CONSUMER INJURIES AND BENEFITS

IN THE DATA-DRIVEN ECONOMY

CONSUMER INJURIES AND BENEFITS

IN THE DATA-DRIVEN ECONOMY

Panel One: Quantifying Injuries & Benefits to Consumers (Part One)

Quantifying Injuries & Benefits to Consumers

Avinash (Avi) Collis

Assistant Professor

Heinz College of Information Systems and Public Policy

Carnegie Mellon University

www.avinash.info

Quantifying injuries to consumers

1. Measuring disutility from targeted online ads on social media platforms
   - The Consumer Welfare Effects of Online Ads: Evidence from a 9-Year Experiment (with Erik Brynjolfsson, Daniel Deisenroth, Haritz Garro, Daley Kutzman, Asad Liaqat, and Nils Wernerfelt). ***American Economic Review: Insights*** 2025.

2. Measuring how much users value their data on social media platforms and if they update their valuations if they are informed about data breach settlements
   - Information Frictions and Heterogeneity in Valuations of Personal Data (with Alex Moehring, Ananya Sen, and Alessandro Acquisti), Working Paper 2025.

3. Measuring users' demand for privacy from data brokers
   - Demand for Privacy from Data Brokers (with Joy Wu and Ananya Sen), Work in progress.

Measuring disutility from targeted online ads on social media platforms

- Facebook's internal A/B testing platform launched in 2013, holdout group maintained continuously since then!
- 0.5% of all users (over 3 billion) randomly assigned to the holdout group and don't see ads (and they don't know this!)
- We recruit a representative sample of users from the no-ads and the ads groups, and measure their Facebook valuations (Willingness to Accept to give up Facebook for 1 month)
- Any differences in valuations should reflect utility/ disutility of ads in the long term

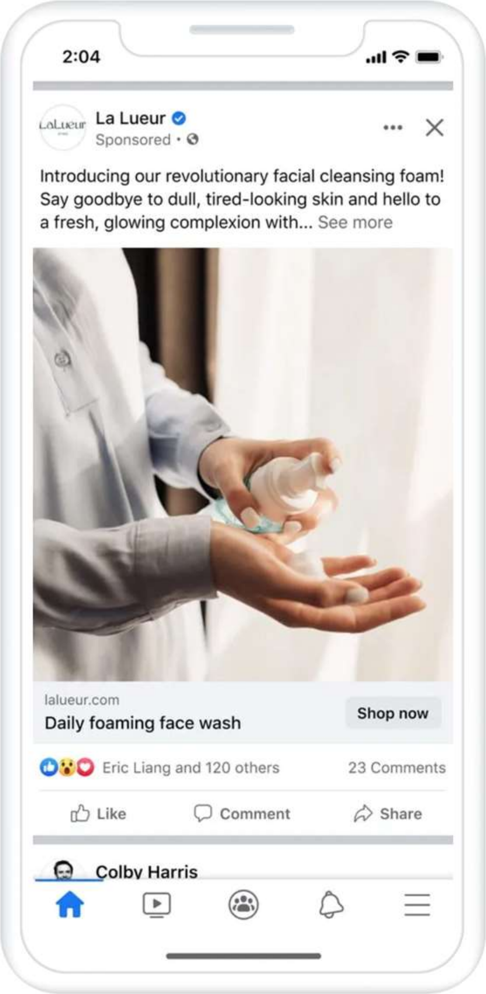

Introducing our revolutionary facial cleansing foam! Say goodbye to dull, tired-looking skin and hello to a fresh, glowing complexion with... See more

