



**Universität  
Zürich<sup>UZH</sup>**

# Fairness in Online Ad Auctions: the Role of the Auction Mechanism

An analysis of how economic competition leads to discrimination  
in the displaying of ads

Simon Giesch

15-733-355

**Bachelor Thesis**

**Completed at the Department of Informatics  
of the University of Zurich**

Head of Research Group: Prof. Dr. Anikó Hannák  
Supervisor: Stefania Ionescu

Department of Informatics at University of Zurich  
Submission date: 20. April 2022



## Abstract

Online advertisement has seen a steep growth in the past two decades due to the rapid digitalization. Different from traditional advertising, the display of online ads is often aided by several algorithms. This includes (a) an ad auction mechanism that allocates the available ad slots to advertisers based on their bids and (b) machine learning (ML) algorithm models and personal user data to match the ads with a receptive audience. Despite the advantages of using such algorithms, recent work showed the resulting system produces unfair outcomes for users.

The aim of this thesis is to understand the role of the auction mechanism in regards to the fairness of the outcome. More precisely, I first look at past literature in order to (a) understand the causes of discriminatory ad displays, (b) overview the current ad auction mechanisms, and (c) review the state of the art solutions aimed to increase fairness. Second, I investigate why the showing of economic opportunity ads exhibits gender discrimination and whether this problem could be overcome by a different ad auction mechanism. To do so, I develop an agent-based model for an unrestrained generalized second-price (GSP) ad auction and compared its performance to my own mechanism: Separated Slots auction.

Separated Slots auction ensures that the users seeing an economic opportunity ads are proportional to the ratio of female and male users. Since I used a simplistic agent-based model, there is room for further auditing the performance of the Separated Slots auction. The Separated Slots auction provides a simple solution, which burdens the platform instead of the advertisers with the cost of fairness.

## Abstrakt

Aufgrund der fortschreitenden Digitalisierung ist die Relevanz von Onlinewerbungen in den vergangenen zwei Jahrzehnten stark gestiegen. Im Gegensatz zur traditionellen Werbung, profitiert Onlinewerbung von mehreren Optimierungsalgorithmen. Einerseits automatisiert ein Online Auction Mechanism die Auktion von freien Werbeslots, andererseits werten Machine Learning ML Algorithmen Benutzerdaten aus, um ein optimales Zielpublikum zu erreichen. In den vergangenen Jahren wurde aufgezeigt, dass diese Algorithmen trotz vieler Vorteile auch die Ursache von unfairer Diskriminierung der Benutzer sind.

Ziel dieser Arbeit ist es, die Rolle von Auction Mechanism in Bezug auf faire Resultate zu untersuchen. Zuerst arbeite ich den bisherigen Stand der Literatur auf, um (a) die Ursachen von Diskriminierung in Onlinewerbung zu verstehen, (b) die Auction Mechanism zu präsentieren, welche momentan im Gebrauch sind, und (c) einen Überblick zu den bereits existierenden Lösungen zu verschaffen. Danach untersuche ich, wieso Job- oder Wohnungswerbung diskriminierend gegenüber Geschlechtern ist und ob diese Diskriminierung mit einem anderen Auction Mechanism verhindert werden kann. Hierzu implementiere ich ein Agent-Based Model für eine unbeschränkte Generalized Second-Price GSP Auktion und vergleiche diese mit den Ergebnissen meiner eigenen Lösung: Separated Slots Auction.

Die Separated Slots Auction garantiert, dass das Geschlecht der Benutzer, welche Werbung für Jobs sehen, proportional zu dem Verhältnis der Plattformbenutzer ist. Da ich ein vereinfachtes Agent-Based Model benutzt habe, ist es notwendig weitere Untersuchungen mit komplexeren Modellen zu machen. Die Separated Slots Auction ist eine einfache Lösung, welche die Kosten für Fairness auf die Plattform anstatt auf die Werbenden abwälzt.

## **Acknowledgement**

I would like to thank Stefania Ionescu for supervising my thesis and helping me with finding a beginning to my thesis. Her support and constructive critique help me to overcome difficult phases with slow progress as she was always reachable and open for questions and help. I would also like to thank Prof. Dr. Anikó Hannák and the whole Social Computing group for accepting my application and letting me, write my thesis at their academic chair. Finally, I thank my parents, family and friends for supporting me throughout my academic career by proof reading my work and discussing important topics with me.



# Contents

|  |            |
|--|------------|
| <b>List of Tables</b>  | <b>iii</b> |
| <b>List of Figures</b>   | <b>v</b>   |
| <b>1 Introduction</b>  | <b>1</b>   |
| <b>2 Research Methodology</b>  | <b>5</b>   |
| 2.1 Structure of the Work and Development of the Research Questions  | 5          |
| 2.2 Collection of the Research Papers and Research Methods . . . . . | 7          |
| <b>3 Online Advertisement</b>  | <b>9</b>   |
| 3.1 Types of Online Advertisements . . . . .                         | 9          |
| 3.2 Phases of an Advertisement Campaign . . . . .                    | 10         |
| 3.2.1 Ad Creation . . . . .  | 10         |
| 3.2.2 Ad Delivery . . . . .  | 12         |
| 3.3 Discrimination Caused by Stakeholders . . . . .                  | 13         |
| 3.3.1 Discrimination by the Advertiser . . . . .                     | 13         |
| 3.3.2 Discrimination caused by the Platform . . . . .                | 15         |
| <b>4 Ad Auction Mechanisms</b>                                       | <b>17</b>  |
| 4.1 Mathematical Definitions & Notation . . . . .                    | 17         |
| 4.1.1 Rank-By-Expected-Value . . . . .                               | 18         |
| 4.1.2 Simplifications & Strategy-Proof Auctions . . . . .            | 19         |
| 4.2 Generalized Second-Price . . . . .                               | 20         |
| 4.3 Vickrey-Clarke-Groves . . . . .                                  | 21         |
| 4.4 Further Development in Auction Algorithms . . . . .              | 22         |
| <b>5 Discrimination &amp; Fairness</b>                               | <b>23</b>  |
| 5.1 Discrimination . . . . .   | 23         |
| 5.2 Discrimination in Online Ad Auctions . . . . .                   | 25         |
| 5.3 Fairness . . . . .   | 26         |
| 5.4 Measures of Fairness . . . . .                                   | 27         |
| 5.4.1 Bidding Strategies for Advertisers . . . . .                   | 28         |
| 5.4.2 Models for Individual Fairness . . . . .                       | 29         |
| <b>6 Simulation</b>  | <b>31</b>  |

|          |  |           |
|----------|--|-----------|
| 6.1      | Design Decisions . . . . .                                       | 31        |
| 6.2      | Separated Slots Auction . . . . .                                | 34        |
| <b>7</b> | <b>Results</b>   | <b>35</b> |
| 7.1      | Discrimination in the Unrestrained GSP . . . . .                 | 35        |
| 7.2      | Proportional Representation in Separated Slots Auction . . . . . | 37        |
| 7.3      | Advantages and Disadvantages . . . . .                           | 38        |
| 7.4      | Limitations . . . . .  | 40        |
| <b>8</b> | <b>Conclusion</b>  | <b>43</b> |
| 8.1      | Future Work . . . . .  | 44        |
|          | Appendix . . . . .   | 49        |



# List of Tables

|     |   |    |
|-----|---|----|
| 2.1 | Starting collection of related work. . . . .                | 7  |
| 2.2 | Papers Collected Through Forward and Backward Propagation . | 8  |
| 2.3 | Search Terms and the Selected Sources . . . . .             | 8  |
| 6.1 | Set of Variables Used for Different Simulations . . . . .   | 31 |
| 6.2 | Classes and their Functionality . . . . .                   | 32 |



# List of Figures

|     |   |    |
|-----|---|----|
| 7.1 | Percentage of Ads Shown to Female and Male Users in Dependence of the Advertiser Ratio - Female to Male Ratio: 50%  | 36 |
| 7.2 | Percentage of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Not Exhausted      | 36 |
| 7.3 | Percentage of Ads Shown to Female and Male Users - Female to Male Ratio: 30% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Exhausted          | 37 |
| 7.4 | Percentage of Ads Shown to Female and Male Users in Dependence of the User Ratio - Retail to Economic Opportunity Ratio: 50%  | 38 |
| 7.5 | Platform Revenue in Dependence of the Advertiser Ratio  | 39 |
| 8.1 | Percentage of Ads Shown to Female and Male Users in Dependence of the Advertiser Ratio - Female to Male Ratio 90%   | 49 |
| 8.2 | Absolute Number of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Exhausted     | 49 |
| 8.3 | Absolute Number of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Not Exhausted | 50 |
| 8.4 | Percentage of Ads Shown to Female and Male Users in Dependence of the User Ratio - Retail to Economic Opportunity Ratio: 90%  | 50 |
| 8.5 | Platform Revenue in Dependence of the User Ratio  | 51 |



# Acronyms

**API** Application Programming Interface

**GSP** generalized second-price

**ML** machine learning

**PII** Personally Identifiable Information

**STEM** Science, Technology, Engineering, Mathematics

**VCG** Vickrey-Clarke-Groves



# Chapter 1

## Introduction

The rapid digitalization of the past two decades also brought a congruent growth in online advertisement. Many big digital technology companies like Alphabet and Meta are generating their main income via online advertisement. In 2020 Meta alone generated \$84.169B from ad revenue (i.e., 98% of its total revenue), increasing it by 21% in comparison to the previous year (Meta 2021). Google earned in the same year \$146.92B from advertisement, increasing its revenue by 8,98% in comparison to 2019 (Johnson 2021). In countries where the society is already strongly embedded in the digital world, online advertisement becomes increasingly important as a tool to precisely reach receptive clients. Online advertisement is not only used for retail and commercial advertising, but also for housing, job or political advertisements. Advertising platforms provide many tools for advertisers to target specific audiences, to maximize the impact of an ad campaign and to receive semi-live statistics on the ad performance (Ali, P., Korolova A., et al. 2019).

Transitioning from conventional to online advertising brings advantages to all parties: the advertisers, the users, and the platform. For the advertisers there are two main advantages. First, they can quickly measure their impact (i.e., the performance of the ad), as platforms often provide related metrics such as the number of clicks, views or impressions. Second, machine learning algorithms

and large available data sets for training allows to aim ads to a fitting target audience. As a result, the advertisers' budgets are used more efficiently. Users of the advertisement platform can also benefit from targeted advertisement, since they receive fitting ads and thus see goods and opportunities in which they are interested in. The platform itself benefits through effective auction algorithms, such as the GSP mechanism, and thus maximize the platform's revenue.

