

Artificial Intelligence and Privacy

Catherine Tucker

Question

How will AI change the economics of privacy?

- ➊ Some general thoughts
- ➋ Specific example (joint work with Amalia Miller)

Varian 1996

Let us think about how privacy concerns enter a basic transaction. Suppose the seller has many different kinds of apples (Jonathan, Macintosh, Red Delicious, etc.) The buyer is willing to pay at most r to purchase a Jonathan, and 0 to purchase any other kind of apple.

The buyer will in general not want the seller to know r , the maximum price that he is willing to pay for the item being sold. If this information were available to the seller, the seller would price the product at the buyer's maximum willingness to pay, and the buyer would receive no surplus from the transaction.

How does AI (and lower costs of prediction) change this?

- ① data persistence
- ② data spillovers
- ③ data repurposing

How does AI (and lower costs of prediction) change this?

- ❶ data persistence: But our privacy preferences change over time leading to dynamics: (Goldfarb and Tucker, 2012)
- ❷ data spillovers: Genetic (Miller and Tucker, 2017) + Visual Data (Augmented Reality, Facial Recognition):
- ❸ data repurposing: Will talk about today

Historic Patterns of Racial Oppression and Algorithms

- with Amalia Miller

What predicts whether an algorithm identifies someone as being of a certain ethnic background?

Privacy debate has shifted to algorithmic-bias - easier to be persuasive about harm

Why might algorithms be biased?

- Biased Programmers
- Biased training data
- Bias is learned from humans

What we do

- Field Test data on Facebook
 - Set up job ads targeted at different ethnic affinities predicted by Facebook Algorithm
 - Find that you are more likely to be identified as African American (relative to population of African Americans) if you live in a state with a history of racial discrimination
 - Shows that history we deplore can end up influencing the predictions of algorithms today
 - Bias originates from data generating process behind the X used for prediction

More details about the experiment

Study this Ad for Federal Pathways program



Facebook and 'Ethnic Affinities'

Detailed Targeting ⓘ INCLUDE people who match at least ONE of the following ⓘ

Add demographics, interests or behaviors ⓘ | Suggestions | Browse

	▶ Education
	▼ Ethnic Affinity
Connections ⓘ	African American (US) <input type="checkbox"/>
	Asian American (US) <input type="checkbox"/>
	Hispanic (US - All) <input type="checkbox"/>
	Hispanic (US - Bilingual) <input type="checkbox"/>
	Hispanic (US - English dominant) <input type="checkbox"/>
ements	Hispanic (US - Spanish dominant) <input type="checkbox"/>

- Targeted at:
 - African-American
 - Asian- American
 - Hispanic (US-All)
 - People who are NOT African-American, Asian-American or Hispanic (US-All)
- We also wanted to match as closely as possible the projections of each race in each county from 2016 census projections. So also stratified by:
 - Gender
 - Age (20-24) (25-29)
- One week

As an Aside

The Civil Rights Act of 1964 also prohibits the 'printing or publication of notices or advertisements indicating prohibited preference, limitation, specification or discrimination' in employment recruitment.

What was supposed to happen



The targeting you chose may not be available for your ad

It looks like your ad may be for housing, employment or credit opportunities. These types of ads can't be run using multicultural affinity targeting. How to fix: Remove multicultural affinity targeting from your ad. If you believe your ad isn't for a housing, employment or credit opportunity, you can request a manual review.

[Request Manual Review](#)

- **Requiring self-certification:** When an advertiser attempts to show an ad that we identify as offering a housing, employment or credit opportunity and uses any other audience segment on Facebook, we will show the advertiser information about our updated anti-discrimination policy. We will then require the advertiser to certify that it is complying with that policy and with applicable anti-discrimination laws.



It looks like your ad may be for housing, employment or credit opportunities. If your ad is for a housing, employment or credit opportunity, click the button below to certify that you will comply with our policy prohibiting discrimination and with applicable anti-discrimination laws. If you believe your ad isn't related to a housing, employment or credit opportunity, you can request a manual review.

[Certify Compliance](#)

[Request Manual Review](#)

Your Certification

Your Certification

It is a violation of Facebook's Advertising Policies to discriminate based on personal characteristics such as race, ethnicity, color, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, medical or genetic condition. It may also violate federal, state, provincial, local or other applicable laws in your area or the area in which the ads will be shown.

Such laws may prohibit discrimination when running ads related to categories such as housing, employment or credit opportunities. For example, when running an ad for an apartment for rent, it may be illegal to exclude people who have children from that opportunity.

You agree that you have reviewed Facebook's [policy](#) prohibiting discrimination and will abide by those policies. You certify that you will not use Facebook advertising to improperly discriminate, and will comply with all applicable laws.



