

ORIE 5355

Discrimination in Platforms

Nikhil Garg

Discrimination and Design Mitigations in online platforms

Special thanks for Muhammad Ali (Northeastern), Amit Datta (CMU), Anupam Datta (CMU), Aleksandra Korolova (USC), and Dennis Zhang (WUSTL) for their slides and important work

Overview

- A large literature has documented discrimination on online platforms
- This discrimination has many causes
- 2 example domains:
 - Online advertising
 - Marketplace platforms
- Ask: What *can* platforms do about this? What *should* they do?
 - What *is* legal? What *should be* legal?
 - Are companies investing enough?
- It's extremely hard to document this kind of thing from the outside

Discrimination in online advertising

Slides by/based on work by:

- *Amit Datta, Anupam Datta, Jael Makagon, Deirdre K. Mulligan, Michael Carl Tschantz. “Discrimination in Online Advertising: A Multidisciplinary Inquiry”*
- *Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, Aaron Rieke. “Discrimination through optimization: How Facebook’s ad delivery can lead to skewed outcomes”*
- *Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, Aaron Rieke. “Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging”*
- *Anja Lambrecht & Catherine Tucker. “Algorithmic Bias? An Empirical Study of Apparent Gender Based Discrimination in the Display of STEM Career Ads”*

Prior study on discrimination detection

Emplo	Is this illegal?				ons s
\$200k careerchange.com					
Why does it happen?					

The Barrett Group

Automated Experiments on Ad Privacy Settings: A tale of opacity, choice, and discrimination (PETS 2015)

Slide Credit: Datta et al. "Discrimination in Online Advertising: A Multidisciplinary Inquiry"

Section 704(b), Title VII of Civil Rights Act

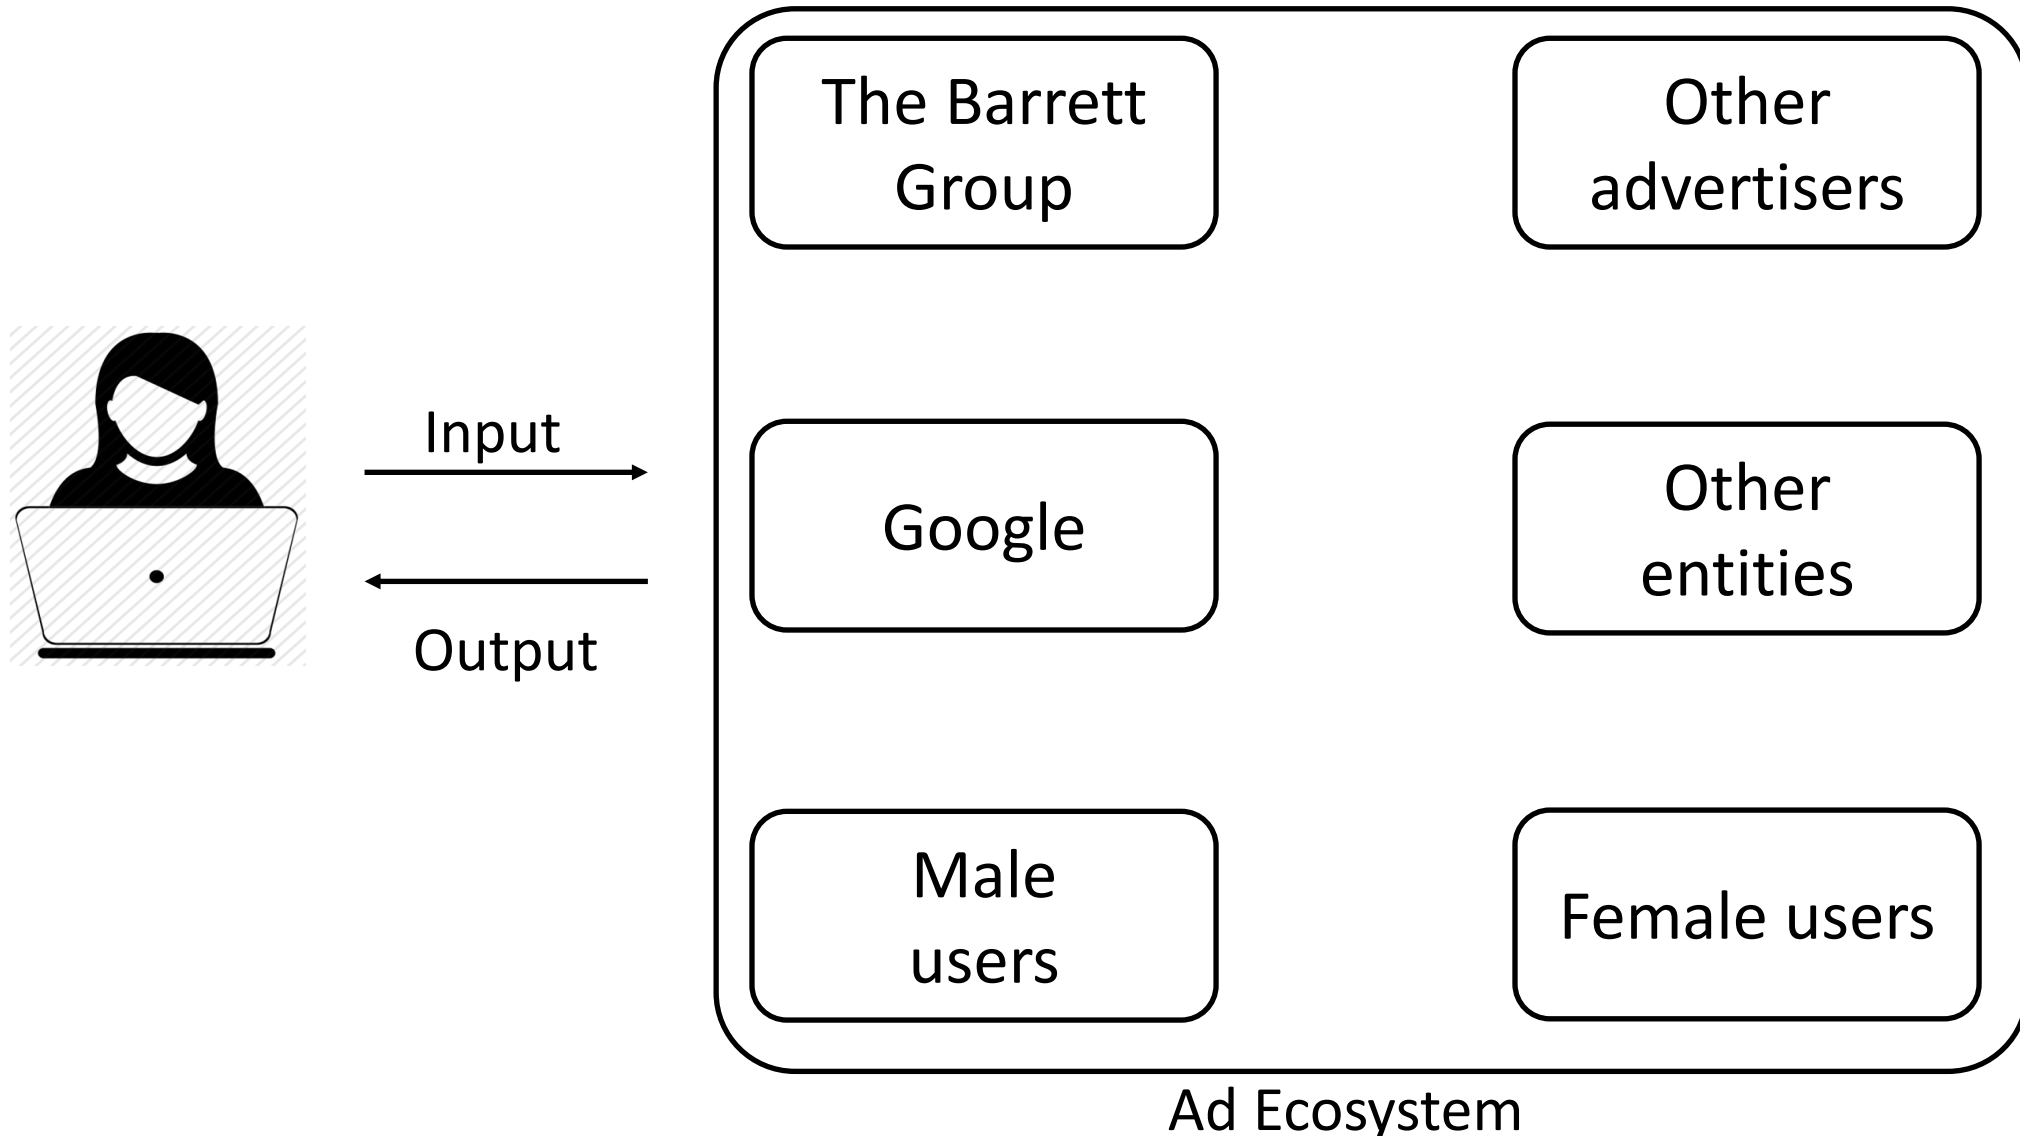
Unlawful “to print or publish or cause to be printed or published any ... advertisement relating to employment ... indicating any preference ... based on ... sex ...”