However, more recent work has uncovered that online advertising also comes with drawbacks. More precisely, the distribution of ads could be discriminatory either due to the behavior of the advertiser, the platform, or the design of the market. The advertiser generating an ad can unintentionally or maliciously use the tools of the platform to discriminate against groups defined through protected attributes such as race, self-declared gender, religion, etc. (Speicher et al. 2018). In addition, the platform itself is not just a neutral provider of advertisement. Instead, it can induce discrimination and skews in the target audience by selecting the audience based on the ads content (even when the advertiser intends to reach a broader audience (Ali, P., B. M., et al. 2019; Ali, P., Korolova A., et al. 2019)). Moreover, even if all these stakeholders themselves are behaving fair, an unregulated ad auction algorithm itself can generate discrimination through competition overflow between different advertisement categories (Ilvento, J. M., and C. S. 2020). Lambrecht and T. C. 2019 show that this competition overflow is caused by advertisers of goods and services (e.g. retailers), who value female users higher than male users on average. This causes a higher competition for female users. Thus advertiser (i.e. job or housing advertisers), who do not distinguish a users value based on gender, are not able to reach as many female as male users.

In this thesis I look into the causes for discrimination in online advertisement and their auction algorithms based on the following two research questions: (a) *what are possible sources for discrimination in online advertisement caused by the stakeholders*, and (b) *why does an unrestrained ad auction mechanism result*



*in discriminatory ad displays for advertisers who do not differentiate between users based on protected attributes.* I then present current measurements for fairness and solutions to discrimination in online ad auction algorithms with help of two other research questions: *what are the current measurements of fairness used in online ad auctions and what are the proposed solutions for solving discrimination in online ad auction algorithms.* Finally, I audit the unrestrained ad auction algorithm in a simulation and apply my own solution to answer the last research question: *Can I reproduce the problem described in the research question 4 in a simulation and provide my own solution (Separated Slot auction) to overcome this problem?*

In the Separated Slot auction, I allocate randomly a user's ad slots either to retail advertisers or to advertisers of economic opportunities, thus creating separate auctions for different advertiser types. The probability for ad slots' allocations to an advertiser type is based on the relative ratio of that advertiser type to the total number of advertisers. Separated Slot auction ensures that the ratio of users of a certain gender reached by advertiser of economic opportunities is proportional to the ratio of users of this gender on the platform. The platform carries the cost of this fairness by losing some of its revenue. However, the agent-based model used for the simulation is simplistic and therefore further investigation on the performance is needed.

In the chapter *Research Methodology*, I present the research structure, the research questions and the development of my thesis. Next, I explain the structure of generating and deploying an ad campaign on an advertisement platform such as Meta in *Online Advertisement*. I further discuss the stakeholders' possibilities for discrimination during these processes, like the usage of look-alike audiences by advertisers or the pre-selecting of the target audience by the platform (Speicher et al. 2018; Ali, P., B. M., et al. 2019). The chapters *Ad Auction Mechanisms* and *Discrimination & Fairness* include the necessary mathematical background for the understanding of online ad auction algorithms, the definitions of discrimination and different

fairness measures used by previous research (Ilvento, J. M., and C. S. 2020; Nasr and T. M. C. 2020). *Simulation* focuses on the implementation of a base algorithm, different GSP measures and my own solution. In *Results* I then show the results of my simulation and evaluate findings. Finally, I complete my thesis with the chapter *Conclusion*, where I discuss the research questions and present directions for future work.

# Chapter 2

## Research Methodology

In this chapter I present the structure of acquiring the necessary sources, the general approach to the topic and the evolution and derivation of the research questions. Since this thesis is not a systematic literature review, I decided to omit the research query, the selection criteria and similar elements.

### 2.1 Structure of the Work and Development of the Research Questions

At the start, my supervisor Stefania Ionescu provided a first outline, where she explained the central problem of the thesis as follows:

*"The problem here is that because some groups of people (e.g. women) are targeted more often than others, the competition for displaying ads is much higher for these heavily targeted groups. As a result, general (i.e. non-targeted) ads have a lower chance of being shown to these groups."*

To get a complete comprehension of the topic of online advertisement, I first had to look into the process of producing online advertisement on a platform. Next, I worded through the formulation and the notations of different ad auction algorithms and their advantages and drawbacks. Then, I read into

the discussion of discrimination and fairness and how to measure them. With the acquired knowledge I was able to understand the causes of and proposed solutions for discrimination through competition overflow. Finally, I simulated a GSP auction algorithm and implemented my own solution for the problem of discrimination due to competition overflow.

As a result, the goal of my thesis is to address the following research questions:

1. *What are possible sources for discrimination in online advertisement caused by the stakeholders (platform, advertiser, and user)?*
2. *What are the current measurements of fairness used in online ad auctions?*
3. *What are the proposed solutions for solving discrimination in online ad auction algorithms?*
4. *Why does an unrestrained ad auction mechanism result in discriminatory ad displays for advertisers who do not differentiate between users based on protected attributes?*
5. *Can I reproduce the problem described in the research question 4 in a simulation and provide my own solution to overcome this problem?*

The first three research questions focus on a in-depth literature review, while the remaining two deal with a practical solution to the problem. The chapter *Online Advertisement* focuses on research question 1. The answer to research question 3 depends on the cause of the discrimination and therefore I am answering it in various chapters. I present the solutions for discrimination caused by stakeholders in the chapter *Online Advertisement* and for discrimination caused by the algorithm in the chapter *Discrimination & Fairness*, which also answers research question 2. The chapters *Simulation*, *Results* and *Conclusion* focus on the answer to research question 4 & 5.

## 2.2 Collection of the Research Papers and Research Methods

Through provided papers and a Social Computing Lecture held by Piotr Sapiezynski I made the following starting collection of sources shown in Table 2.1.

| Starting Collection                               |
|---|
| Ali, P., B. M., et al. <a href="#">2019</a>       |
| Ali, P., Korolova A., et al. <a href="#">2019</a> |
| Ilvento, J. M., and C. S. <a href="#">2020</a>    |
| Lambrech and T. C. <a href="#">2019</a>           |
| Nasr and T. M. C. <a href="#">2020</a>            |
| Speicher et al. <a href="#">2018</a>              |

Table 2.1: Starting collection of related work.

From search terms, forward and backward propagation I then identified additional resources. Table 2.2 shows the sources, which I collected through forward and backward propagation and Table 2.3 shows the search terms which I used to find the other sources. For collecting I mainly used Google Scholar and if the source was not completely obtainable on Google Scholar, I continued my search with the search engine Swisscovery.

I collected the statistical data for Meta and Google from four websites. The main topics which I focused on during the collection were: *Types of Online Ad Auction Algorithms*, *Procedure of Online Advertising*, *Discrimination and Fairness in Algorithms* and *Discrimination and Fairness in Online Advertising*.

| Starting Source                         | Collected Sources  |
|---|--|
| Corbett-Davies et al. 2017 <sup>1</sup> | Schubert and H. M. 2019  |
| Finocchiaro et al. 2021                 | Chawla and J. M. 2020,<br>Corbett-Davies et al. 2017,<br>Yaghini, H., and Krause A. 2019 |
| Nasr and T. M. C. 2020                  | Ekstrand et al. 2022,<br>Finocchiaro et al. 2021   |
| Speicher et al. 2018                    | Angwin and T. 2016   |
| Yaghini, H., and Krause A. 2019         | Heidari et al. 2019  |

Table 2.2: Papers Collected Through Forward and Backward Propagation

| Used Search Terms        | Collected Sources    |
|--------------------------|----------------------|
| Ad Revenue Google 2020   | Johnson 2021         |
| Contextual Ads           | Li and J.-L. J. 2009 |
| Facebook Ad Revenue 2020 | Meta 2021            |
| Online Ad Auction        | Varian 2009          |
| Sponsored Search Ads     | Ghose and Y. S. 2009 |

Table 2.3: Search Terms and the Selected Sources

---

<sup>1</sup>This source was first obtained through source Finocchiaro et al. 2021 and afterwards used for back propagation.

# Chapter 3

## Online Advertisement

In this chapter I explain the procedure of generating and delivering an online ad campaign on an online platform. After presenting different types of online ads, I explain the two phases: *ad creation* and *ad creation*, based on the findings of Ali, P., B. M., et al. [2019](#); Ali, P., Korolova A., et al. [2019](#) and Speicher et al. [2018](#). These sources all focused on Meta's (former Facebook) ad business, thus the structure of Online Advertisement is also based on Meta's structure. However, other big advertisement platforms' structures share similarities with Meta's structure.

### 3.1 Types of Online Advertisements

There are three different types of online ads:

- *Sponsored search ads* are textual ads, which are displayed in search engines alongside non-sponsored results. Normally these sponsored ads are displayed on top of the results (Ghose and Y. S. [2009](#)). They are shown, based on certain keywords and/or on targeted attributes of the users.
- *Display ads* are banner or pop-up ads on top of a web page that are adapted to the content of the web page and the demographics of the users

(Li and J.-L. J. 2009). They vary in size and can contain text, images, videos and sound.

- *Contextual ads* can be visual or text-based (Li and J.-L. J. 2009). In comparison to display ads, they are embedded into the web site. They can also be based on the web page’s content, certain key words or targeting of a specific audience.

*Targeted* advertisement uses user information collected onsite or bought from third parties to evaluate a user and estimates (e.g. with a ML model) how fitting ads and users are. This allows to distribute ads to specific audiences, which are receptive for the advertisement. In traditional ads, advertisers can only targeting based on the location of the ad (e.g. billboard) or if it is displayed in another media based on the user demographic of this media.

Different platforms have different online advertisement procedures. However, their procedures will be more or less similar, because user-friendliness and familiarity limit the design. Ali, P., B. M., et al. 2019 divides the advertisement process into two high-level phases: *ad creation* and *ad delivery*.

## 3.2 Phases of an Advertisement Campaign

### 3.2.1 Ad Creation

During the *ad creation*, the advertiser

- defines the content of the ad, by providing a headline, text, image or video, and a link to go to, when a user clicks on the ad (Ali, P., B. M., et al. 2019).
- selects the targeting criteria for the audience (Ali, P., B. M., et al. 2019).
- specifies the budget, the advertisement duration and a bidding strategy (Ali, P., B. M., et al. 2019).

For targeting an audience, Meta offers advertisers three different possibilities



(Speicher et al. 2018).

- *Attribute-based targeting*: Meta offers over 1000 binary attributes, which range from demographic to interest and behavioral categories (e.g. vegetarian, liberal, race-affiliation, etc.). The other type of attributes are so called free-form attributes, which are labels derived from user data. They range from interests in specific websites and apps (e.g. myGayTrip.com) to religion, food or news preferences (e.g. evangelicalism, donuts, Marie Claire) to any kind of niche interest (e.g. trainspotting, scuba diving (Speicher et al. 2018)). These free-form attributes are generated through collection of data on the Meta websites, apps or by buying it from data aggregators.
- *Personally Identifiable Information (PII) based targeting*: An advertiser could upload a list of PII such as names, birthdays, addresses, email addresses, phone numbers, etc. Meta will then search its user base for the people specified in the list and target them with the ads (Speicher et al. 2018). This allows an advertiser to directly target specific people with their ads, if they are using a Facebook, Instagram or any other of Meta's platforms.
- *Look-alike audience targeting*. Similar to PII targeting, the advertiser uploads a list of PII of user profiles for which Meta searches for similar users to generate a broader audience (Speicher et al. 2018).