lalueur.com

Daily foaming face wash

Users in the ads group spend 9.4% less time

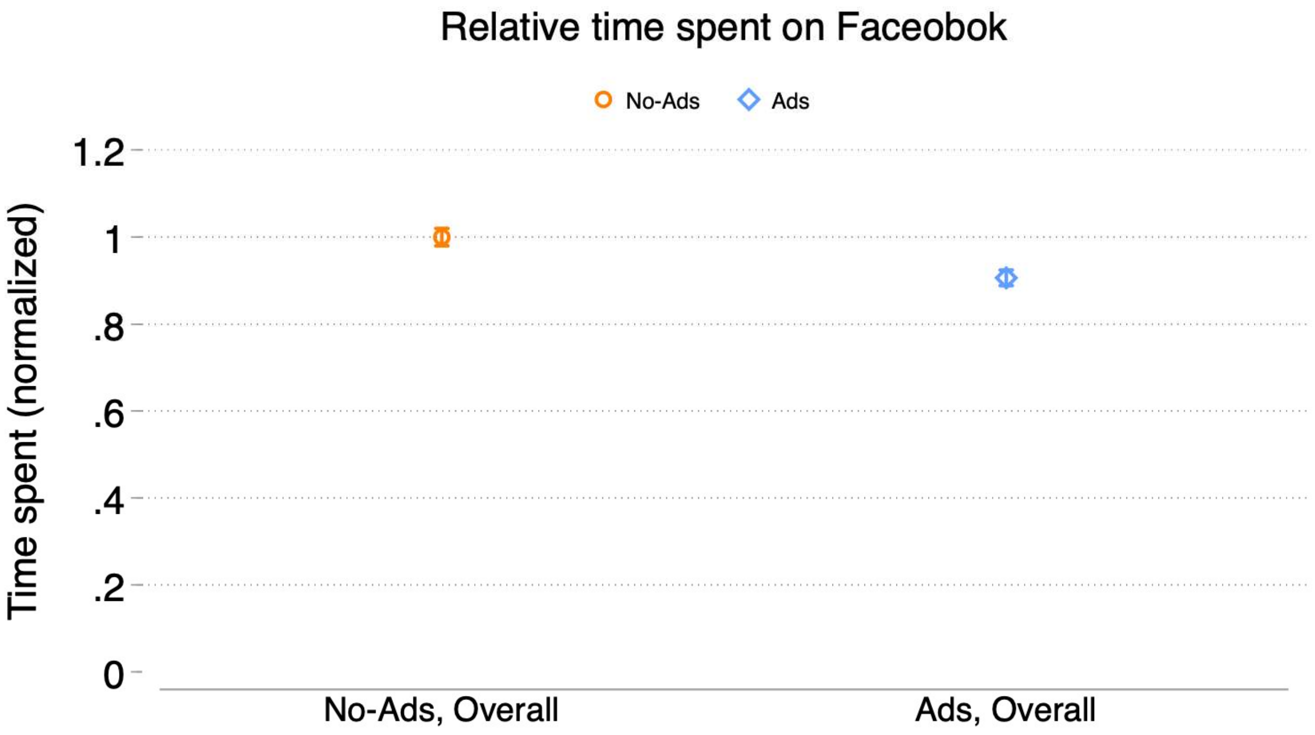

No significant differences in valuations!

Control (no-ads) WTA = $31.95

Treatment (ads) WTA = $31.04

Difference = -0.9 [-3.06, 1.25]

Minimum detectable difference = **$3.18/month** (~10% of FB value)

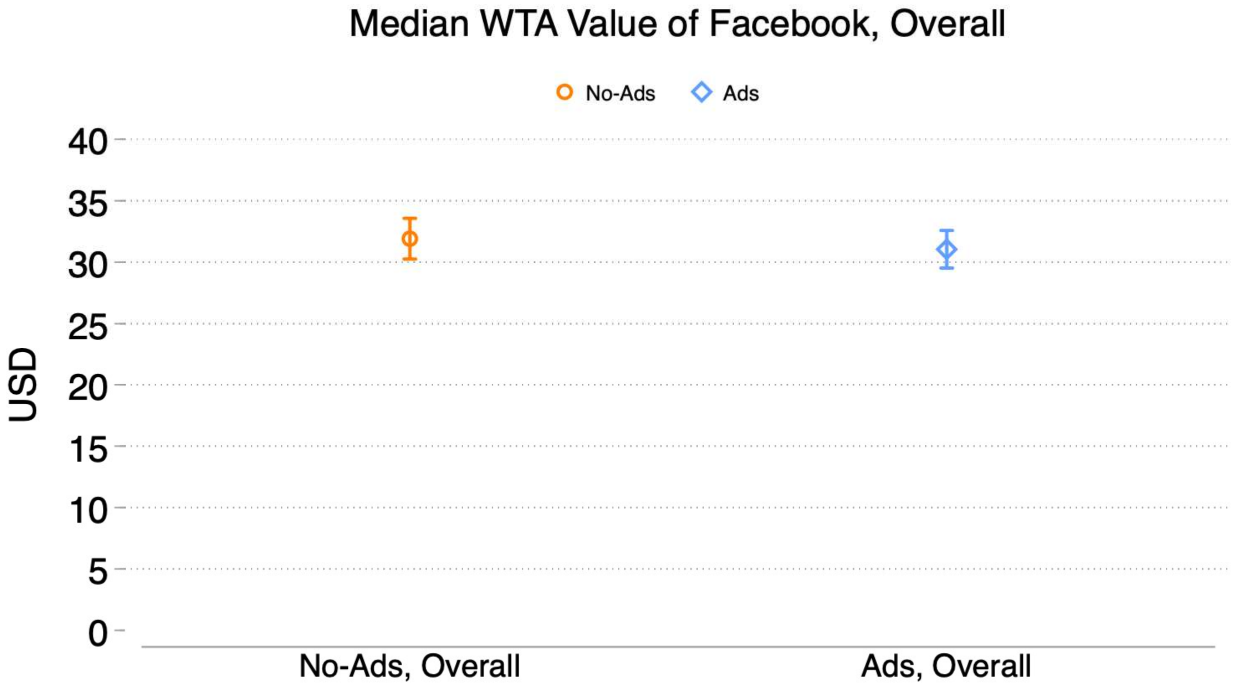

Measuring how much users value their data on social media platforms, and if they update their valuations if they are informed about data breach settlements