350
HELP WANTED, MEN

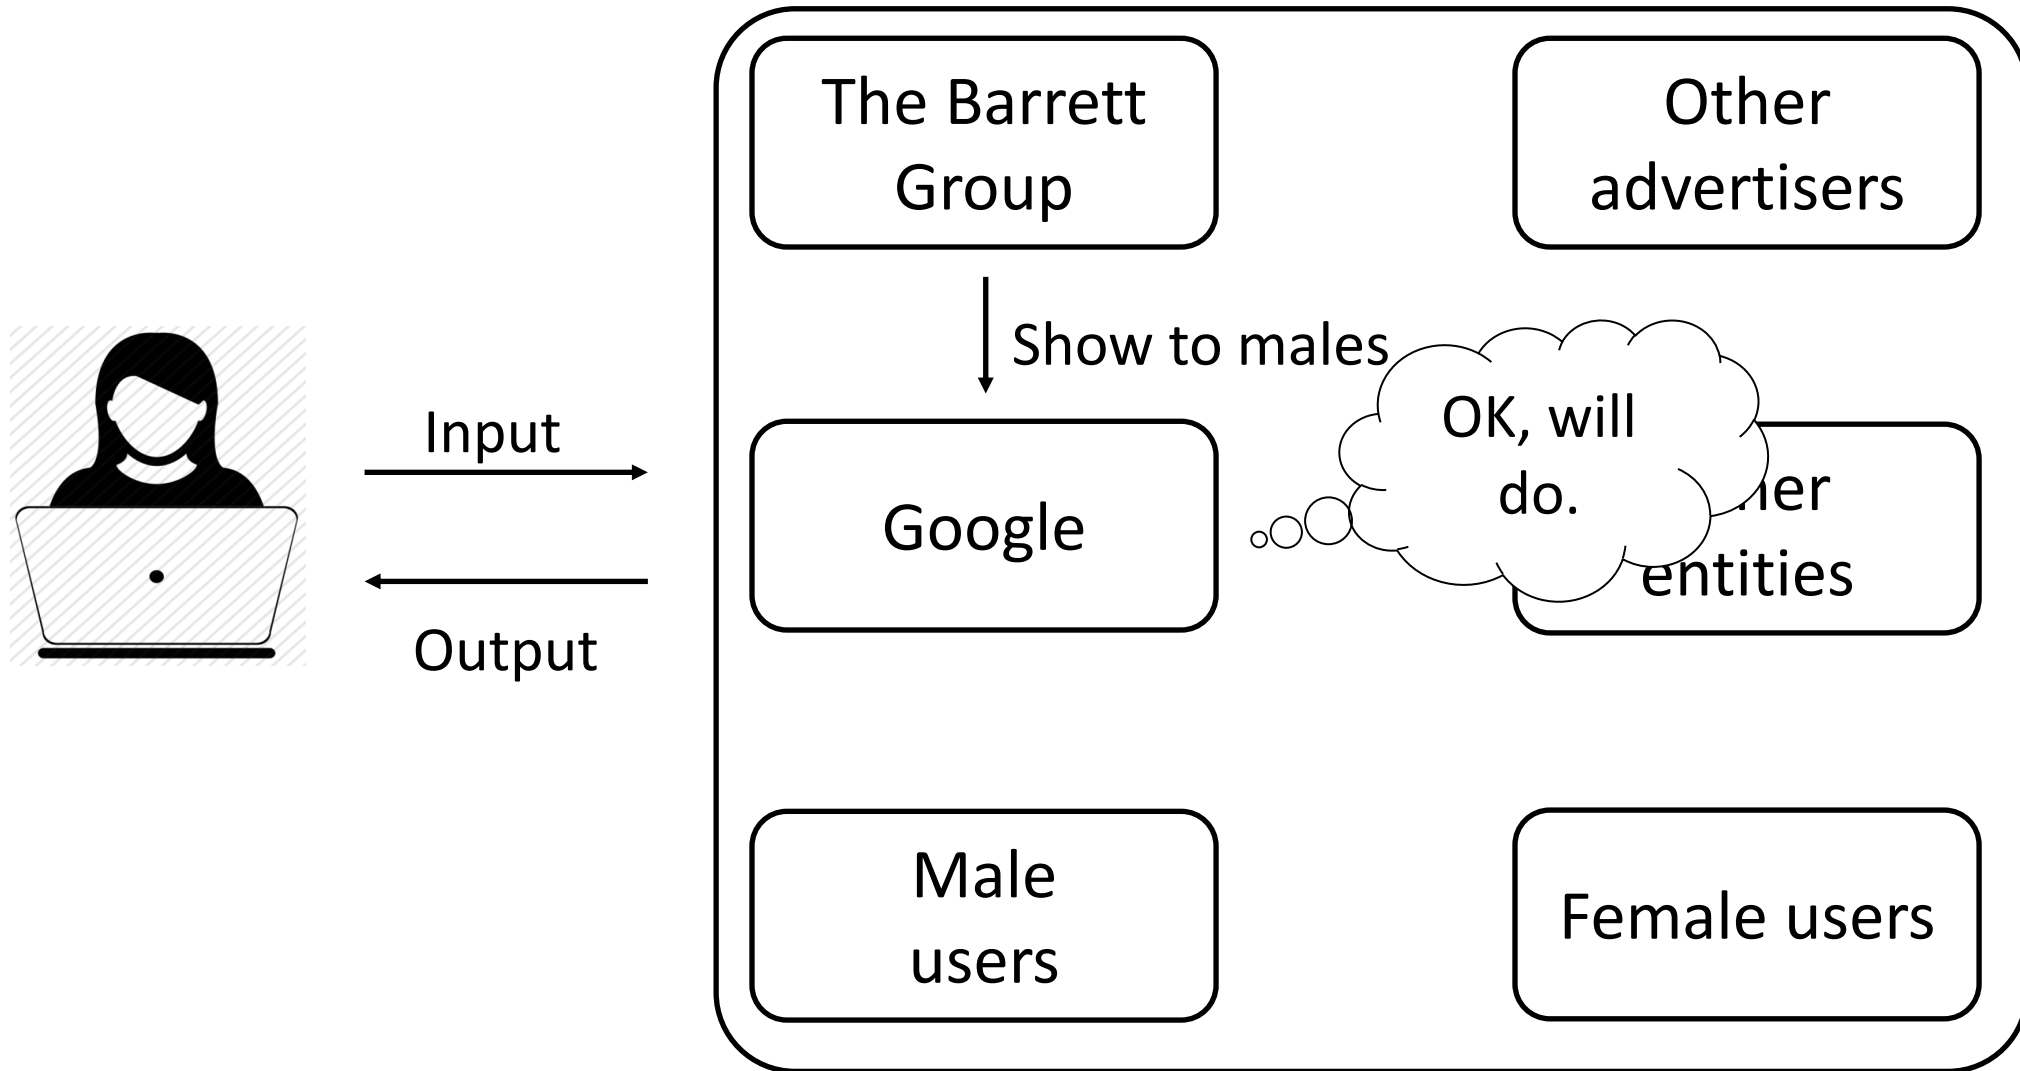
**ACCOUNTANT — CPA or CPA
candidate for small S.F. finan. dist.
CPA firm perm. Resume to this
paper AD No. 54081.¹**

Source: San
Francisco
Chronicle, Jan.
21, 1972

The ad ecosystem has many parties



Possible cause: direct advertiser targeting



Google allows targeting on gender

Choose how to target your ads

- Demographics
- Interests & remarketing (affinity audiences) – show ads to people based on their interests. [4 ideas](#)
- Use a different targeting method