The advertiser can specify the budget by selecting a total budget, a daily budget, and, if desired, a bid cap (Ali, P., Korolova A., et al. 2019). For the bidding strategy there are also multiple strategies available:

- *Impression optimization*: This option maximize the amount of users seeing the ads. The advertiser can further specify if a user should see the ad multiple times or if the impression should be unique (Ali, P., B. M., et al. 2019; Ali, P., Korolova A., et al. 2019).
- *Engagement optimization*: This option maximizes the amount of social

media interaction with the ads (e.g. comment, share, like, etc.). The ads will be shown to users who are more likely to react and engage with the ad (Ali, P., B. M., et al. 2019).

- *Sales optimization*: This option maximizes the amount of advertisement revenue by showing it to users who are more likely to be interested in the good or service which is promoted. Sales are optimized through estimated clicks and landing page views by users (Ali, P., Korolova A., et al. 2019).

Once the advertiser finishes the ad creation and submits his ad, the platform delivers the ad to the target audience.

### 3.2.2 Ad Delivery

In the ad delivery phase the platform delivers the ad to the target audience. Either the platform itself provides an algorithm, which bids for each user visiting the platform during the delivery or a third party agency provides this service to the advertiser. The algorithm estimates a value for each user and simulates the bidding to maximize the auctioneers revenue under the constraint of the chosen bidding strategy.

Platforms also assign a score to the ads based on different factors such as their design quality and the fit with chosen target audience (Ali, P., Korolova A., et al. 2019). This score is also taken into consideration during the bidding for a new user.

In addition, platforms, like Meta, also provide semi-live statistics on the ad performance (Ali, P., Korolova A., et al. 2019). In the feedback statistics the advertiser normally gets an overview of the demographics reached over time; e.g., the amount of men and women reached with an ad.

### 3.3 Discrimination Caused by Stakeholders

In this section I explain the different stakeholders and their roles in discrimination as pointed out by Ali, P., Korolova A., et al. [2019](#); Ali, P., B. M., et al. [2019](#); Speicher et al. [2018](#).

#### 3.3.1 Discrimination by the Advertiser

The most obvious cause of discrimination would be the selection of protected attributes such as racial-affinity, gender, etc. to specifically target an audience or exclude a certain group from receiving the advertisements. Until 2016, advertisers on Meta could for example exclude certain users from receiving housing ads, based on their racial-affinity, which is a users self-perceived race Angwin and T. [2016](#). Because of social backlash and a lawsuit, Meta and several other platforms removed the option for selecting certain protected attributes, when advertising for houses, job and other economic opportunities (Speicher et al. [2018](#)). However Speicher et al. [2018](#) show, that an advertiser is still able to discriminate deliberately or unconsciously with any of the three targeting options.

#### Discrimination in Attribute-Based Targeting

Even if protected attributes are removed from the options, Speicher et al. [2018](#) show that there are plenty of other (free-form and binary) attributes, which can be used as a proxy for discrimination. First, they show using seemingly neutral binary attributes, like *conservative*, *liberal*, *vegetarianism*, *mountain biking*, *primary OS mac OS X*, etc. could be used to generate racially exclusive and inclusive audiences for the races *Asian*, *Black*, *Indian* and *White* in the demography of the USA. Next they also show that there are free-form attributes, which can be misused as proxies for protected attributes such as: *BlackNews.com*, *myGayTrip.com*, *Marie Claire*. Many protected groups such as religious groups, addicted people or the LGBT community can easily be directly targeted and are thus vulnerable for discrimination through free-form

attributes. E.g. *REHAB*, *LGBT community*, *Evangelicalism*, *Islam* (Speicher et al. 2018).

Finally, Meta also offers a marketing Application Programming Interface (API), which suggests similar attributes based on a chosen free-form attribute. This can be misused for discrimination in the following ways (Speicher et al. 2018):

1. An advertiser could use the API, to search for discriminatory free-form attributes, which taken by themselves appear to be neutral.
2. An advertiser could exploit the API suggestions to search for more and more biased free-form attributes until finishing with an extremely biased free-form attribute.

### **Discrimination in Targeting based on PII**

Discrimination through personal identifiable information is very simple: an advertiser creates a list of PII from people representing a protected group (Speicher et al. 2018). Getting information on such a protected group is also very easy. There are many public data sources as for example in the USA many states release public voter and criminal records, which are accessible for everyone (e.g.: voterrecords.com 2022; StateRecords.org 2022). An advertiser could also use his own collected customer information, if he is collecting it or buy information from data brokers. Note that not all people on a list of PII will be using the platform (with an identifiable account). Speicher et al. 2018 showed an advertiser could discriminate as such on Meta by using public voter records from North Carolina. They observed a high percentage of targetable users.

### **Discrimination in Targeting based on Look-Alike Audiences**

Starting with a source audience, which is highly discriminatory, an advertiser can generate a much bigger, but still discriminatory look-alike audience, according to Speicher et al. 2018 findings. They observed that, given a base audience where certain attributes are over- or under-represented, the look-alike audience

generated by Meta will also mimic the over- or under-representation of these attributes.

### 3.3.2 Discrimination caused by the Platform

Ali, P., B. M., et al. [2019](#) and Ali, P., Korolova A., et al. [2019](#) empirically examined Meta's ad delivery phase for discrimination. Through various black box testing they found that:

- Ads were shown to different demographics even when the advertiser does not target a specific audience. (There exists no option for fair targeting among all user demographics. An advertiser can only skip the targeting options.)
- After the creation of the ad, Meta evaluates the ad through machine learning and selects the best fitting audience. The main influence for the selected audience is hereby the ad's picture. So if for example an advertiser selects a picture of cosmetics, Meta decides to show this ad to significantly more women than men, even if no target audience was specified.
- If an advertiser of political ads selects a target audience with opposite political views, the cost for advertising is higher and the reach of the ad is lower, than when the advertiser selects a fitting audience. This is caused by the score estimating fit between the ad and the audience. If the model estimates the audience will not be interested in or like the ad, the score assigned to the ad will be worse.

The papers show that an advertisement platform, like Meta, is not just a "neutral" provider of ads, but itself induces and influences discrimination. Therefore, they argue this should be taken into account for legal regulations of these platforms.

The discrimination caused by the platform not only discriminates against the users, but also puts the advertiser at a disadvantage. Because even when the

advertiser themselves, has no discriminatory intent, the output from the ad campaign can still be discriminatory. Additionally, forcing a for example political advertiser to only advertise for a predetermined "fitting" audience through scoring, heavily limits the decisions an advertiser can make in comparison to traditional advertisement. It also tampers the reachability of audiences and can lead to filter bubbles for the users.

# Chapter 4

## Ad Auction Mechanisms

This chapter covers the functionality and notation of different ad auction mechanisms. This notation is common in the Ad Auction literature. While writing this chapter I oriented myself on formal models from Decarolis, G. M., and P. A. 2020; Easley and K. J. 2010; Varian 2009.

Below, I introduce the Vickrey-Clarke-Groves (VCG) and GSP models. I elaborate on the GSP model further as I decided to use it for my simulation of an online ad auction. I chose to focus on GSP since it is used by many big platforms such as Google, Microsoft Bing and Yahoo! for sponsored search ads. In contrast, VCG is used for example by Facebook, Twitter and by Google, which switched to VCG for sponsored search ads in 2012 (Decarolis, G. M., and P. A. 2020). The two algorithms only vary in their payment scheme and they share the same mathematical notations and definitions.

### 4.1 Mathematical Definitions & Notation

In online ad actions a set of agents (i.e., bidders or advertisers)  $i \in I = \{1, \dots, n\}$  is assigned to a set of positions (i.e. ad slots on a website)  $slot \in S = \{1, \dots, m\}$  for a single auction (i.e. auction for one user's ad slots). Without loss of generality one can assume  $m < n$  (i.e. the number of bidders is bigger than the number of available ad slots) because the auctioneer can introduce any

number "dummy bidders" with a reserve price as a minimum bid, if there are less agents than positions. The slots are ordered from the best to the worst position, where  $slot_1 = 1$  is the best position.

Each bidder's value for a certain user can be expressed in terms of a *value-per-click*  $v_i \geq 0$ . We assume that clicks on the ads generate value and every click from a user has the same value for the advertiser  $i$  independent of the position. We further introduce a *click-through-rate (CTR)*, which models the probability of a click on an ad in dependence of the ad's position  $slot_j$  and quality  $\hat{Q}_i$  (Varian 2009). The *click-through-rate* of an agent  $i$  in a slot  $slot_j$  can therefore be described as

$$CTR_{i,j} = slot_j \cdot Q_i, \quad (4.1)$$

where  $slot_j \in [0, 1]$  is a position-specific effect with  $1 = slot_1 \geq slot_2 \geq \dots \geq slot_m > 0$  and the quality  $Q_i \in [0, 1]$  is an add-specific effect and depends on the relevance of the ad to the user and the quality of the agent's product or service.

#### 4.1.1 Rank-By-Expected-Value

The agent's true value and his expected value for a certain slot are defined as

$$v_{i,j} = CTR_{i,j} \cdot v_i = slot_j \cdot Q_i \cdot v_i, \quad (4.2)$$

$$\hat{v}_{i,j} = \widehat{CTR}_{i,j} \cdot \hat{v}_i = slot_j \cdot \hat{Q}_i \cdot \hat{v}_i, \quad (4.3)$$

where  $v_{i,j}$ ,  $v_i$ ,  $CTR_{i,j}$ ,  $Q_i$  are the true (unknown) values, Click-Through-Rate and quality and  $\hat{v}_{i,j}$ ,  $\hat{v}_i$ ,  $\widehat{CTR}_{i,j}$ ,  $\hat{Q}_i$  are the advertiser's expected values, Click-Through-Rate and Quality. To estimate the different expected variables the agent uses ML algorithms. The expected ad quality depends on ad-specific



factors such as ad design, while the expected value-per-click depends on user-specific factor such as his interests, age, etc. Even though the individual agent’s value  $v_i$  for a click and quality  $Q_i$  are different from each other they will all have the same relative values for  $slot_j$  in comparison to  $slot_{j+1}$ .

In an auction each agent also submits a per-click bid  $b_i$  which is transformed analogously to the value into an *effective bid* for each slot  $s$ :

$$b_{i,j} = slot_j \cdot \hat{Q}_i \cdot b_i. \quad (4.4)$$

Since the position-effect of a  $slot_j$  is the same for all agents, the final ranking in the auction depends on the *effective bid*  $b_i$  and the expected quality  $\hat{Q}_i$ . Instead of agents bidding individually on each position, the bids can be ranked by  $\hat{Q}_i \cdot b_i$ . The generated ranking is called *rank-by-expected-value*, which always finds an optimal assignment  $z^* = z_1^*, z_2^*, \dots, z_m^*$  of bids to positions, where  $z_i^* = \hat{Q}_i \cdot b_i$ .