- Experiment: Create markets for personal data

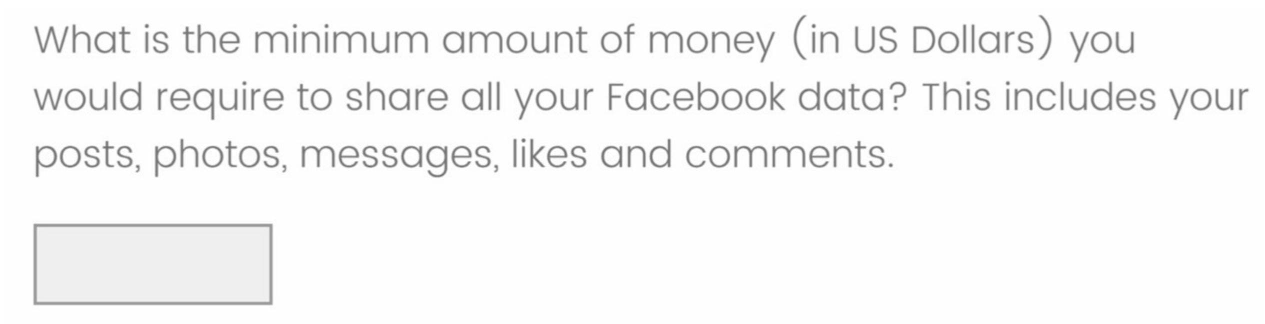

What is the minimum amount of money (in US Dollars) you would require to share all your Facebook data? This includes your posts, photos, messages, likes and comments.

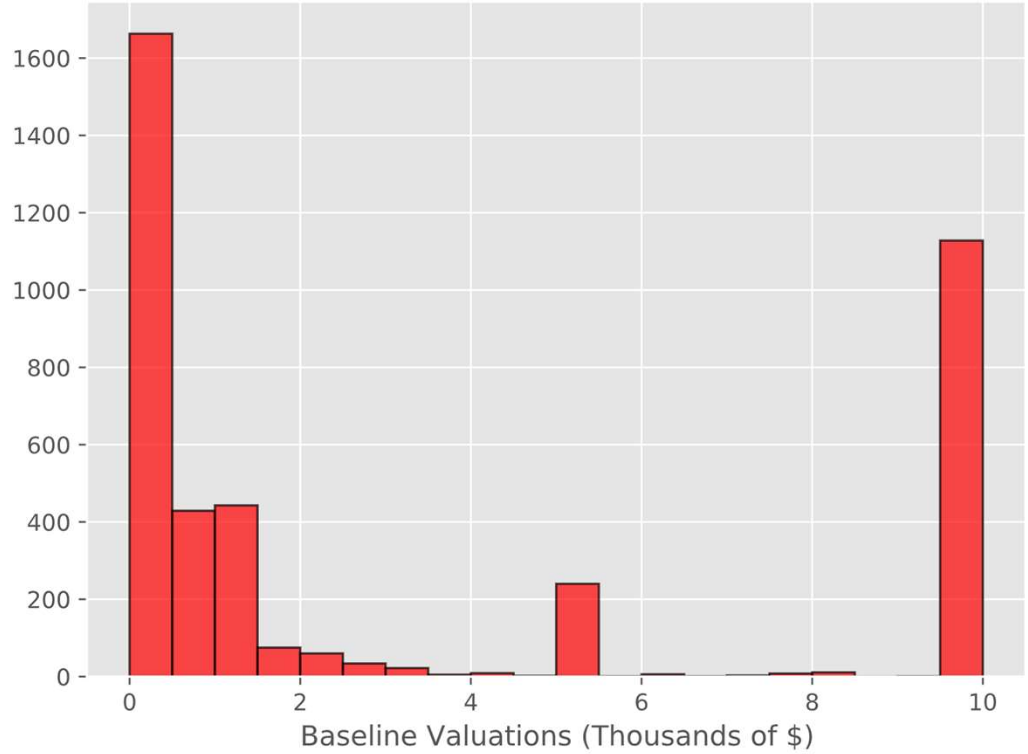

The figure shows a histogram with the distribution of valuations at $250 intervals.

Median is $750 with 18% less than $100 and 25% greater than 10000.
High valuations an expression of unwillingness to part with data

Information Treatments

Do valuations change when users are provided information about legal settlements?

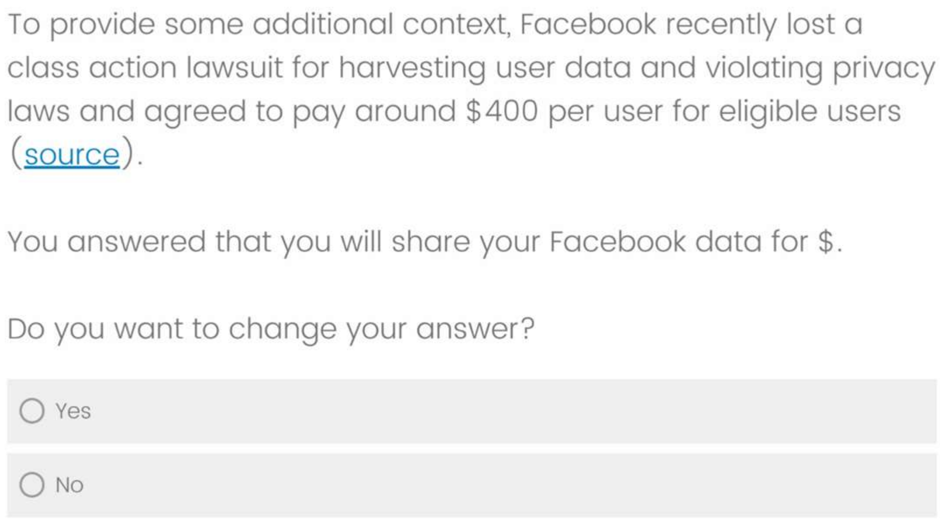

To provide some additional context, Facebook recently lost a class action lawsuit for harvesting user data and violating privacy laws and agreed to pay around $400 per user for eligible users (source).