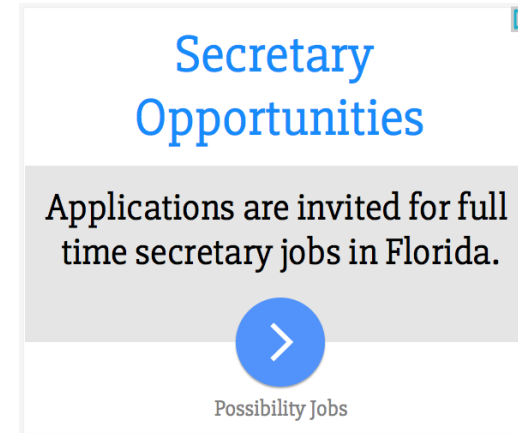
Demographics [?](#)

<p>GENDER</p> <p><input type="checkbox"/> Male</p> <p><input checked="" type="checkbox"/> Female</p> <p><input type="checkbox"/> Unknown ?</p>	<p>AGE</p> <p><input checked="" type="checkbox"/> 18-24</p> <p><input checked="" type="checkbox"/> 25-34</p> <p><input checked="" type="checkbox"/> 35-44</p> <p><input checked="" type="checkbox"/> 45-54</p> <p><input checked="" type="checkbox"/> 55-64</p> <p><input checked="" type="checkbox"/> 65 or more</p> <p><input type="checkbox"/> Unknown ?</p>	<p>PARENTAL STATUS</p> <p><input checked="" type="checkbox"/> Parent</p> <p><input checked="" type="checkbox"/> Not a parent</p> <p><input type="checkbox"/> Unknown ?</p>
---	--	---

Reach a significantly wider audience by showing ads to people whose Age, Gender, and Parental status we do not know.

Sexist ad targeting

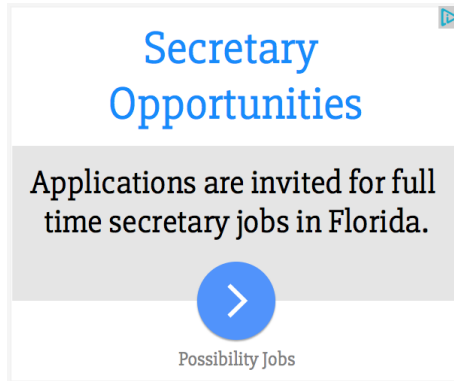
Targeted to females



Targeted to males



Sexist ad delivery



56,497 impressions, all to females



73,607 impressions, all to males

Google is trying to limit direct targeting

“Update to Personalized advertising policies: Housing, employment, and credit (October 2020)”