### 4.1.2 Simplifications & Strategy-Proof Auctions

Ties in the *rank-by-expected-value* can be broken at random or according to the bidder’s index, where I decided on the latter to minimize my simulation’s runtime. To further simplify the notation and simulation of the algorithm, we can set the estimated quality of each agent to 1 (Varian 2009). We also assume that the expected  $\hat{v}_i$  and  $\widehat{CTR}_{i,j}$  of each advertiser is common knowledge (Decarolis, G. M., and P. A. 2020) by arguing that the different agents are able to learn each others expected values in the first few rounds of auctions. We can look at the auction after the advertisers have learned each others expected values.

An auctioneer prefers to maximize his profit, therefore he wants to force the agent to bid *truthfully*, where  $b_{i,j} = \hat{v}_{i,j}$ . An auction algorithm, which forces the bidders to always bid truthfully is called *strategy-proof*. The agent is charged less than the truthful bid in both algorithms. Therefore, the agent will also be

able to generate a utility as follows:

$$u_{i,j} = \begin{cases} v_{i,j} - c_{i,j} = slot_j \cdot Q_i(v_i - b_{i+1}), & \text{if } j = z_i^* \\ 0, & \text{else} \end{cases} \quad (4.5)$$

where  $c_{i,j}$  is the cost for the ad.

## 4.2 Generalized Second-Price

GSP has a very simple payment method. After bidding the agents are sorted by *rank-by-expected-value* as described before. Then the top ranking bidders, which win a slot in the auction are paying the price bidden by the advertiser on the next lower position multiplied with the ratio of the qualities of the two advertisers:

$$p_{gsp,i}(b) = \frac{\hat{Q}_{i+1} \cdot b_{i+1}}{\hat{Q}_i}, \quad (4.6)$$

if bidder  $i \leq m$ , where  $m$  is the last slot. Since I decided to simplify the estimated quality  $\hat{Q}_i$  by setting it equal to 1 for each advertiser, this corresponds to  $p_{gsp,i}(b) = b_{i+1}$ . The advertiser's effective bid  $b_i$  should be equal to  $\hat{v}_i$  if the advertiser is bidding truthfully.

The pricing of GSP is very straight-forward and can easily be explained to any customer, which is an advantage for a platform, however it is not strategy-proof. In other words, truthful bidding is not optimal for the agents because they learn each others values for the users. Therefore the advertisers will lower their bids to maximize their utilities, while remaining in the same slot position. Easley and K. J. [2010](#) shows that even though truthful bidding does not generally provide a Nash Equilibrium, there always exists an optimal envy-free Nash Equilibrium, which generates at least as much revenue as the VCG mechanism. With the introduction of *balanced bidding*, where every agents  $i$  bids in a way such that he is indifferent between position  $slot_i$  and  $slot_i + 1$  (ref eq. [4.7](#)),

the outcome of the GSP mechanism is exactly the same as the outcome of the VCG mechanism.

**Balanced Bidding:** A set of bids  $b = b_1, b_2, \dots, b_n$  is sorted descending and satisfies *balanced bidding* if  $b_1 = v_1$  and for all other bidders  $i \in 2, \dots, m$  the following equation holds:

$$slot_{j-1} \cdot (v_i - b_i) = slot_j \cdot (v_i - b_{i+1}) \quad (4.7)$$

I used this equation to adjust the advertisers' bidding strategies in my simulation.

### 4.3 Vickrey-Clarke-Groves

In the payment rule of the VCG mechanism the agent is charged the cost imposed on others as a result of demoting all advertisers below him by one rank:

$$p_{vcg,i}(b) = \sum_{k \neq i} \hat{v}_k(z^{-i}) - \sum_{k \neq i} \hat{v}_k(z^*) = \sum_{k=i+1}^{m+1} (slot_{k-1} - slot_k) \hat{Q}_k \cdot b_k, \quad (4.8)$$

where  $z^{-i}$  is the optimal assignment without the agent  $i$  and  $z^*$  the optimal assignment including the agent  $i$ . While for VCG it is more complicated to explain and implement the pricing system, it comes with the benefits of being strategy-proof and allocates the bidders efficiently Easley and K. J. 2010.

I decided to use a GSP mechanism in my simulation because of the following reasons:

- Although some platforms (e.g. Google) partly shifted from GSP to VCG pricing, GSP still is highly used by many platforms.
- In terms of implementation it seemed easier and more straightforward to implement GSP and most papers focused on solving discriminatory

outcomes in GSP mechanism.

- A VCG mechanism and a GSP mechanism in an envy-free Nash Equilibrium and under balanced bidding provide the same outcome. In this situation, the platform's revenue won't change in dependence of the algorithms.
- By inserting the payment scheme of the VCG mechanism in my simulation, the GSP mechanism can be translated to a VCG.

## 4.4 Further Development in Auction Algorithms

Many advertiser tend to use digital marketing agencies (DMA) to do the bidding for them. These DMAs belong to agency networks which are big. Therefore it can happen that these agencies bid for several of their client's in the same auction. This allows these network to use different strategies, where they can coordinate different clients bids to minimize their cost (Decarolis, G. M., and P. A. 2020). This can alter the platform's revenue and change the auction algorithm's functionality. Decarolis, G. M., and P. A. 2020 show that VCG outperforms GSP, if agency bidding occurs. There is also further research ongoing in making GSP and VCG resistant to these strategies.

# Chapter 5

## Discrimination & Fairness

In this chapter I present possible definitions of discrimination and fairness. I further elaborate the causes of discrimination inside an online ad auction when stakeholders (i.e. job advertisers) do not actively behave discriminatory against protected groups. I conclude the chapter with different solutions to solve discrimination in online ad auctions.

Since neither discrimination nor fairness have a universally accepted definition and are defined and discussed differently in multiple fields, I have to provide my understanding of these terms based on my research. I do not intend to provide a universal definition for these terminologies, but rather want to elaborate how I understood them to ease further comprehension of this thesis.

### 5.1 Discrimination

The Oxford Dictionary provides three different definitions for the word discrimination (OxfordLearnersDictionaries.com [2022a](#)):

- *the practice of treating somebody or a particular group in society less fairly than others (i.e. discrimination based on age/race/sex etc.)*

- *the ability to judge what is good, true, etc. synonym discernment (e.g. He showed great discrimination in his choice of friends.)*
- *the ability to recognize a difference between one thing and another; a difference that is recognized (e.g. to learn discrimination between right and wrong)*

While in the third definition of discrimination refers to a relatively neutral differentiation of status or things, the second one even has a positive connotation. The first definition is however the most common association with the word discrimination, which is generally negatively associated. The European Union has enshrined the right to non-discrimination in the treaty establishing the union and in the charter of fundamental rights, where it states:

*"Any discrimination based on any ground such as sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation shall be prohibited." (Article 21 (1) European-Parliament 2000)*

When I refer to *discrimination against a protected group*, I refer to groups defined by attributes, mentioned in the quote above. However, there is socially and legally accepted discrimination based on age (e.g. prohibition to consume alcohol, drive or consume certain media for minors). The understanding and sensitivity of discrimination is undergoing constant change. Today, discrimination is a growing research topic in Computer Science and researchers look into machine learning algorithms, mechanism designs and automated algorithms to discover and mitigate possible sources of discrimination.

In regards to advertising, it is sometimes socially accepted if an advertiser sells and advertises goods only to specific demographic groups (e.g. clothes or hygiene products only for men or women, hearing aid for elderly people, etc.). On the other hand it is legally prohibited for advertisers of economic

opportunities (i.e. housing or job ads) to only target one specific group (e.g. men).

For a better understanding, I work with the following framework. Advertisers are divided into two distinct binary sets. Advertisers of specific goods and services, which I refer to as *retailers*, are allowed to target specific group based on their legally protected attributes mentioned in the charter of fundamental rights European-Parliament 2000. To advertisers for housing, jobs or similar things, I refer as *advertisers of economic opportunities* or *economic opportunity advertisers*. These advertisers should not discriminate against protected groups. I refer to *fair behaviour* in an online ad auction, when the stakeholders act in the following way:

- *The platform does not pre-process the ad and thus does not limit the audience, before deploying the audience.*
- *The economic opportunities advertiser do not use any tools provided by the platform to discriminate against the users. They value users independent of protected attributes, but are allowed to value users based on other attributes (e.g. preferring users with a certain education for a job ad.)*
- *The retailers are allowed to target groups based on protected attributes.*

## 5.2 Discrimination in Online Ad Auctions

As I have already shown in chapter 3, that the platform and the advertiser can discriminate against the user. However, even when all stakeholders *behave fairly* the structure of the auction algorithm itself can lead to discriminatory outcomes. Lambrecht and T. C. 2019 find that gender-neutral Science, Technology, Engineering, Mathematics (STEM) ads are shown more than 20% more to

men than women. The differences are even higher for young adults. Lambrecht and T. C. 2019 also show that this discrimination does not come from consumer behaviour, such as women being less likely to click on a STEM ad. Women were even more likely to click on the STEM ads. They then show that the algorithms' behaviour is not caused by discriminatory training. For example, the algorithm just reflecting a lack of female workers in STEM jobs of the host country. With the usage of country-specific data from the World Bank (i.e. the extent of female labour participation), they showed a general lack of significance in correlation of these data with the algorithms behaviour.

The last remaining explanation is that this discrimination is caused by economic competition. Lambrecht and T. C. 2019 show that female users are on average more expensive than men to advertise to. On one hand, (young) female adults are more likely to engage with advertising and, on the other hand, women traditionally still control the majority of household expenses. Therefore (young) female users are more valuable for retailers which results in a higher competition in auctions of female users. This *competition overflow* makes it harder for *economic opportunity advertisers* to reach an even amount of women and men, if they do not adjust their bidding strategies.

### 5.3 Fairness

The definition given by the Oxford Dictionary on fairness is the following (OxfordLearnersDictionaries.com 2022b):

- *the quality of treating people equally or in a way that is reasonable*
- *a pale colour of skin or hair*

Since the second definition has a completely different meaning, I am only interested in the first definition. There is a lot of research being done on the



topic how to measure and achieve "equal" or "reasonable" treatment. Thus there exist many different notions of fairness, however multiple papers Ekstrand et al. 2022; Finocchiaro et al. 2021; Corbett-Davies et al. 2017 acknowledge that universal fairness is not an achievable concept and there are trade-offs between different fairness measures.

## 5.4 Measures of Fairness

There are multiple ways to differentiate the various fairness measures. Fairness can be divided into *ex-ante* fairness measures which are applied before competition (i.e. equality of opportunity or individual fairness) and *ex-post* measures, which are applied after competition (i.e. equality of outcomes, group fairness (Heidari et al. 2019)).