You answered that you will share your Facebook data for $.

Do you want to change your answer?

○ Yes

○ No

Revision of valuations

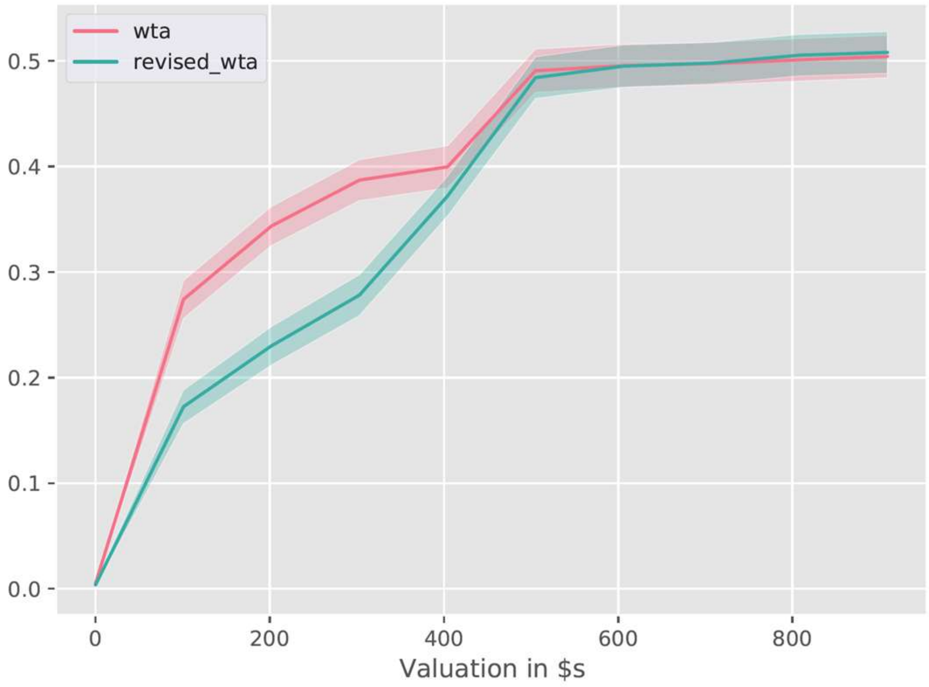

- 55% of users update their valuations
- Revisions happen below $400, making the distribution less dispersed.
- Follow up study shows significant uncertainty about the data breach lawsuit

Demand for Privacy from Data Brokers

- Experiment: Elicit privacy preferences when individuals become aware of data brokers' activities that involve harvesting government or commercial-sourced personal data
- Information interventions informing users about government or commercial data sources
  - Interventions borrowed from the FTC 2014 report on data brokers
- Key outcome variables:
  1. Beliefs about privacy from data brokers
  2. Incentive compatible willingness to pay (WTP) to delete data from data brokers
     - We buy users a data deletion service with some of their earnings from the experiment

Results

- 40% of Americans are willing to pay an average of $25 to maintain their privacy from data brokers for one year
- Information intervention informing users about government or commercial data sources shifts beliefs about data leakage to brokers, but does not change WTP for data broker deletion services

Conclusion

- These studies show examples of measuring injuries (and benefits) to consumers using experiments (online and in the field)
- Rather than focusing on absolute $$ figures, comparing incentivized valuations for platforms, or data, or other aspects of the digital economy (e.g. network effects) across different groups could be useful for these measurement exercises

Thank you

www.avinash.info

Informational Injury Workshop

Catherine Tucker

Agenda

Challenges to Studying Privacy as an Economist

First Steps

What is Privacy?

My Favorite Definition: Freedom from Unwarranted Intrusion

However, for economists, privacy may be either a ‘taste for privacy’ or a concern that data, as an input, can have adverse economic consequences.

What Started This All

HARVARD

LAW REVIEW.

VOL. IV. DECEMBER 15, 1890. NO. 5.

THE RIGHT TO PRIVACY.

> "It could be done only on principles of private justice, moral fitness, and public convenience, which, when applied to a new subject, make common law without a precedent; much more when received and approved by usage."
>
> WILLES, J., in Millar v. Taylor, 4 Burr. 2303, 2312.

THAT the individual shall have full protection in person and

Therefore Concepts of Informational Injury Should be Dynamic

- Privacy as a concept is bound up in technological change.
- Brandeis and Warren’s seminal article was written in response to the portable camera and the new privacy concerns it created.
- In 2026, we are facing unprecedented digital change and more data than ever.
- This means that a static approach to measuring informational injury will always be fraught.

Agenda

Challenges to Studying Privacy as an Economist

First Steps

Any Measure of Informational Injury is Contextual

- I like my privacy framework where I point out that we should be most worried in our algorithmic era about data that:
  1. Has huge economic consequences
  2. Has spillovers
  3. Has persistence
- For example, if my genomic data is released, it could affect my ability to obtain health insurance, as well as my sister's ability to obtain health insurance, decades into the future.

