03)	(82.85, 0.0)	50.0
2	(102.4, 0.0)	(80.59, 0.0)	(99.13, 0.0)	(86.78, 0.0)	(121.49, 0.0)	(80.14, 0.0)	(77.16, 0.01)	(83.76, 0.0)	(45.42, 0.66)	(103.1, 0.0)	50.0
3	(80.91, 0.0)	(90.0, 0.0)	(129.87, 0.0)	(43.34, 0.74)	(124.5, 0.0)	(72.14, 0.02)	(95.83, 0.0)	(116.18, 0.0)	(87.31, 0.0)	(89.58, 0.0)	50.0
4	(93.65, 0.0)	(137.15, 0.0)	(60.01, 0.16)	(107.83, 0.0)	(78.88, 0.01)	(53.07, 0.36)	(84.38, 0.0)	(98.94, 0.0)	(61.23, 0.13)	(52.87, 0.36)	50.0
5	(86.26, 0.0)	(80.23, 0.0)	(85.87, 0.0)	(101.13, 0.0)	(99.32, 0.0)	(101.56, 0.0)	(99.09, 0.0)	(122.46, 0.0)	(87.79, 0.0)	(69.13, 0.04)	50.0
6	(112.19, 0.0)	(110.53, 0.0)	(90.26, 0.0)	(109.3, 0.0)	(84.24, 0.0)	(69.09, 0.04)	(84.98, 0.0)	(77.28, 0.01)	(56.45, 0.25)	(59.43, 0.17)	50.0
7	(106.44, 0.0)	(132.26, 0.0)	(95.13, 0.0)	(122.91, 0.0)	(72.3, 0.02)	(74.6, 0.01)	(79.02, 0.01)	(61.13, 0.13)	(90.89, 0.0)	(64.1, 0.09)	50.0
8	(60.02, 0.16)	(60.08, 0.16)	(88.67, 0.0)	(122.13, 0.0)	(67.03, 0.05)	(89.8, 0.0)	(167.53, 0.0)	(93.99, 0.0)	(84.59, 0.0)	(80.59, 0.0)	50.0
9	(74.45, 0.01)	(104.81, 0.0)	(58.04, 0.2)	(82.17, 0.0)	(62.61, 0.11)	(90.72, 0.0)	(54.09, 0.32)	(105.85, 0.0)	(75.41, 0.01)	(79.02, 0.01)	50.0
10	(114.68, 0.0)	(199.3, 0.0)	(179.95, 0.0)	(73.14, 0.02)	(92.24, 0.0)	(112.91, 0.0)	(88.13, 0.0)	(63.02, 0.1)	(75.46, 0.01)	(89.4, 0.0)	50.0
11	(139.0, 0.0)	(99.67, 0.0)	(131.25, 0.0)	(92.24, 0.0)	(91.44, 0.0)	(98.07, 0.0)	(133.64, 0.0)	(81.82, 0.0)	(46.02, 0.63)	(55.47, 0.28)	50.0
12	(136.44, 0.0)	(63.68, 0.09)	(93.33, 0.0)	(92.44, 0.0)	(84.12, 0.0)	(129.34, 0.0)	(92.0, 0.0)	(83.06, 0.0)	(65.37, 0.07)	(51.25, 0.42)	50.0
13	(40.51, 0.83)	(172.55, 0.0)	(95.83, 0.0)	(96.4, 0.0)	(85.92, 0.0)	(85.23, 0.0)	(71.49, 0.02)	(57.52, 0.22)	(76.78, 0.01)	(68.3, 0.04)	50.0
14	(71.67, 0.02)	(130.78, 0.0)	(164.91, 0.0)	(87.88, 0.0)	(72.98, 0.02)	(118.84, 0.0)	(47.32, 0.58)	(74.04, 0.02)	(101.62, 0.0)	(84.82, 0.0)	50.0
15	(96.68, 0.0)	(98.11, 0.0)	(82.54, 0.0)	(69.58, 0.03)	(97.61, 0.0)	(85.45, 0.0)	(116.35, 0.0)	(79.89, 0.0)	(104.76, 0.0)	(31.7, 0.98)	50.0
16	(74.01, 0.02)	(138.75, 0.0)	(98.81, 0.0)	(158.47, 0.0)	(126.15, 0.0)	(152.81, 0.0)	(96.34, 0.0)	(67.31, 0.05)	(103.95, 0.0)	(106.62, 0.0)	50.0
17	(82.5, 0.0)	(128.22, 0.0)	(106.67, 0.0)	(77.7, 0.01)	(109.42, 0.0)	(88.15, 0.0)	(99.85, 0.0)	(114.56, 0.0)	(69.81, 0.03)	(107.25, 0.0)	50.0
18	(76.04, 0.01)	(74.81, 0.01)	(109.8, 0.0)	(111.11, 0.0)	(81.8, 0.0)	(75.17, 0.01)	(108.48, 0.0)	(99.75, 0.0)	(76.01, 0.01)	(49.49, 0.49)	50.0
19	(94.75, 0.0)	(90.15, 0.0)	(144.54, 0.0)	(62.77, 0.11)	(49.68, 0.49)	(96.47, 0.0)	(88.63, 0.0)	(125.28, 0.0)	(82.67, 0.0)	(59.48, 0.17)	50.0
20	(124.37, 0.0)	(63.24, 0.1)	(50.85, 0.44)	(106.72, 0.0)	(80.8, 0.0)	(79.84, 0.0)	(76.14, 0.01)	(66.45, 0.06)	(76.0, 0.01)	(143.05, 0.0)	50.0
21	(109.56, 0.0)	(133.78, 0.0)	(87.25, 0.0)	(136.85, 0.0)	(84.68, 0.0)	(66.16, 0.06)	(60.3, 0.15)	(72.58, 0.02)	(50.99, 0.43)	(58.4, 0.19)	50.0
22	(83.9, 0.0)	(47.85, 0.56)	(60.39, 0.15)	(99.85, 0.0)	(70.22, 0.03)	(90.22, 0.0)	(89.96, 0.0)	(147.39, 0.0)	(91.55, 0.0)	(76.6, 0.01)	50.0