Equality of opportunity or individual fairness wants to ensure that only an individual's personal qualifications are taken into account (e.g. education, experience) in a decision-making problem. Thus minimizing the effect of circumstances (e.g. race, social background, gender, etc.) on the outcomes (Heidari et al. 2019). Ensuring this in the real world is difficult, because circumstances often effect someone's personal qualifications (e.g. growing up in a certain social class effects education). Only taking individual fairness into account without regarding the effect on circumstances on personal qualifications, will lead to discriminatory outcomes. For example minorities who are part of the lower income class and have less access to higher and private education, are at disadvantage when competing for high paying jobs.

Equality of outcomes or group fairness ensures that individual of different (protected) groups have the same (relative) representation in the outcomes of a competition. In this measurement of fairness the outcomes are independent of someone's circumstances. However, it cannot guarantee to take all individual's effort into considerations, leading to unequal treatment of individual from different groups with the same personal qualifications. For example a job as a

teacher could be given to an individual representing a minority even though another individual who does not represent this minority would be better fitting for this job.

These two interpretations of fairness are contradictory and discussed in many scientific fields. Which of these fairness measures a certain individual prefers, strongly depends on the individual’s world view. Applying any of these two fairness measures in a complete way in the real world is not feasible, because the real world is too complex and it is not possible to strictly divide an individual’s attributes into attributes for personal qualification and for circumstances.

Regarding online ad auctions, different stakeholders have partly contradictory interest and applying fairness constraint to the auctions to ensure group fairness in the algorithms will increase the costs for other stakeholders (e.g. less revenue for the platform or more bidding costs for the economic opportunities advertisers).

### 5.4.1 Bidding Strategies for Advertisers

Nasr and T. M. C. 2020 provide bidding strategies for advertisers in GSP auctions to reach the same amount of male as female users for a given advertiser. They provide two measures for the advertiser’s utility and two corresponding algorithms. The expected utility for the advertiser  $i$  takes the type  $\theta$  (i.e. gender) of the user into consideration.

$$\hat{u}_i = \sum_{\theta} p_i^{\theta} \cdot q(v_i^t; g_i^{\theta}) \cdot (v_i^t - c_i^t) \quad (5.1)$$

$p_i^{\theta}$  is the probability which advertiser  $i$  assigns to the user belonging to the type  $\theta$  (i.e. male, female).  $v_i^t$  &  $c_i^t$  are the value and cost of the current auction round  $t$  for advertiser  $i$  if he wins the round.  $q(v_i^t; g_i^{\theta})$  is the probability of winning the current round given the advertiser’s bid and cumulative density function  $g_i^{\theta}$  of the other advertisers’ bids. Nasr and T. M. C. 2020 further adjust the advertiser’s value function to allow him to *overbid* for female users

and to *underbid* for male ones in comparison to his true value for both of the user types.

Nasr and T. M. C. 2020 introduce an absolute and a ratio constraint for the advertiser  $i$ . The *absolute* parity constraint  $K$  ensures that the amount of female and male users won by the advertiser  $i$  deviates at most by the number  $K$  after each auction. I.e.  $n_i^m - n_i^f \leq K$ , where  $n_i^m, n_i^f$  are the number of male resp. female users won. The respective value function for the advertiser is defined as (Nasr and T. M. C. 2020):

$$V(k, \theta; g_i) = R^\theta(\Phi_i^\theta(k; g_i); g_i^\theta) + \delta N^\theta(\Phi_i^\theta(k; g_i), k; g_i) \quad (5.2)$$

On the other hand the *ratio* parity constraint  $K$  ensures that the ratio between female and male users won, is always smaller than  $K$  after each auction. Where the value function is defined as follows:

$$V(n_i^m, n_i^w, \theta; g_i) = R^\theta(\Phi_i^\theta(n_i^m, n_i^w); g_i^\theta) + \delta N^\theta(\Phi_i^\theta(n_i^m, n_i^w), n_i^m, n_i^w; g_i) \quad (5.3)$$

#### 5.4.2 Models for Individual Fairness

Ilvento, J. M., and C. S. 2020 introduce multi-category fairness based on envy-freeness and individual fairness. Users on a platform are envy-free, if they prefer their own allocations of ads over every one's else. It therefore ignores a users qualification and focuses only their preferences. Individual fairness as already mentioned, ignores preferences and focuses on qualifications. They further want to ensure that individual fairness is satisfied separately in each different ad category but also simultaneously for all categories. For this they introduce *compositional fairness*, which combines inter-category envy-freeness and total-variation fairness into one single mechanism. Implementing this measure would be rather difficult, since they did not provide any implementation themselves and only provide a mechanism for a single slot auction.

Chawla and J. M. [2020](#) introduce a mechanism to ensure individual fairness in one category of ads through inverse proportionality with nice properties such as being extendable for multiple different categories. For simplicity of design, I did not audit the GSP auction for individual fairness in my simulation.

All these different fairness measures incentivize further research of performance and compatibility between them.

# Chapter 6

## Simulation

In this chapter I outline my design decisions in implementing an ad auction algorithm for simulation purposes. I present the structure of my algorithm and explain different parameters used in my simulation.

### 6.1 Design Decisions

I decided to first implement a simple truthful GSP ad auction. For this I designed seven classes with different functionalities (ref. table 6.2).

The class *auction* was later divided into an unrestrained GSP auction and a Separated Slot auction, which contained my solution. For the generation of the users, advertisers and the estimate values I defined variables, for which I inserted different values during the simulations (ref. table 6.1).

| Variable Name                                       | Minimum | Maximum | Step Size |
|---|---------|---------|-----------|
| ratio of users to advertisers                       | 10      | 100     | .10       |
| size of advertisers                                 | 10      | 100     | .10       |
| ratio of retail to economic opportunity advertisers | 0.1     | 0.9     | +0.1      |
| ratio of female to male users                       | 0.1     | 0.9     | +0.1      |
| budget for the retailers                            | 100     | 1000    | .10       |

Table 6.1: Set of Variables Used for Different Simulations

| Class         | Functionality   |
|---------------|---|
| main          | <ul style="list-style-type: none"> <li>• reads the simulation's configuration CSV file</li> <li>• runs one or multiple simulations</li> <li>• saves the collected statistics into a new CSV file</li> </ul>   |
| simulation    | <ul style="list-style-type: none"> <li>• runs a single auction with 25 different seeds for randomisation.</li> <li>• tracks the statistics of each simulation and passes them to the main class</li> </ul>  |
| auction       | <ul style="list-style-type: none"> <li>• generates the users</li> <li>• generates the advertisers</li> <li>• runs an iteration over each user and every advertiser bids on them</li> <li>• selects the winning advertisers and calculates the payment scheme</li> <li>• updates the statistics with the winning advertisers' types and the user's type</li> </ul> |
| advertiser    | <ul style="list-style-type: none"> <li>• is either of type <math>r</math> or <math>e</math> (retailer or economic opportunity)</li> <li>• has a budget and bids on every user based on the user's type and its own type</li> </ul>  |
| user          | <ul style="list-style-type: none"> <li>• is either of type <math>f</math> or <math>m</math> (female or male)</li> </ul>   |
| csv_generator | <ul style="list-style-type: none"> <li>• generates the different configuration files for multiple simulation with different parameters</li> </ul>   |
| graph         | <ul style="list-style-type: none"> <li>• reads the output files of a simulation and plots a graph with the statistics</li> <li>• collects and averages the different seeds of one single simulation</li> </ul>  |

Table 6.2: Classes and their Functionality

Advertisers of economic opportunities value female and male users equally and therefore draw their estimated *value-per-click* from a log-normal distribution with variance  $\sigma^2 = 0.7$  and expected value  $\mu = -2.8$ . Retailers use the same

variance when drawing their estimated *value-per-click*, but their expected value depends on the user's gender:  $\mu_f = -3.5$  &  $\mu_m = -2.4$ . I decided to take these values based on the simulations of Nasr and T. M. C. 2020. To simplify the bidding even further, I set the estimated ad quality  $\hat{Q}_i = 1$  for all users, as discussed in [Ad Auction Mechanisms](#). I decided to use 10 slots per user. Lowering the amount of available slots would result in even higher competition between retailers and advertisers of economic opportunities. Similarly to decreasing the user size, while keeping the budget and the advertiser size fixed or increase the ratio of advertiser types. I distributed the position effect with even step size between the slots in the interval 0, 1, i.e.  $slot_1 = 1, slot_2 = 0.9, \dots, slot_{10} = 0.1$ . Distributing them with an uneven step sizes will only influence the balanced bidding, since only the equation of balanced bidding depends on the position effect (ref. equation 6.1). I also kept the reserve price constant at a value of 0.1. Increasing or decreasing the reserve price would mainly result in a change of revenue for the platform, without influencing the other results. If the reserve price is too high or the advertisers are close to deplete their budget, the reserve price will cause unfilled ad slots.

After implementing a base version of my algorithm with truthful GSP bidding, I adjusted it for the *balanced bidding* strategy. I used equation 4.7 to calculate the adjusted bidding from lowest to highest:

$$b_i = v_i - \frac{slot_i}{slot_{i-1}} \cdot (v_i - b_{i+1}) \quad (6.1)$$

The algorithm first calculates the last users bid adjustment, where  $b_{i+1} = \text{reserve price} = 0.1$ . Next, the algorithm iterates backwards over the other advertisers adjusting their bid accordingly. The first advertisers bid will not change and be truthful, since there exists no position effect from  $slot_{i-1}$ , if  $i = 1$ .

## 6.2 Separated Slots Auction

When implementing my own solution, I randomly assigned each user's ad slot to either retailers or advertisers of economic opportunities. The probability of assigning the ad slot to one of the two advertisers types is equal to the type's ratio to the total amount of advertisers:

$$slot_{i,k} \in \begin{cases} S_{retailers} \text{ with probability: } p_r = \frac{n_r}{n_{tot}} \\ S_{economic} \text{ with probability: } p_e = \frac{n_e}{n_{tot}} \end{cases}, \quad (6.2)$$

where the  $slot_{i,k}$  is the slot at position  $i$  of user  $k$  and  $n_r$ ,  $n_e$  and  $n_{tot}$  are the number of retailer or economic opportunity advertisers and the total amount of advertisers. After assigning each slot for a user, I then simulate two unrestrained GSP auction on the two separated sets of users. This separation of advertiser slots nullifies the competition overflow from different user valuation of different advertiser types.

After finishing the implementation of these auction algorithms, I generated all combinations of the variables shown in table 6.1 as configuration CSV files. I then ran all 972 CSV files with both auction methods, collected the statistics for them and plotted graphs where I show the ratio of the two ad types to female and male users. To limit the computation length for my simulation, I ran a lower number of randomized seeds for big user and advertiser sizes. The large sizes reduce the error significantly enough to produce a high accuracy even with a low amount of randomization seeds. My code is publicly available on my GitHub account <sup>1</sup>.