Can we identify cases where a ‘taste for privacy’ is economically rational?

Informational Injury Due to the Spread of Data Associated With Unfounded Stigma

- Mental Health
- Reproductive Health
- Past Crimes

Informational Injury From the Spread of Data Associated With Addiction

- Health
- Spending
- Gambling

Thank You

cetucker@mit.edu

CONSUMER INJURIES AND BENEFITS

IN THE DATA-DRIVEN ECONOMY

Panel Two: Quantifying Injuries & Benefits to Consumers (Part Two)

The Social Costs of Privacy Laws

- Telematics in Auto Insurance
- Facial Recognition in Law Enforcement

Telematics in Auto Insurance

- Tracking: Embedded devices record real-time driving patterns:
  - Speeding, miles driven, hard breaks, aggressive turns, unsafe following, abrupt lane changes, distractions (texting), location, night driving . . .
- Prediction: A personalized safety score is generated by the insurer's algorithm. The score adjusts in real time.
- Premiums: Insurance premiums reflect each policyholder's safety score and adjust periodically (monthly).
- Enrollment: Optional; usually with upfront discount

  22% market share (2022)

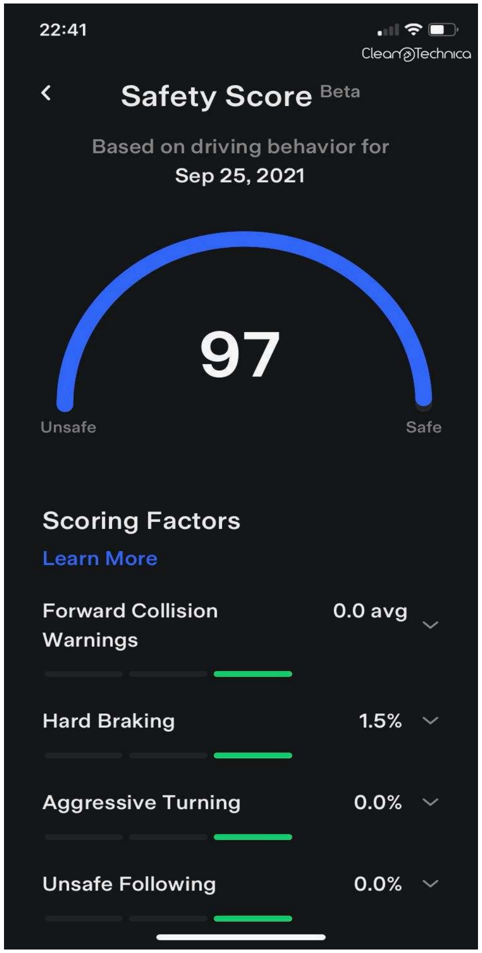

Safety Score Beta

Based on driving behavior for
Sep 25, 2021

97

Scoring Factors

Learn More

Forward Collision Warnings 0.0 avg

Hard Braking 1.5%

Aggressive Turning 0.0%

Unsafe Following 0.0%

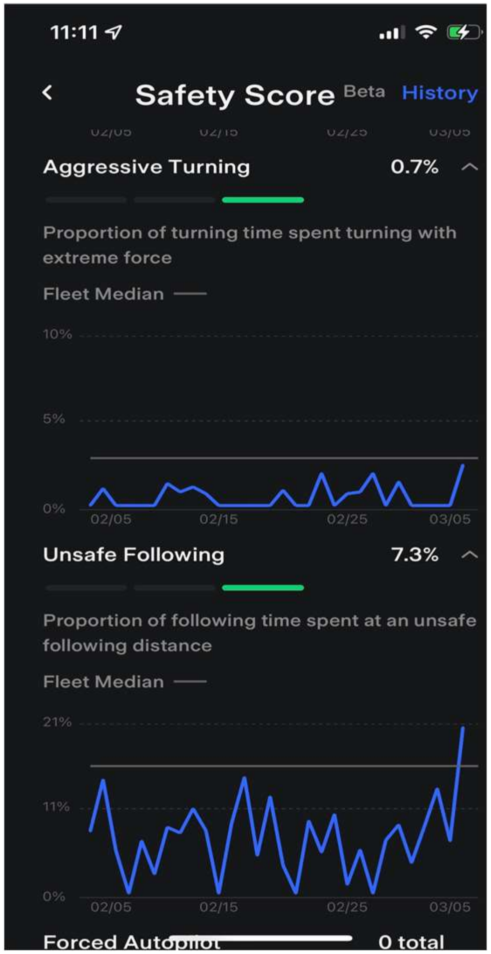

Safety Score Beta History

Aggressive Turning 0.7%

Proportion of turning time spent turning with extreme force

Fleet Median

Unsafe Following 7.3%

Proportion of following time spent at an unsafe following distance

Fleet Median

Forced Autopilot 0 total

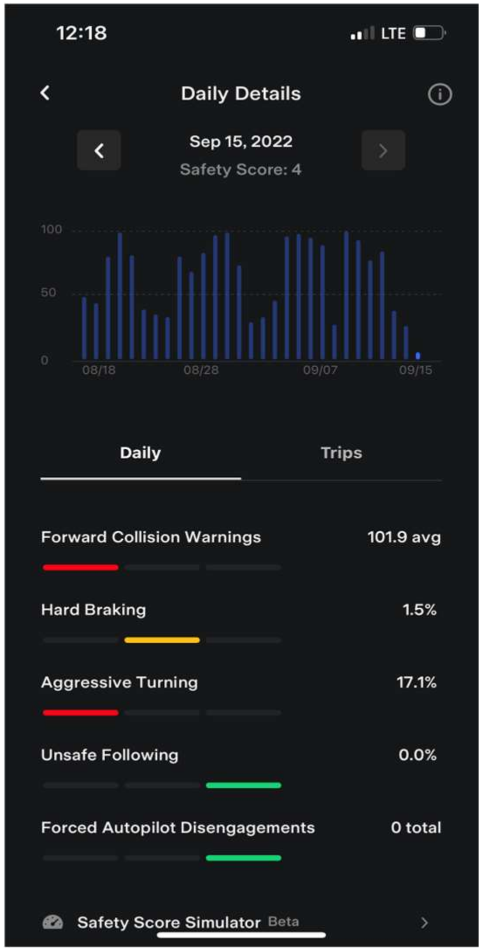

Daily Details

Sep 15, 2022
Safety Score: 4

Daily Trips

Forward Collision Warnings 101.9 avg

Hard Braking 1.5%

Aggressive Turning 17.1%

Unsafe Following 0.0%

Forced Autopilot Disengagements 0 total

Safety Score Simulator Beta

The Benefits (1): Safety

Theory:

- Rewards: Premium discounts incentivize safer driving
- Coaching: Drivers receive real-time alerts when engaging in dangerous maneuvers
- Being "watched": Drivers are aware that they are being measured, triggering more deliberate driving decisions

The Benefits (1): Safety

Reduction of up to 30% in fatal accidents:

- Reimers & Shiller (JLE 2020): roughly <u>50% reduction in accidents</u>
  - Data on early adoption of "Snapshot" by Progressive
- Jin & Vasserman, (NBER 2021): Opting in to telematics makes drivers <u>30% safer</u>
  - Data on safety score of enrolled drivers over a period of time
- Soleymanian et al, (MS 2019): <u>21% reduction in hard breaking</u>
  - Data on driving patterns of participants and non-participants, over a period of time
  - Only 1% of policyholders exhibited no improvement in driving
  - Improvement occurred over the first month, and persisted afterwards

The Benefits (2): Fairness

- Lower premiums
  - Early estimate (Progressive.com): 12% average discount (not due to adverse selection)