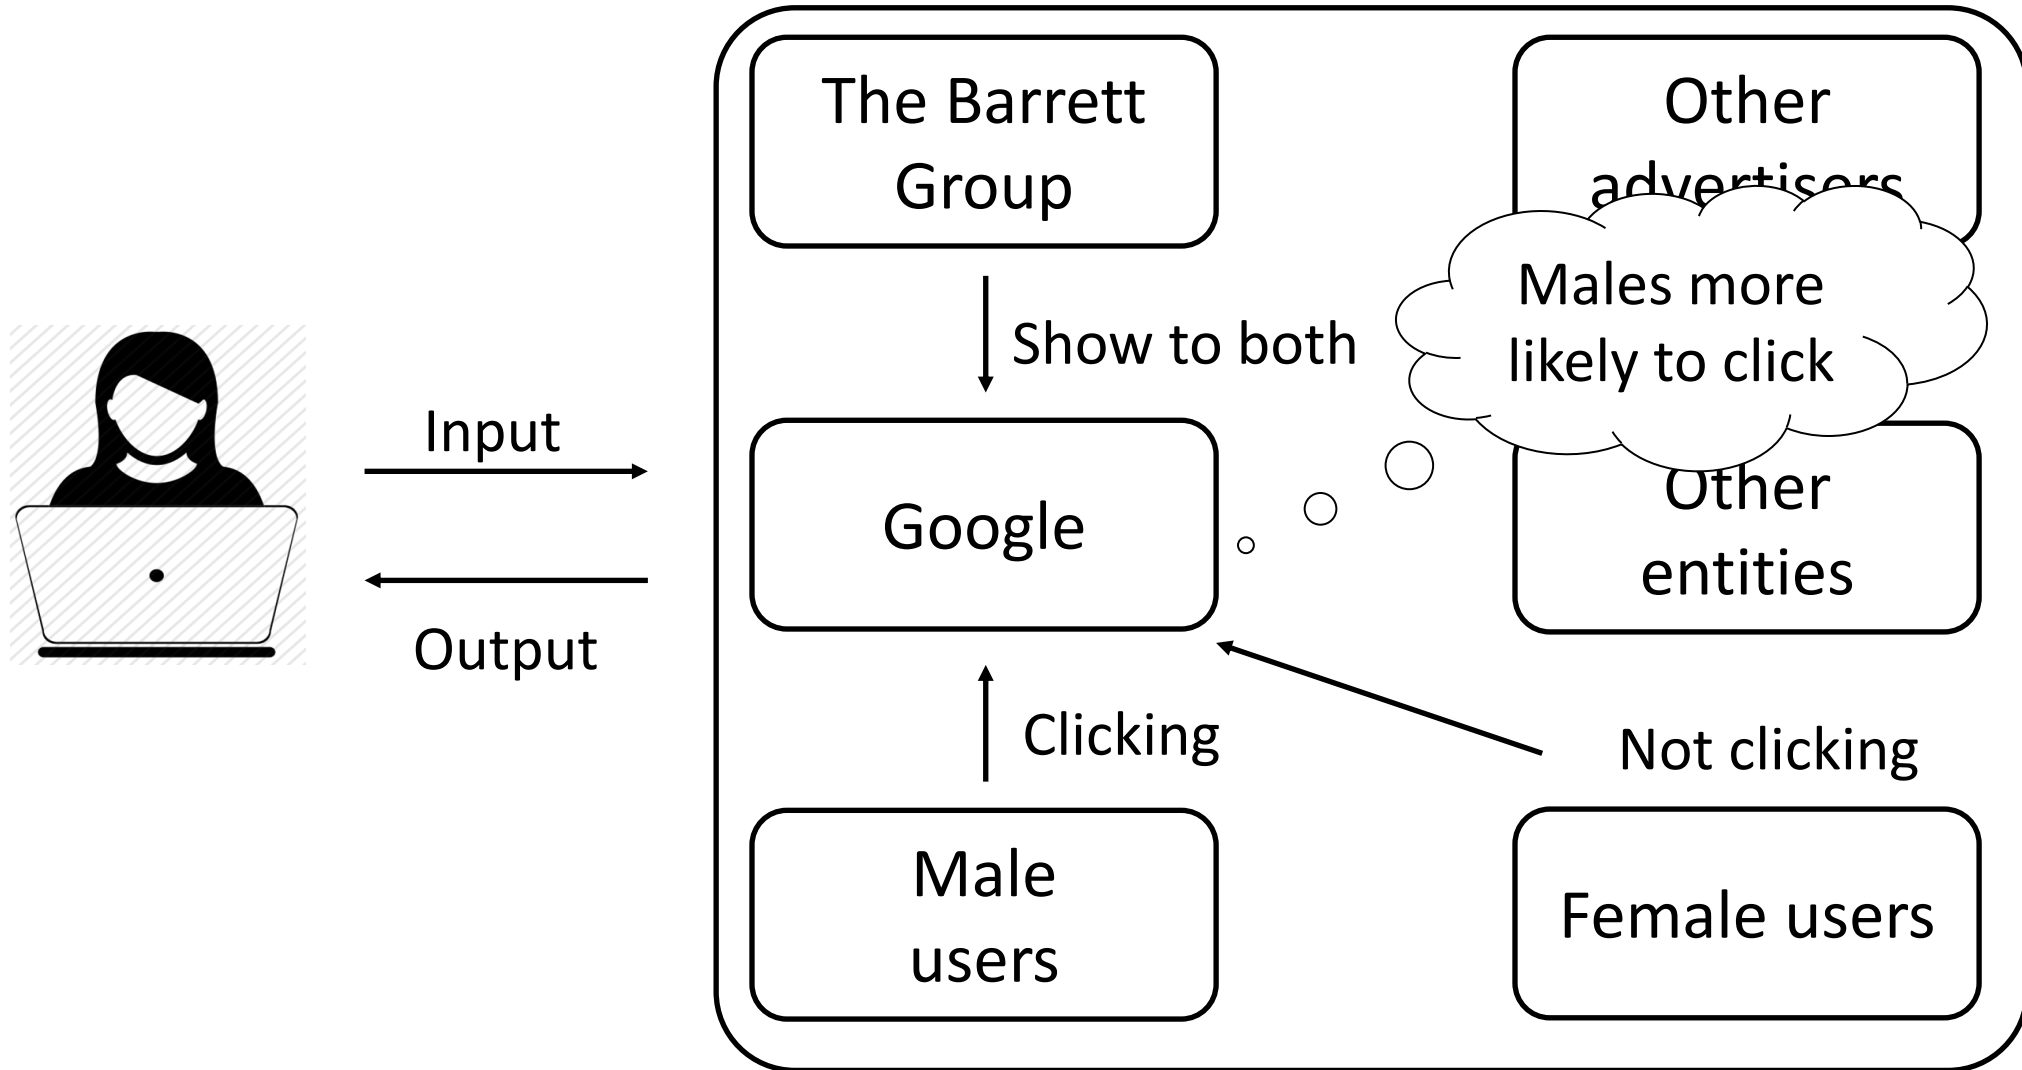
On October 19, 2020, Google will update its [Personalized advertising policies](#) to introduce new targeting restrictions. In an effort to improve inclusivity for users disproportionately affected by societal biases; **housing, employment, and credit products** or services can no longer be targeted to audiences based on **gender, age, parental status, marital status, or ZIP code**.

Examples of impacted products or services include, but are not limited to

- Housing: Housing listing sites, individual houses for sale or rental, real estate services
- Employment: Ads for jobs, job recruitment sites, job listing sites
- Credit: Credit cards, loans including home loans, car loans, appliance loans, short-term loans