---

<sup>1</sup>[https://github.com/Wahlbar/Ad\\_Auction\\_Algorithm](https://github.com/Wahlbar/Ad_Auction_Algorithm)



# Chapter 7

## Results

In this chapter I discuss my results of the simulation. To start, I show that the discrimination of users based on gender is clearly visible in an unrestrained GSP auction. Then, I show that my solution of Separated Slots auction ensures that the ratio of female and male users reached by economic opportunity advertisers always reflects the ratio of total female and male users on the platform. However, there are limitations to my solution, which I discuss in the last section.

### 7.1 Discrimination in the Unrestrained GSP

I was able to generate a heavily discriminatory outcome with the simulation of an unrestricted GSP auction. Figure 7.1 shows a constant disadvantage for female users in comparison to male users in the case of an unrestrained GSP. Even in an auction, where only 10 % of the advertisers are retailers, the percentage of economic ads shown to female users is less than those shown to male users. This behaviour also happens, when the female to male ratio diverges from an even ratio of female and male users (ref. 8.1).

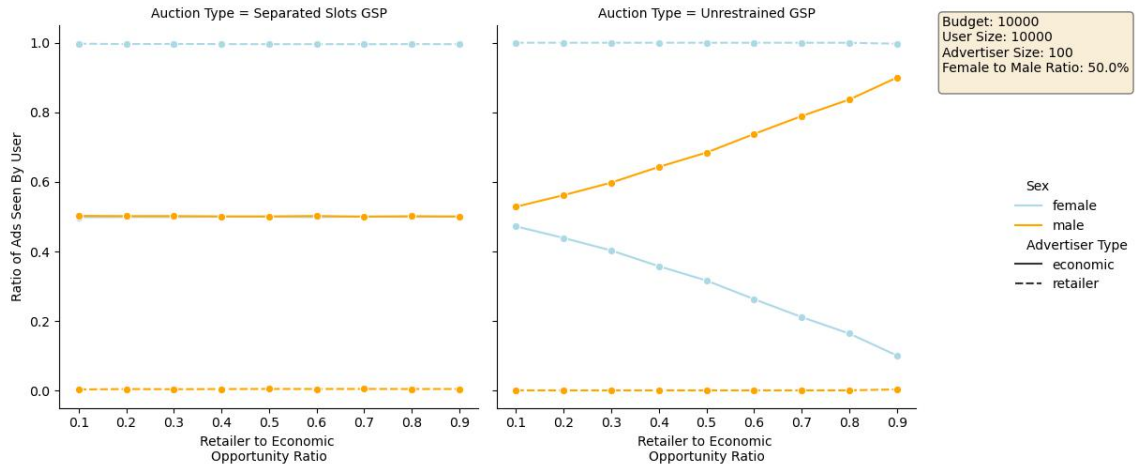


Figure 7.1: Percentage of Ads Shown to Female and Male Users in Dependence of the Advertiser Ratio - Female to Male Ratio: 50%

When looking at an averaged simulation such as Figure 8.3, the discriminatory behaviour is clearly visible. For an even ratio of female to male users and a ratio of 90% retailers to 10% economic opportunity advertisers, only around 10% of the economic opportunity ads are shown to female users.

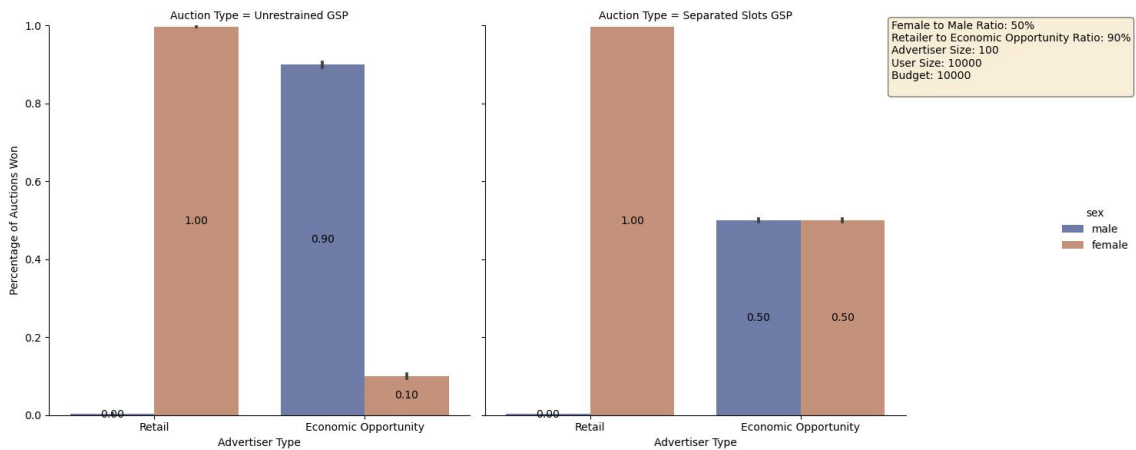


Figure 7.2: Percentage of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Not Exhausted

The discrimination measured in this simulation is much higher than in the real world. It stems from multiple factors explained in 7.4 Limitations. Further examples are displayed in the Appendix. From nearly 2000 plots generated in

my simulations, I selected the smallest set of Graphs, which displays the most information<sup>1</sup>.

## 7.2 Proportional Representation in Separated Slots Auction

All figures display the results of a Separated Slots auction done with the exact same variables as the unrestrained GSP auction. Figure 7.1 shows, that the percentage of different users seeing economic opportunity ads represents the ratio of users on the platform and does not change in dependence of the ratio of the advertiser types. Figure 7.2 shows a fair and even reach for economic opportunity ads for an even audience. Comparing Figure 7.2 and 7.3, we see that the ratios of economic opportunity ads seen by female and male users reflect the ratio of female and male users on the platform.

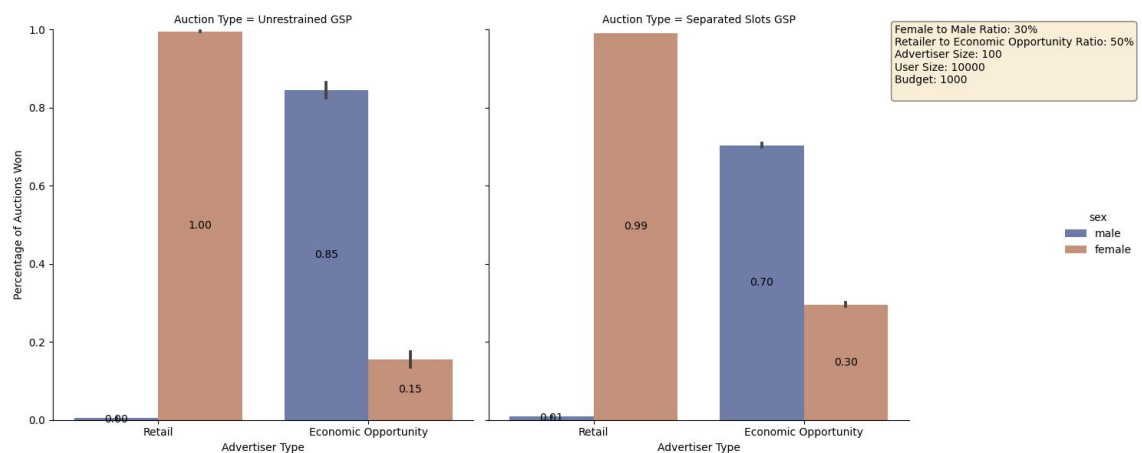


Figure 7.3: Percentage of Ads Shown to Female and Male Users - Female to Male Ratio: 30% - Retailer to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Exhausted

Figure 7.4 also visualizes the proportionality of the ratio of economic opportunity ads to ratio of gender of the users on the platform. It is clearly visible that the ratio of female seeing economic opportunity ads is unfair in comparison to the male ratio in the unrestrained GSP.

<sup>1</sup>Upon request I can provide additional plots.

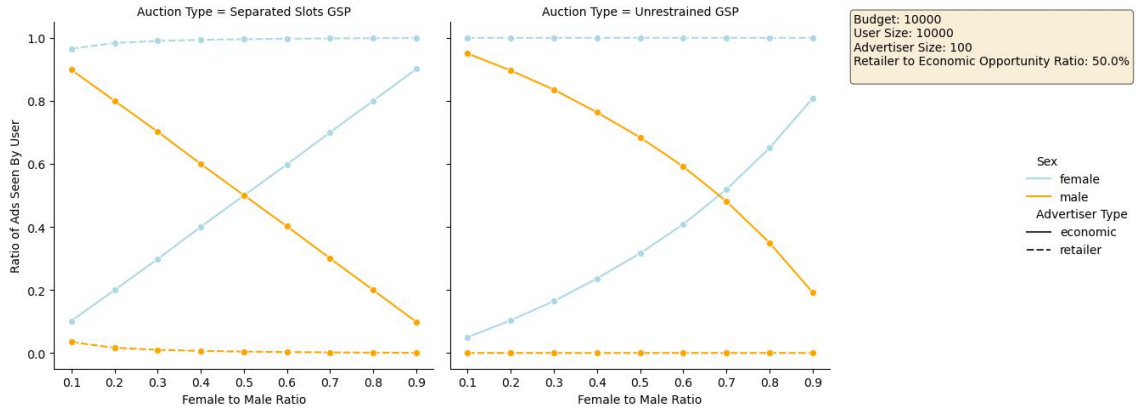


Figure 7.4: Percentage of Ads Shown to Female and Male Users in Dependence of the User Ratio - Retail to Economic Opportunity Ratio: 50%

### 7.3 Advantages and Disadvantages

The clear and most forward advantage of the Separated Slots auction is the proportionality of user ratios for economic opportunity. Since, I use an *ex-ante* restraint for the auction algorithms, advertisers are allowed to freely compete in their separated set of slots. Therefore each advertiser can still bid for the best fitting user and users can still receive relevant and interesting ads. Further the algorithm is straight-forward to implement and it only needs to be implemented by the platform to ensure non-discriminatory results, if the advertisers of economic opportunities behave fair.

However, as pointed out in chapter [Discrimination & Fairness](#) there is always a trade-off in fairness and utility. In the Separated Slots model, the retail advertisers will lose some utility because they are bidding on less available slots. Therefore individual retailers reach less fitting users for the same number of auctions. Through the limited slots, the competition inside the separated auctions will rise slightly. However, this loss of utility for the retailer is socially acceptable and will occur in any solutions for discrimination in ad auctions.

Users could also have a loss in utility. For example a user, who has no interests in economic opportunity, still has a chance that one of his slots is allocated to an economic opportunity auction. He will have a loss in his utility if it

happens. I did not implement any variables to track the utility of individual users. Therefore it is hard to say, if the users utility would decrease in general. I tend to argue, that there would at most be a insignificant loss of utility for users, since the separated GSP auctions have no further restrictions and advertisers with a very high value (and therefore a good match) for a certain user can still win their auction. Also, if we assume that all users are interested in some kind of economic opportunity advertisements, there will be a loss of utility for male users (because they receive on average less economic opportunity adds) and an equal increase of the female users' utilities.