Table 14: Batch2: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different U.S. states is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook. The Chi Square test was not able to be performed for the highlighted values.

Ad ID	Solar Cells	Solar Cells (Copy)	Solar Cells + Contrarian Logo	Solar Cells + Advocacy Logo	Oil rigs	Oil rigs (Copy)	Oil rigs + Contrarian Logo	Oil rigs + Advocacy Logo	Controls	Controls (Copy)	df
1	(56.15, 0.26)	(76.29, 0.01)	(120.82, 0.0)	(129.39, 0.0)	(82.96, 0.0)	(67.05, 0.05)	(111.55, 0.0)	(144.61, 0.0)	(82.41, 0.0)	(75.7, 0.01)	50.0
2	(103.46, 0.0)	(45.43, 0.66)	(69.82, 0.03)	(91.97, 0.0)	(88.38, 0.0)	(112.53, 0.0)	(56.23, 0.25)	(81.43, 0.0)	(83.0, 0.0)	(51.76, 0.41)	50.0
3	(100.08, 0.0)	(82.08, 0.0)	(94.15, 0.0)	(106.96, 0.0)	(35.4, 0.94)	(99.9, 0.0)	(87.68, 0.0)	(74.96, 0.01)	(93.73, 0.0)	(96.57, 0.0)	50.0
4	(87.17, 0.0)	(75.94, 0.01)	(72.36, 0.02)	(59.21, 0.17)	(68.65, 0.04)	(173.72, 0.0)	(63.46, 0.1)	(104.6, 0.0)	(98.66, 0.0)	(56.61, 0.24)	50.0
5	(175.35, 0.0)	(87.64, 0.0)	(106.15, 0.0)	(56.92, 0.23)	(108.4, 0.0)	(103.01, 0.0)	(81.05, 0.0)	(124.17, 0.0)	(73.78, 0.02)	(82.22, 0.0)	50.0
6	(85.33, 0.0)	(78.0, 0.01)	(62.29, 0.11)	(73.52, 0.02)	(90.17, 0.0)	(95.55, 0.0)	(65.63, 0.07)	(58.39, 0.19)	(87.8, 0.0)	(107.88, 0.0)	50.0
7	(114.37, 0.0)	(65.28, 0.07)	(80.3, 0.0)	(37.1, 0.91)	(80.11, 0.0)	(91.93, 0.0)	(148.37, 0.0)	(55.74, 0.27)	(46.31, 0.62)	(42.23, 0.77)	50.0
8	(119.44, 0.0)	(83.23, 0.0)	(96.1, 0.0)	(162.83, 0.0)	(58.64, 0.19)	(117.64, 0.0)	(140.48, 0.0)	(82.33, 0.0)	(137.94, 0.0)	(72.56, 0.02)	50.0
9	(80.43, 0.0)	(167.22, 0.0)	(105.45, 0.0)	(78.67, 0.01)	(85.61, 0.0)	(65.75, 0.07)	(88.78, 0.0)	(72.07, 0.02)	(100.58, 0.0)	(178.92, 0.0)	50.0
10	(98.4, 0.0)	(114.45, 0.0)	(58.7, 0.19)	(56.52, 0.24)	(123.32, 0.0)	(82.62, 0.0)	(83.53, 0.0)	(105.16, 0.0)	(115.62, 0.0)	(99.46, 0.0)	50.0
11	(145.62, 0.0)	(132.34, 0.0)	(84.09, 0.0)	(67.4, 0.05)	(98.36, 0.0)	(78.12, 0.01)	(81.24, 0.0)	(76.28, 0.01)	(107.2, 0.0)	(71.97, 0.02)	50.0
12	(108.4, 0.0)	(82.18, 0.0)	(124.75, 0.0)	(111.42, 0.0)	(99.05, 0.0)	(77.79, 0.01)	(60.49, 0.15)	(79.04, 0.01)	(121.28, 0.0)	(68.45, 0.04)	50.0
13	(76.28, 0.01)	(139.4, 0.0)	(84.2, 0.0)	(109.85, 0.0)	(84.53, 0.0)	(106.61, 0.0)	(89.0, 0.0)	(99.93, 0.0)	(70.95, 0.03)	(83.73, 0.0)	50.0
14	(45.35, 0.66)	(101.0, 0.0)	(79.78, 0.0)	(51.63, 0.41)	(60.56, 0.15)	(50.03, 0.47)	(102.07, 0.0)	(71.49, 0.02)	(99.34, 0.0)	(120.22, 0.0)	50.0
15	(74.58, 0.01)	(88.26, 0.0)	(100.78, 0.0)	(126.71, 0.0)	(66.83, 0.06)	(79.3, 0.01)	(88.17, 0.0)	(78.04, 0.01)	(137.73, 0.0)	(64.26, 0.08)	50.0
16	(68.78, 0.04)	(62.84, 0.1)	(59.53, 0.17)	(75.44, 0.01)	(94.0, 0.0)	(63.42, 0.1)	(114.76, 0.0)	(62.78, 0.11)	(64.39, 0.08)	(57.58, 0.22)	50.0
17	(79.97, 0.0)	(134.0, 0.0)	(72.11, 0.02)	(71.85, 0.02)	(97.8, 0.0)	(73.4, 0.02)	(43.85, 0.72)	(119.27, 0.0)	(47.49, 0.57)	(61.72, 0.12)	50.0
18	(50.12, 0.47)	(125.92, 0.0)	(94.28, 0.0)	(79.73, 0.0)	(88.83, 0.0)	(87.96, 0.0)	(73.34, 0.02)	(74.46, 0.01)	(68.28, 0.04)	(76.84, 0.01)	50.0
19	(62.31, 0.11)	(52.73, 0.37)	(93.47, 0.0)	(46.61, 0.61)	(117.6, 0.0)	(60.83, 0.14)	(136.66, 0.0)	(75.28, 0.01)	(88.84, 0.0)	(96.01, 0.0)	50.0
20	(91.08, 0.0)	(126.32, 0.0)	(118.75, 0.0)	(90.96, 0.0)	(85.47, 0.0)	(91.01, 0.0)	(86.94, 0.0)	(93.59, 0.0)	(80.74, 0.0)	(79.47, 0.01)	50.0
21	(108.18, 0.0)	(130.64, 0.0)	(94.78, 0.0)	(80.13, 0.0)	(85.01, 0.0)	(66.4, 0.06)	(84.7, 0.0)	(105.86, 0.0)	(95.75, 0.0)	(99.19, 0.0)	50.0