Unclear how effective this will be: discrimination has many other causes

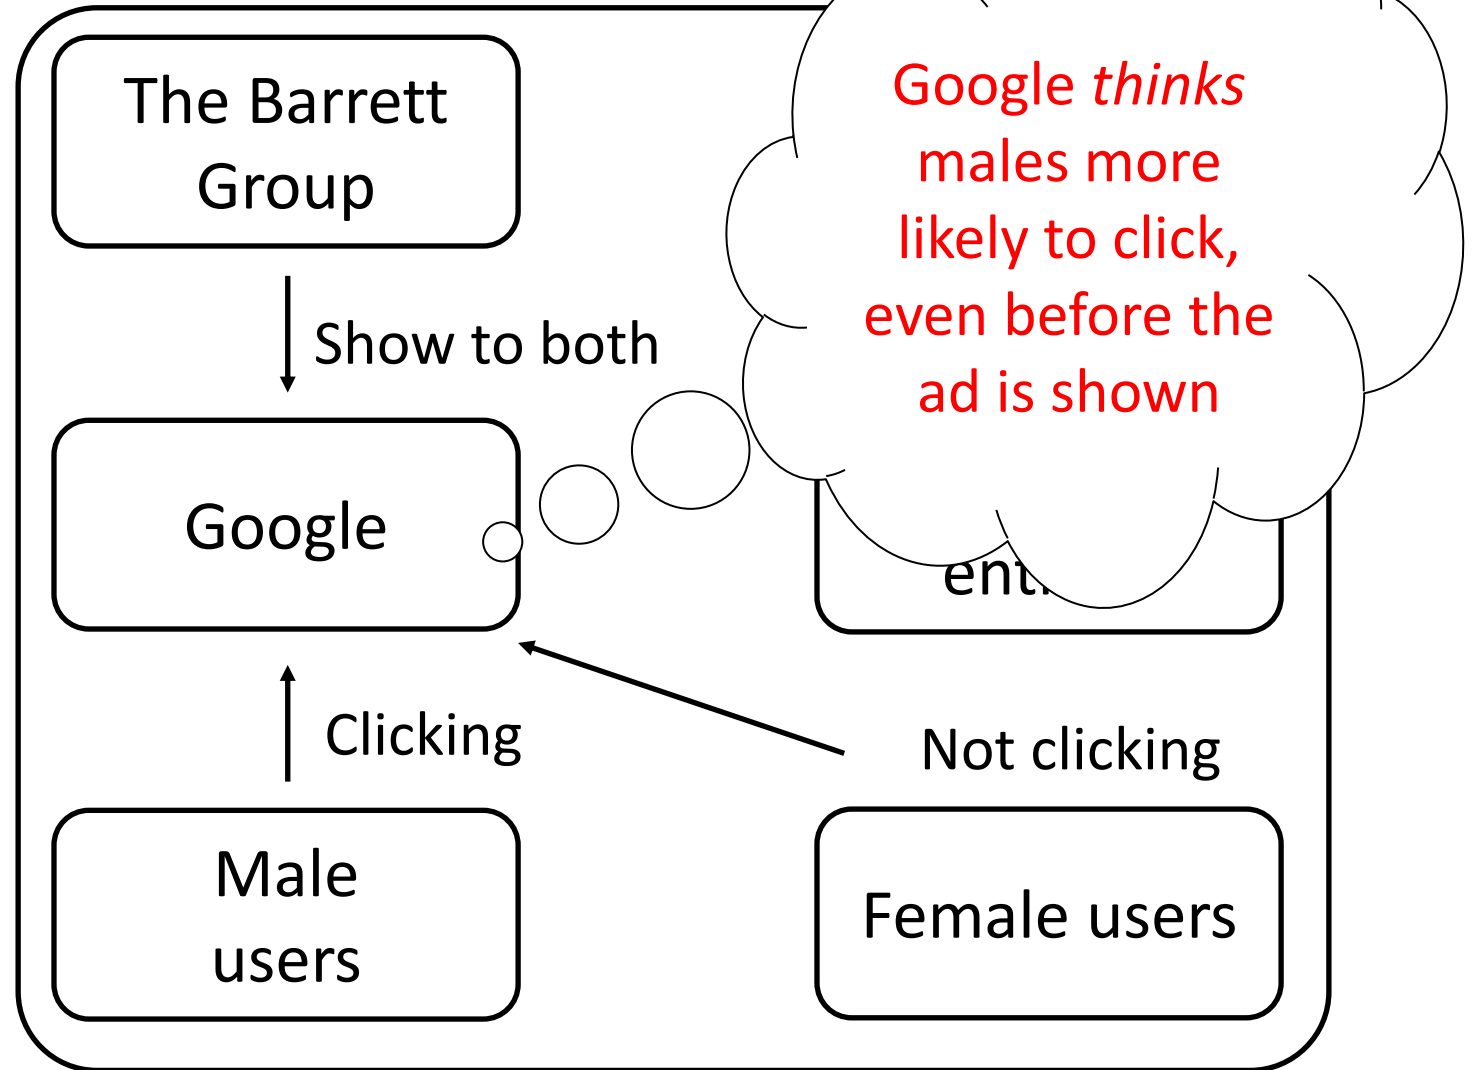
Possible cause: behavior of other users



Possible cause: behavior of other users




Input →
← Output



First experiment [Facebook]: Stereotypical (terrible) ads



 **Fashion Folk**
Sponsored - 



Try these 9 muscle gaining tips to combat your fast metabolism and achieve the mass you want




BODYBUILDING.COM
9 Killer Ways To Gain Muscle Naturally!
Tired of being known as the 'skinny guy'? Then try th...

2 Shares



 Like  Comment  Share




 **Fashion Folk**
Sponsored - 

Find out what essentials build the makeup kits of celebrity makeup artists.



ELLE.COM
How to Build a Makeup Kit, According to Three Celebrity Makeup Artists

 1  1 Comment

 Like  Comment  Share


Delivered to:

85% men

5% men


...despite *identical* targeting parameters

What leads to this disparity on Facebook? images



Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"



The skew in delivery is observable from the very beginning and stays constant!

Can skews arise in the **ad delivery** phase?