The platform experiences a rather significant loss of utility in form of revenue loss. As visible in figure 7.5 the platform revenue is lower in the Separated Slots auction.

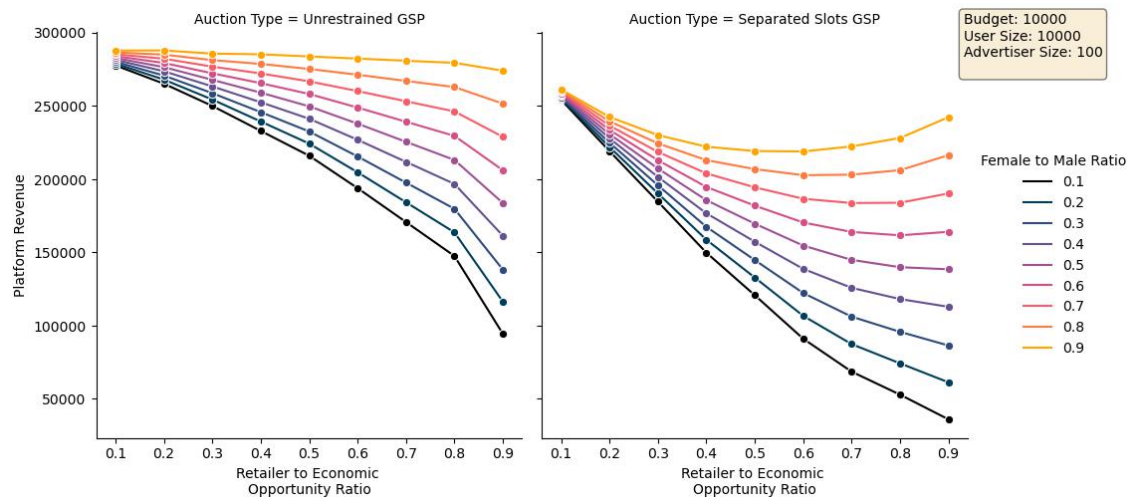


Figure 7.5: Platform Revenue in Dependence of the Advertiser Ratio

There are multiple explanations for this loss. Through the division of one auction into two, there will be two last slots, one for the retailers and one for the economic opportunity advertisers. In my simulation, I therefore let both last positions pay the reserve price. A better solution would be to adjust the reserve price according to the position effect of the last slot. This solution could help to mitigate a part of the platform's loss. However, there is also a loss stemming from the annihilated competition overflow. Through the separation economic

advertisers can bid less for women than before and win better positions. The overall loss of revenue for the platform will be a reason for the platform not to implement this solution.

## 7.4 Limitations

While I provide a simple and straightforward solution to discrimination in ad auction algorithms, there are still many limitations to my findings. First, I implemented a very simple base model for solving discrimination. It gives an easy and comprehensible overview for the main interactions between the auction, the types of ads seen by users and the platform's revenue. However, I did not track many additional useful metrics such as:

- The individual user's utility for the set of ads she/he is viewing.
- The individual advertiser's utility for the slots and users they win.
- The interaction with more complex users, which not only have one binary attribute, but also multiple categorical or continuous attributes (e.g. age, education, interests, etc.).
- The interaction of advertisers with not only binary preferences (i.e. retailers preferring men over women and economic opportunity advertisers having no preferences.).
- The interaction of not only binary advertisers, but the inclusion of additional retailers with new and different target audiences, similar to more complex users.

I did not implement other solutions for discrimination introduced in chapter 5 such as  $K$ -parity or  $K$ -ratio. Nasr and T. M. C. 2020 approach was already implemented by themselves. They even provided an iterative algorithm to solve a rather costly optimization of their Markov-Decision-Problem. However, the lack of compatibility between their approach and my simulation was quite big. They provided an advertiser oriented solution, where the non-discriminatory

advertiser can adjust his bidding based on the user's gender. To calculate their advertiser's utility and value function the cumulative density function (CDF) of all other users is needed. This is feasible for a single advertiser acting strategically. But when multiple economic opportunity advertisers start bidding strategically based on their solution, it becomes a difficult decision to choose in which order they will be bidding and adjusting their bid based on the CDF of each others. Additionally they focused on a single slot auction, while I used multiple slots. For a slot size of ten, I would have needed ten concurrent K-ratio or K-parity for each fair advertiser running. Even with their rather fast algorithm, this gets very costly in terms of computation with higher advertiser and user sizes. I think it is still possible to implement their measure into my simulation, but regarding to the time constraints for my thesis, I decided to rather focus on a working and complete base implementation.

Ivento, J. M., and C. S. [2020](#) on the other hand provided a measure for inter-category and intra-category fairness as a combination of individual fairness and envy-freeness. They do provide a theoretical mechanism. However they did not implement an algorithm themselves. Additionally their model is also only for one slot and could get more complicated with multiple positions. In order to simulate their model, I would have had not only to correctly implement and test their theoretical mechanism, but also to adapt it for multiple slots. Again, a high computational cost and time constraints of my thesis kept me from implementing it.





# Chapter 8

## Conclusion

In this thesis I have shown that there are many different causes for discrimination in online ad auction algorithms. Platforms provide tools to the advertisers for unintentional or deliberate discrimination against users. The platform also tampers the target audience, without the knowledge of the advertiser, which can cause harm for users and advertisers. An unregulated auction algorithm itself causes discrimination against heavily targeted audiences, resulting in discriminatory ad campaigns for advertiser, who want to behave fair. Additionally, the usage of machine learning algorithm to predict the users receptivity for ads can also be a cause of bias and discrimination (Favaretto, E., and E. B. S. 2019).

The cost of this discrimination bears not only the individuals, who are discriminated against, but also society itself. If not mitigated, discrimination in algorithms solidifies existing stigmata and stereotypes against protected groups. Thus, discovering and mitigating discrimination inside algorithms is important to support ongoing evolution of a society's self-perception.

The discrimination and skew caused by the online ad auction algorithm does not stop at the protected groups themselves but also splits these groups according to the click-through-rate. E.g. a STEM job advertisement is still reaching female users, but those reached are likely to be less interested in conventional retail ads and have a below average click-through-rate for retailers. While

a female user with the same qualification, but a high click-through-rate will be very expensive and therefore it's unlikely that she will see the STEM ad. Therefore stereotypical "nerdy" female are reached, who are less interested in stereotypical "female" consumption, thus reinforcing the stereotype of "nerdy" female STEM workers.

I introduced different models of fairness measures and different methods to mitigate discrimination in ad auction algorithms. I simulated a GSP mechanism and my own Separated Slots auction model as a possible base for ex-ante fairness measures. Since the platform already tracks the advertiser type to ensure certain legislation and for its own statistics, it is easy to implement the Separated Slots auction. Because the Separated Slots auction leads to less revenue, platforms will unlikely implement this solution. Nevertheless, this could be enforced through legislation.

## 8.1 Future Work

This thesis can serve as a base for many future studies. There is room for a more in-depth analysis of the Separated Slots auction in a more complex simulation. With the set of existing solution growing, comparing different approaches becomes inevitable. Until now most research has only be done in regards to binary division of gender. On one hand this was definitely a simplification for the models, but on the other hand there is also a lack of information in regards to other gender. Platforms such as Meta only track the groups: female, male and others, where others includes user who did not specify their gender or selected another gender. Collecting empirical data on discrimination against various protected groups is important to better understand the field. There is also a lack of attention in terms of harmful ads, such as gambling or scam ads, which no user wants to see but inevitably will. These ads will probably also discriminate against different protected groups.

# Bibliography

- Ali, M., Sapiezynski P., Korolova A., et al. (2019). “Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging”. In: arXiv: [1912.04255](https://arxiv.org/abs/1912.04255).
- Ali, M., Sapiezynski P., Bogen M., et al. (Nov. 2019). “Discrimination through optimization: How Facebook’s Ad delivery can lead to biased outcomes”. In: *Proceedings of the ACM on Human-Computer Interaction* 3.CSCW, pp. 1–30. DOI: [10.1145/3359301](https://doi.org/10.1145/3359301).
- Angwin, J. and Parris Jr. T. (Oct. 2016). *Facebook Lets Advertisers Exclude Users by Race*. URL: <https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race?token=oXaDaCvsdX3ZY7-YJd8F3L-6fSTJ6BUj> (visited on 01/22/2022).
- Chawla, S. and Jagadeesan M. (Nov. 2020). “Fairness in ad auctions through inverse proportionality”. In: *arXiv preprint arXiv:2003.13966*. arXiv: [2003.13966](https://arxiv.org/abs/2003.13966).
- Corbett-Davies, S. et al. (Aug. 2017). “Algorithmic decision making and the cost of fairness”. In: *Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining*. New York, NY, USA, pp. 797–806. DOI: [10.1145/3097983.3098095](https://doi.org/10.1145/3097983.3098095).
- Decarolis, F., Goldmanis M., and Penta A. (2020). “Marketing agencies and collusive bidding in online ad auctions”. In: *Management Science* 66.10, pp. 4433–4454.
- Easley, D. and Kleinberg J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge university press.

- Ekstrand, M. D. et al. (2022). “Fairness and Discrimination in Information Access Systems”. In: *arXiv preprint arXiv:2105.05779*. arXiv: [2105.05779](https://arxiv.org/abs/2105.05779).
- European-Parliament (2000). *Charter of fundamental rights of the European Union*.
- Favaretto, M., De Clercq E., and Elger B. S. (Feb. 2019). “Big Data and discrimination: perils, promises and solutions. A systematic review”. In: *Journal of Big Data* 6.1, p. 12. DOI: [10.1186/s40537-019-0177-4](https://doi.org/10.1186/s40537-019-0177-4).
- Finocchiaro, J. et al. (Mar. 2021). “Bridging Machine Learning and Mechanism Design towards Algorithmic Fairness”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: Association for Computing Machinery, pp. 489–503. ISBN: 9781450383097. DOI: [10.1145/3442188.3445912](https://doi.org/10.1145/3442188.3445912).
- Ghose, A. and Yang S. (Oct. 2009). “An empirical analysis of search engine advertising: Sponsored search in electronic markets”. In: *Management science* 55.10, pp. 1605–1622. DOI: [10.1287/mnsc.1090.1054](https://doi.org/10.1287/mnsc.1090.1054).
- Heidari, H. et al. (Jan. 2019). “A moral framework for understanding fair ml through economic models of equality of opportunity”. In: *Proceedings of the conference on fairness, accountability, and transparency*. New York, NY, USA, pp. 181–190. DOI: [10.1145/3287560.3287584](https://doi.org/10.1145/3287560.3287584).
- Ilvento, C., Jagadeesan M., and Chawla S. (Jan. 2020). “Multi-category fairness in sponsored search auctions”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Barcelona, Spain, pp. 348–358. DOI: [10.1145/3351095.3372848](https://doi.org/10.1145/3351095.3372848).
- Johnson, J. (Feb. 2021). *Google: ad revenue 2001-2020*. URL: <https://www.statista.com/statistics/266249/advertising-revenue-of-google/> (visited on 01/19/2022).
- Lambrecht, A. and Tucker C. (July 2019). “Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads”. In: *Management science* 65.7, pp. 2966–2981. DOI: [10.1287/mnsc.2018.3093](https://doi.org/10.1287/mnsc.2018.3093).