Table 15: Batch3: χ^2 statistic and p values to test the hypothesis that the observed ad delivery in different U.S. states is proportional to the population estimates provided by Facebook. In a majority of ads, the observed ad delivery is not proportional to the population estimates from Facebook. The Chi Square test was not able to be performed for the highlighted values.

Gender	Estimated Audience Size (Lower Bound)	Estimated Audience Size (Upper Bound)
male	100725600	118361400
female	121112000	142313900
unknown	1997800	3190900

Table 16: Estimated Facebook ad audience size estimates for different genders

Ages	Estimated Audience Size (Lower Bound)	Estimated Audience Size (Upper Bound)
18-24	42538900	50135000
25-34	55250800	65054000
35-44	42453300	50032600
45-54	31922400	37540100
55-64	26205100	30632600
65+	26166800	30603100

Table 17: Estimated Facebook ad audience size estimates for different age groups

State	Estimated Audience Size (Lower Bound)	Estimated Audience Size (Upper Bound)
Alabama	3400000	4000000
Alaska	530500	624100
Arizona	4900000	5800000
Arkansas	2000000	2400000
California	27200000	32000000
Colorado	3800000	4400000
Connecticut	2400000	2800000
Delaware	644900	758700
Florida	16500000	19400000
Georgia	7500000	8900000
Hawaii	978900	1200000
Idaho	1200000	1400000
Illinois	8200000	9600000
Indiana	4400000	5200000
Iowa	2000000	2300000
Kansas	1900000	2300000
Kentucky	3000000	3500000
Louisiana	3200000	3700000
Maine	889400	1000000
Maryland	4100000	4800000
Massachusetts	4700000	5600000
Michigan	6300000	7500000
Minnesota	3400000	4100000
Mississippi	1900000	2300000
Missouri	3900000	4600000
Montana	656300	772100
Nebraska	1200000	1500000
Nevada	2300000	2700000
New Hampshire	895800	1100000
New Jersey	6200000	7300000
New Mexico	1200000	1400000
New York	13800000	16300000
North Carolina	7300000	8600000
North Dakota	493100	580100
Ohio	7500000	8800000
Oklahoma	2700000	3200000
Oregon	2700000	3200000
Pennsylvania	8000000	9400000
Rhode Island	761100	895400
South Carolina	3500000	4200000
South Dakota	548100	644900
Tennessee	4800000	5600000
Texas	21300000	25100000
Utah	2200000	2500000
Vermont	392800	462100
Virginia	5800000	6900000
Washington D. C.	694200	816700
Washington	4800000	5700000
West Virginia	1100000	1300000
Wisconsin	3700000	4300000
Wyoming	350300	412100

Table 18: Estimated Facebook ad audience size estimates for different U.S. states