- what elements of the ad creative drive it? **Mostly the image.**
- is it driven by the performance of each ad, or decided a priori? **At least partially a priori.**
- is the decision made by humans or machines? **???.**

idea: make ads that look the same for humans but different for machines

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

height: 1080 px



width: 1080 px

R: 152
G: 135
B: 151

R: 232
G: 118
B: 4

R: 69
G: 49
B: 74

R: 69
G: 49
B: 74

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

height: 1080 px



width: 1080 px

R: 152
G: 135
B: 151
A: 255

R: 232
G: 118
B: 4
A: 255

R: 69
G: 49
B: 74
A: 255

R: 69
G: 49
B: 74
A: 255

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

height: 1080 px



width: 1080 px

R: 152
G: 135
B: 151
A: 128

R: 232
G: 118
B: 4
A: 128

R: 69
G: 49
B: 74
A: 128

R: 69
G: 49
B: 74
A: 128

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

height: 1080 px



width: 1080 px


R: 152
G: 135
B: 151
A: 3

R: 232
G: 118
B: 4
A: 3

R: 69
G: 49
B: 74
A: 3

R: 69
G: 49
B: 74
A: 3

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"



No difference in delivery between visible and invisible "male" ads

Similarly, no difference between visible vs invisible "female" ads

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"


Can skews arise in the **ad delivery** phase?

- what elements of the ad creative drive it? **Mostly the image.**
- is it driven by the performance of each ad, or decided a priori? **At least partially a priori.**
- is the decision made by humans or machines? **At least partially by machines.**

Slide Credit: Ali and Sapiezynski et al. "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes"

Same pattern with political ads!

Ran ads to balanced R/D audience, at same time, from same account, with same budget. Only difference was content of the ad:



Ad platform steers content towards “aligned” users

Slide credit: Ali, Sapiezynski, Korolova et al. “Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging”

Political ad delivery: Conclusions

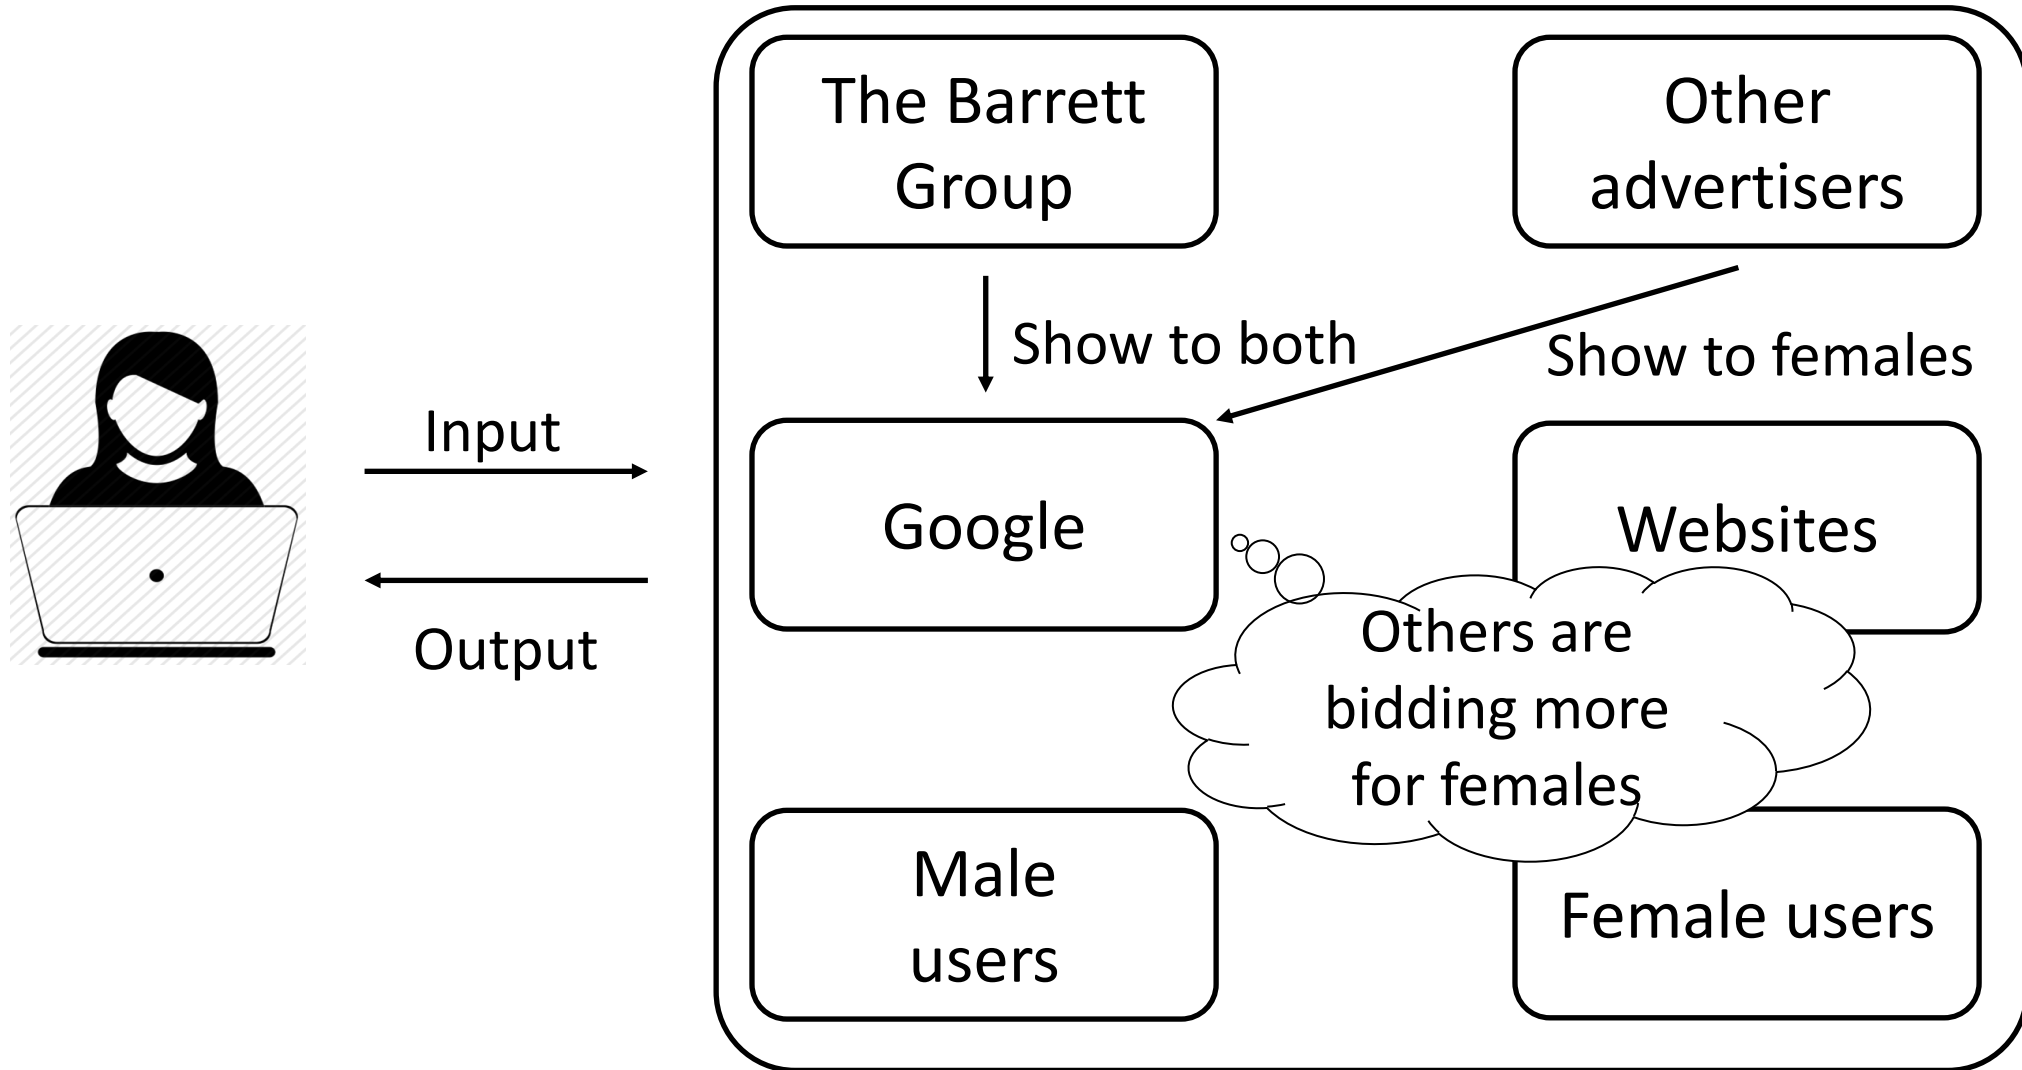
Facebook's political ad delivery is echo-chamber like