- Li, Y. and Jhang-Li J. (Jan. 2009). “Pricing display ads and contextual ads: Competition, acquisition, and investment”. In: *Electronic Commerce Research and Applications* 8.1, pp. 16–27. DOI: [10.1016/j.elerap.2008.06.001](https://doi.org/10.1016/j.elerap.2008.06.001).
- Meta (Jan. 2021). *Facebook reports fourth quarter and full year 2020 results*. URL: <https://investor.fb.com/investor-news/press-release-details/2021/Facebook-Reports-Fourth-Quarter-and-Full-Year-2020-Results/default.aspx> (visited on 01/18/2022).
- Nasr, M. and Tschantz M. C. (Jan. 2020). “Bidding strategies with gender nondiscrimination constraints for online ad auctions”. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: Association for Computing Machinery, pp. 337–347. ISBN: 9781450369367. DOI: [10.1145/3351095.3375783](https://doi.org/10.1145/3351095.3375783).
- OxfordLearnersDictionaries.com (2022a). *discrimination noun - Definition, pictures, pronunciation and usage notes*. URL: <https://www.oxfordlearnersdictionaries.com/definition/english/discrimination?q=discrimination> (visited on 04/13/2022).
- (2022b). *fairness noun - Definition, pictures, pronunciation and usage notes*. URL: <https://www.oxfordlearnersdictionaries.com/definition/english/fairness?q=fairness> (visited on 04/14/2022).
- Schubert, C. and Hütt M. (Mar. 2019). “Economy-on-demand and the fairness of algorithms”. In: *European Labour Law Journal* 10.1, pp. 3–16. DOI: [10.1177/2031952519829082](https://doi.org/10.1177/2031952519829082).
- Speicher, T. et al. (Nov. 2018). “Potential for discrimination in online targeted advertising”. In: *Conference on Fairness, Accountability and Transparency*. PMLR. New York, NY, USA, pp. 5–19.
- StateRecords.org (2022). *Access State Records Online*. URL: <https://staterecords.org/criminal.php> (visited on 04/17/2022).
- Varian, H. R. (May 2009). “Online ad auctions”. In: *American Economic Review* 99.2, pp. 430–434. DOI: [10.1257/aer.99.2.430](https://doi.org/10.1257/aer.99.2.430).
- voterrecords.com (2022). *10,209,949 Voter Records in Ohio*. URL: <https://voterrecords.com/voters/oh/1> (visited on 04/17/2022).

Yaghini, M., Heidari H., and Krause A. (2019). “A Human-in-the-loop Framework to Construct Context-dependent Mathematical Formulations of Fairness”. In: *CoRR* abs/1911.03020. arXiv: [1911.03020](https://arxiv.org/abs/1911.03020).

# Appendix

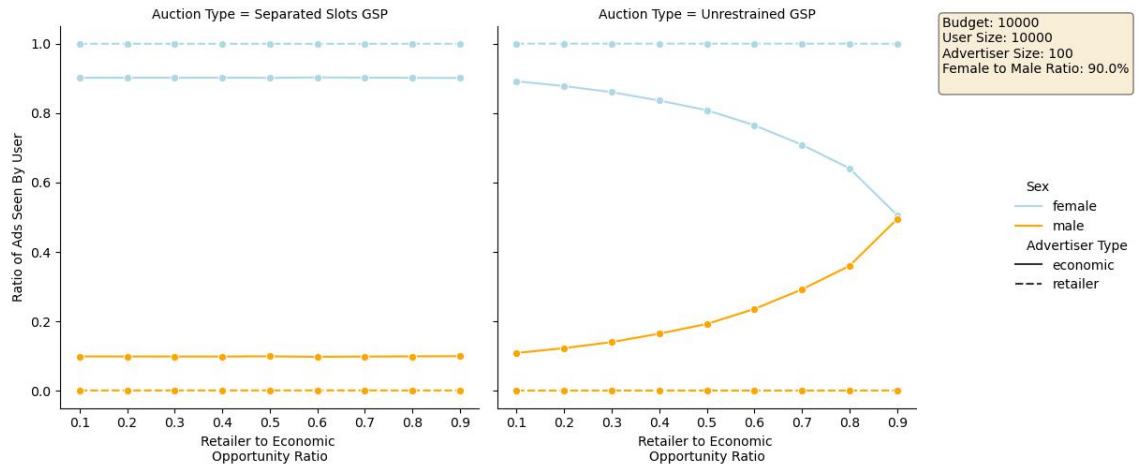


Figure 8.1: Percentage of Ads Shown to Female and Male Users in Dependence of the Advertiser Ratio - Female to Male Ratio 90%

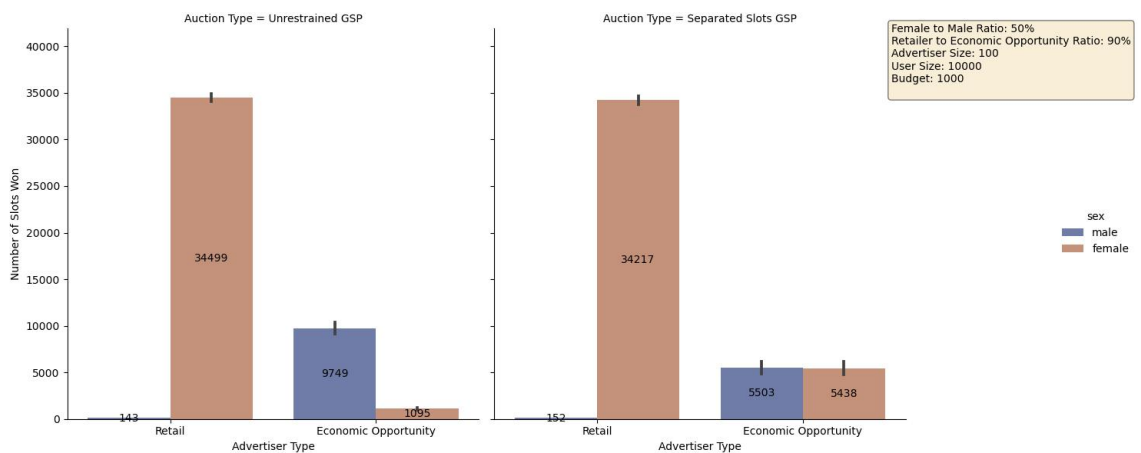


Figure 8.2: Absolute Number of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retailer to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Exhausted

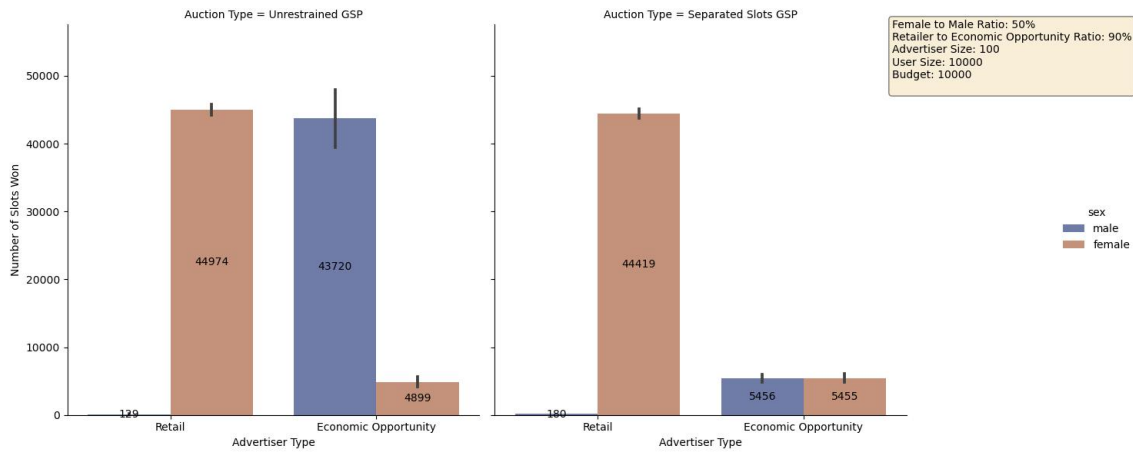


Figure 8.3: Absolute Number of Ads Shown to Female and Male Users - Female to Male Ratio: 50% - Retail to Economic Opportunity Ratio: 90% - Advertisers' Budgets: Not Exhausted




Figure 8.4: Percentage of Ads Shown to Female and Male Users in Dependence of the User Ratio - Retailer to Economic Opportunity Ratio: 90%



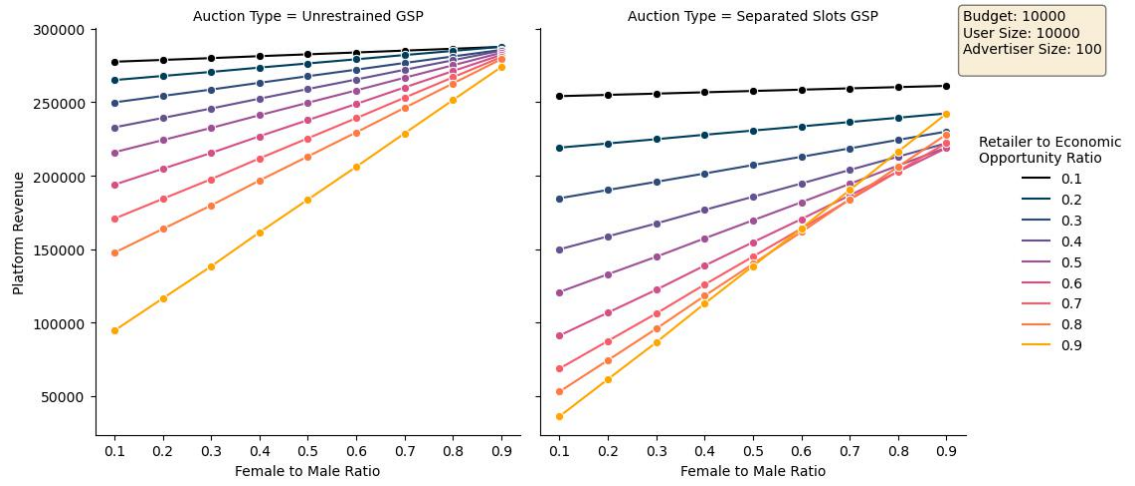


Figure 8.5: Platform Revenue in Dependence of the User Ratio