  - Discounts are greater when more insurers enter UBI market
  - Discounts are personalized, based on safety score
- Reduced reliance on non-driving social-demographic rating factors
  - Eliminates factors like credit score, home ownership, and education
- Actuarial justice: no group-based rating
  - E.g., gender
  - People who drive more miles pay higher premiums

The Law

- California prohibits auto insurers from using telematics to price insurance
  - “an insurer shall not use a technological devise to collect or store information about the location of the insured vehicle.”
  - California Insurance Commissioner: “We won't bend on protecting consumer data, privacy, and fair rates.”
- Some states (e.g., NY) restrict how telematics data are collected and used
  - Prohibited factors: “A company may collect distracted driving statistics; however, such statistics may not be used in the algorithm to determine the final score”
- Massive literature in law, sociology, and policy on the privacy loss resulting from telematics; zero interest in the benefits

California Data Privacy Protection

4000 annual auto accident fatalities in California

30% reduction = 1200 lives

Facial Recognition: The Technology

AI-powered computer vision technology

- Private facial recognition software services are trained on billions of publicly accessible images of humans
- Detect, recognize, classify, and extract features from people’s images
- Used in law enforcement to identify individuals: compare an image of a victim or suspect to a reference set of images

Facial Recognition: The Benefits

- Enforcement against human trafficking
  - 28 million people are in forced labor and sex trafficking, 20% are children
    - One of the most difficult crimes to prevent and solve
  - Opportunities to identify: during travel; images posted online
- 21 State Attorney Generals: FRT has been used by law enforcement to identify 18,000 trafficking victims and 6,000 traffickers.
- Thorn developed a non-profit tool—”Spotlight”—used by law enforcement in 40,000 cases, where investigators rescued more than 9,000 children, and arrested 10,000 traffickers.
- An NSF-funded company developed with Carnegie Mellon’s Robotics Institute a tool matching social media data on missing persons with online escort ads. Within two years, it contributed to the identification of 6,800 sex trafficking victims, saving 70,000 investigative hours per year

Facial Recognition: Privacy Laws

EU AI Act: “Real-time remote biometric identification” which recognizes people at a distance is “unacceptable risk” and is banned.

- Exceptions allowed in investigation of serious crimes, “targeted in terms of the individuals to be identified, the location, [and] temporal scope” + individually approved by a judge

Privacy Watchdog: “the exchange of freedom and privacy for some early anecdotal evidence that it might help some people is wholly insufficient to trade away civil liberties.”