Facebook (non-transparently) puts its thumb on the scale of distribution of paid political messaging

Disabling microtargeting ability for campaigns gives Facebook more control over distribution of political messaging

Black-box auditing is costly and time-consuming!

Possible cause: being outbid by others



Some evidence of “competitive spillovers”

“Algorithmic Bias? An Empirical Study of Apparent Gender Based Discrimination in the Display of STEM Career Ads”

By Anja Lambrecht & Catherine Tucker

- Studied STEM career ads on Facebook in 191 countries
- Men indeed were more likely to see the Ads
 - ... even though Women were more likely to click on Ad conditional seeing it
- They attribute difference to “competitive spillovers” – women are more sought-after ad targets for retail
 - ...and so are more expensive targets
 - ...and so optimizing for price (budget) led to showing the ad more to men

Illegal? Unethical?

- In “*Discrimination in Online Advertising: A Multidisciplinary Inquiry*,” the authors in detail explain the (current) legal status
- Laws/Regulation actively being written about this stuff
- Need external audit capability
- Which of the above causes do you consider illegal/unethical? Who should be held responsible for each cause?

Discrimination in online marketplaces

Discrimination exists on online marketplaces

“Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment.” (2015). Benjamin Edelman, Michael Luca, and Dan Svirsky.

- Ran classic “resume-test” experiment on AirBnB
- “Guests with distinctively black names received positive responses 42% of the time, compared to 50% for white guests.”
- Hosts lost \$65 to \$100 of revenue for each black guest rejected.

“Racial and Gender Discrimination in Transportation Network Companies” (2016). Yanbo Ge, Christopher R. Knittel, Don MacKenzie & Stephen Zoepf

- “Across all trips, the cancellation rate for African American sounding names was more than twice as frequent compared to white sounding names. ”
- We also find evidence that drivers took female passengers for longer, more expensive, rides in Boston

Marketplace design can help (partially) mitigate discrimination

“When Transparency Fails: Bias and Financial Incentives in Ridesharing Platforms” (2020). Jorge Mejia & Chris Parker

- In response to findings about bias, platforms reduced “operational transparency through **removing information about riders’ gender and race from the ride request presented to drivers.**”
- **Does this work?** “Our results confirm that bias at the ride request stage has been eliminated. However, after acceptance, racial and LGBT biases are persistent, while we find no evidence of gender biases.”

AirBnB similarly pushed “instant booking” and partially removed guest names/profile pictures from the host portal

Special thanks to Dennis J. Zhang for slides/images

Ratings systems can also reduce discrimination

“Discrimination with Incomplete Information in the Sharing Economy: Evidence from Field Experiments on Airbnb” (2020) Ruomeng Cui, Jun Li, Dennis J. Zhang