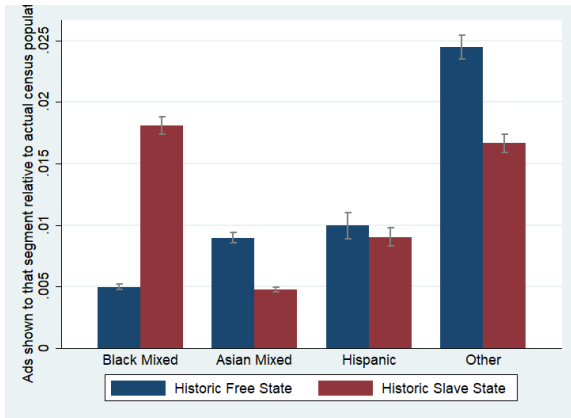
I acknowledge that I've read Facebook's policies and will comply with Facebook's policies and applicable laws.

[Cancel](#)

[Accept](#)

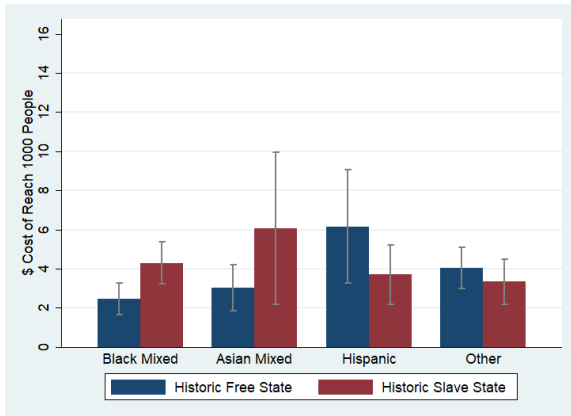
What we found

Show to far more people identified as African American in Historic Slave States than predicted by Census data



We looked for an explanation in the economics of advertising algorithms but not much evidence

This isn't driven by lower costs of advertising to African Americans in Historic Slave States



So what is driving this?

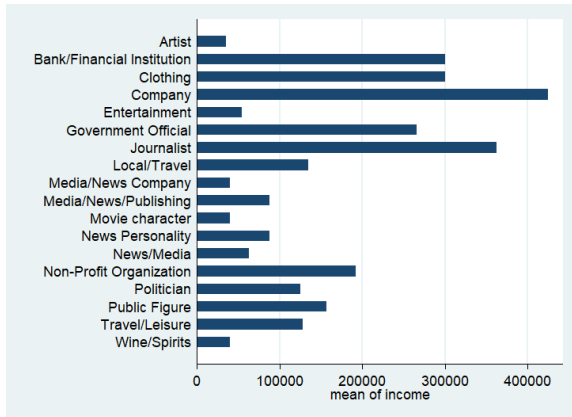
We think that three features of the algorithm are driving this.

Feature 1: Facebook uses 'liking' cultural artifacts to predict ethnic affinity.

'Multicultural Affinity' is 'the quality of people who are interested in and likely to respond well to multicultural content.'

- African-American: Liking 'Rasheeda'
- Hispanic: Liking 'Chiquis'
- Asian-American: Liking 'Manny Pacquiao'

Feature 2: Poorer People are more likely to like
'Cultural' objects - such as Celebrities on Facebook.
Rich people like Companies, Banks, and Journalists.



Feature 3: States with historic patterns of discrimination exhibit lower incomes for African Americans

Research in economics has suggested that African Americans are more likely to have lower incomes in states which have exhibited historic patterns of discrimination (Sokoloff and Engerman, 2000; Bertocchi and Dimico, 2014).

Further Tests

- Effect replicates for other historical wrongs (voting rights)
- Tests controlling for income at county level provide support (though facing some missing data challenges)
- There may of course be additional explanations (perhaps liking 'ethnic-cultural' entities is also a function of ethnic concentration.)

Punchline

- Algorithm designed to predict ethnic affinity appears to be influenced by historic injustice
- Not because algorithm responds to click behavior or costs
- Instead because
 - ① Algorithm predicts ethnic affinity based on whether you like certain 'cultural objects'
 - ② People with lower incomes more likely to like 'cultural' objects. Higher incomes engage with news media
 - ③ People who live in states with histories of injustice are more likely to be lower income

Implications

- New explanation for apparent bias: Algorithms don't think about the data generating process behind the X.
- Matters because if you choose to exclude on the basis of predicted ethnic affinity will end up excluding who have historically been most excluded
- Also shows issues facing economics of privacy regarding 'data ownership': It isn't clear that when someone likes 'Fresh off the Boat' or writes a message in Spanish they are aware that may be used to predict their ethnicity (or other things) in a world of AI.

Big Picture

Traditional models of the economics of privacy need to also reflect

- Data persistence and dynamics implied
- Data repurposing and uncertainty implied
- Data spillovers and lack of control or choice implied

Thank you!

cetucker@mit.edu

Bertocchi, G. and A. Dimico (2014). Slavery, education, and inequality. *European Economic Review* 70, 197–209.

Goldfarb, A. and C. Tucker (2012). Shifts in privacy concerns. *American Economic Review: Papers and Proceedings* 102(3), 349–53.

Miller, A. and C. Tucker (2017). Privacy protection, personalized medicine and genetic testing. *Management Science*.

Sokoloff, K. L. and S. L. Engerman (2000). History lessons: Institutions, factors endowments, and paths of development in the new world. *The Journal of Economic Perspectives* 14(3), 217–232.