Sept 2024: Dutch government fines Clearview AI $33 million for violation of GDPR (failure to get consent from people for use of their images in training the algorithm)

Local bans on the use of the software by law enforcement in U.S. cities

- Also, significant restrictions on collecting images data to train the system

EU Data Protection “Czar”:

“Given the nature of the personal data at stake—sensitive biometric data—and that vulnerable people may be involved—migrants” more attention must be paid to the “impact on the fundamental rights to privacy and data protection…”

Facial Recognition: “Data Minimization” Laws

Big concern: disparate accuracy across racial groups

- Cases in which black men were mistakenly identified and arrested

Why?

- “Other race effect” due to smaller number of images of members of minority groups in the training phases.
- Data minimization laws restrict the ability to tag images demographic classifications, slowing down the recognition of group membership and the pace of intra-group learning. It also eliminates the ability to audit the algorithm for inter-group fairness.

This problem is gradually lessening:

- Kashmir Hill (NYTimes): “the window of time for [the racial disparity] criticism to be effective is closing as top developers have focused on addressing the problem of biased algorithms.” Privacy advocates worry that the “greater accuracy across diverse groups” would be used “as a justification to deploy the technology more widely.”

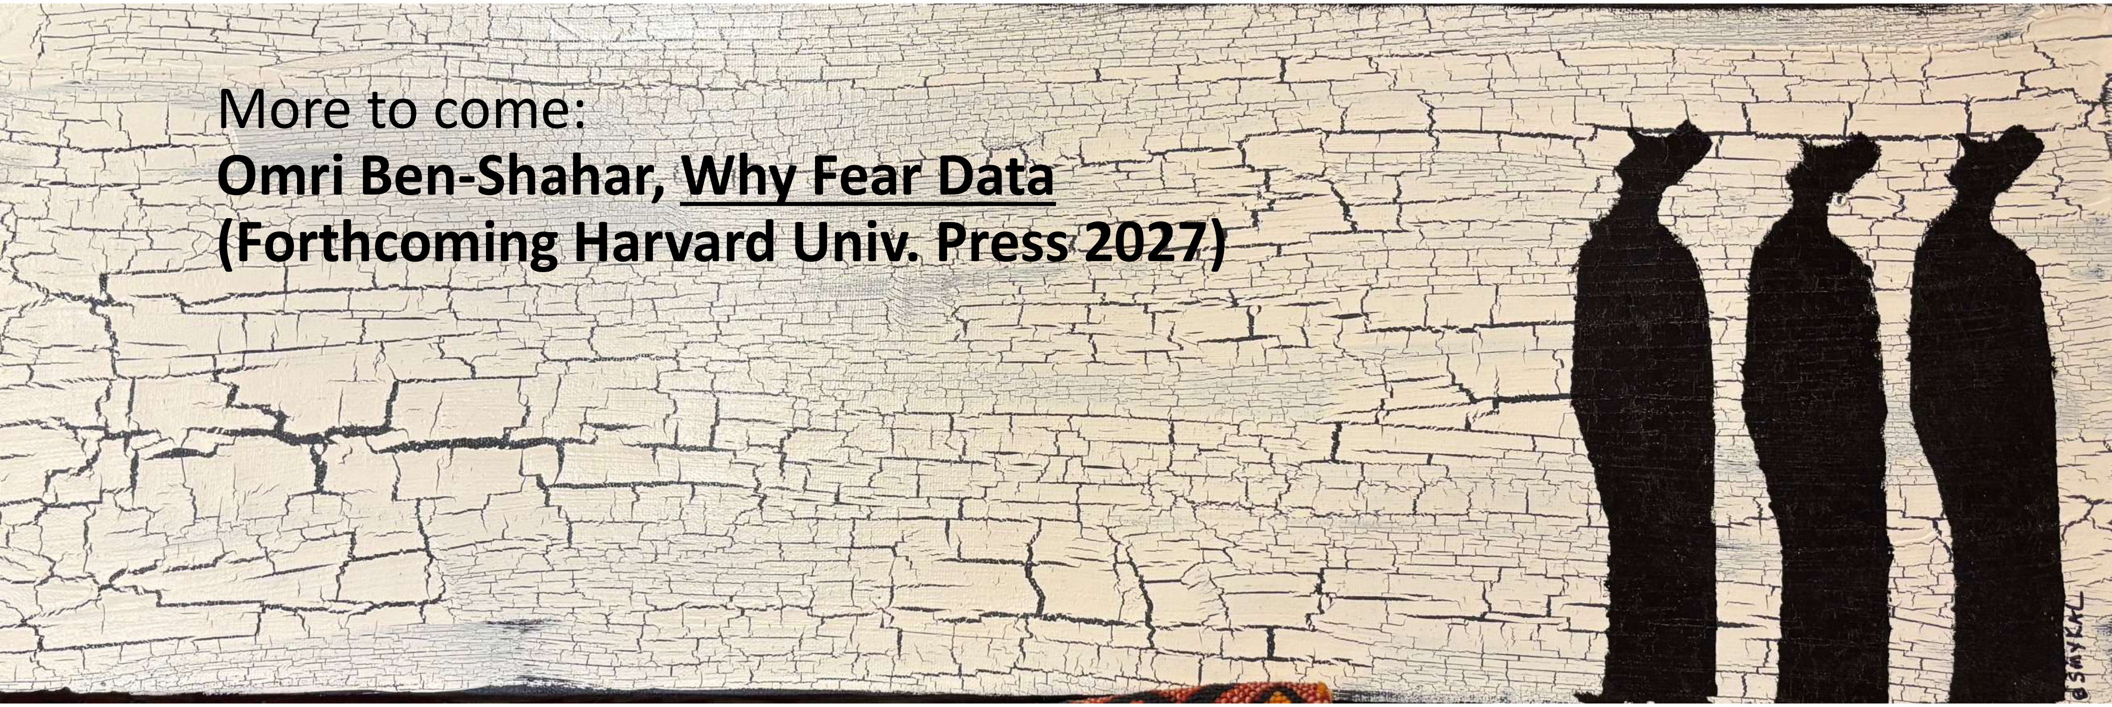

More to come:

Omri Ben-Shahar, Why Fear Data
(Forthcoming Harvard Univ. Press 2027)

Privacy and Data Security Economics at the FTC: Differences Across Consumers in Values for Privacy

Daniel H. Wood

February 26, 2026

The views expressed in this presentation are those of the author and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner.

Much of this talk taken from “The Costs and Benefits of Digital Privacy Protections” joint with James Thomas

FTC Consumer Protection Economics

- Deception
  - Question: what is injury?
- Unfairness & Policy Analysis
  - Question: do costs outweigh benefits?

Answering Questions: Values

- Consumer values necessary
  - Deception: value of misrepresented attribute
  - Cost benefit: total value of policy / practice to consumers
- Measured by
  - Willingness to pay (WTP): $ consumer willing to pay to get something
  - Willingness to accept (WTA): $ consumer willing to receive to give up something

Values in Privacy Economics

4 Problems

- Context
- Asymmetric information
- Measurement
- WTA versus WTP

These Remarks: Heterogeneity in Values

5th Problem

- For a given privacy tradeoff, WTP or WTA varies across consumers
- Important fact for policy
- Opportunity for research

Value Concentrated in People With High Values

- Much of the value of privacy concentrated in people with high values
- Example: Lin and Strulov-Shlain (2024)
  - Incentivized measurement of WTA to share several forms of Facebook data
  - 20% report WTA > $100, ≈ 2x median WTA
- Conservatively 50% of total value of privacy captured by 25% of people with highest values

Analyzing Costs and Benefits of Novel Privacy Protections

High Value Consumers

- Most of social benefits of privacy protections located in high value consumers
- Behavior of high value consumers has outsized importance for benefits analysis:
  - Baseline: Do they use existing privacy-enhancing technologies or behaviors? Do they understand existing privacy choices accurately?
  - Counterfactual: Would they use new privacy protections at high rate?

GDPR and High Value Consumers

- People who opt out under GDPR substituting away from other privacy protective behaviors (Aridor, Che, and Salz RAND 2023)
  - Counterfactual: high value consumers do use new privacy protections
  - Baseline: but high value consumers did have some privacy protection before

Targeted policies

GDPR, ATT, universal opt out

- Targeted policies can capture most of benefits and avoid some costs
- GDPR: consent banners impose large time costs (Farronato, Fradkin, and Lin 2024)
- ATT: 80% take-up largely eliminates any social benefits of behavioral advertising from Apple users
- Universal opt out avoids both these costs

Deception Harm

- Calculations of harm that assume every consumer has mean WTA will underestimate harm (Wood and Stone 2018)
- Example
  - Platform
    - Platform offers service for $p = 0$
    - Claims it will not share user data with 3rd parties
    - In fact it sells user data to a data broker
  - Consumers
    - Values v for service, that vary across consumers
    - WTA c for sharing data

Deception Harm

Quantifying injury

- Consumers with $v < c$ use service, but only because deceived
- Share of consumers harmed is share of consumers with v in $(0, c)$
- Injured consumers suffer injury between 0 and c

Deception Harm

Comparison

- Homogeneous WTA: $c = 10$
  - Injury = (Mean $v \in (0, 10)$) $*$(Share of $v \in (0, 10)$)
- Heterogeneous WTA: 50% of users have $c = 20$, others have WTA 0
  - Injury = 0.5 $*$([Mean $v \in (0, 20)$) $*$(Share of $v \in (0, 20)$)
- With v uniform, injury with heterogeneous WTA is 2x that calculated with homogeneous
- Intuition: high-WTA consumers more likely to be marginal, and if marginal suffer more injury
- Heterogeneity in privacy values implies higher privacy deception harm

CONSUMER INJURIES AND BENEFITS

IN THE DATA-DRIVEN ECONOMY

Panel Three: Data Breaches, Impacts on Consumers, & Efforts to Minimize Injuries

CONSUMER INJURIES AND BENEFITS
IN THE DATA-DRIVEN ECONOMY

Panel Four: The Costs and Benefits of Behavioral & Contextual Advertising

The Costs and Benefits of Behavioral and Contextual Advertising

Cristobal Cheyre | Assistant Professor

Information Science

Cornell Ann S. Bowers College of Computing and Information Science

Cornell University

Two Theoretical Frameworks

Theory 1: Cost-Reduction Framing

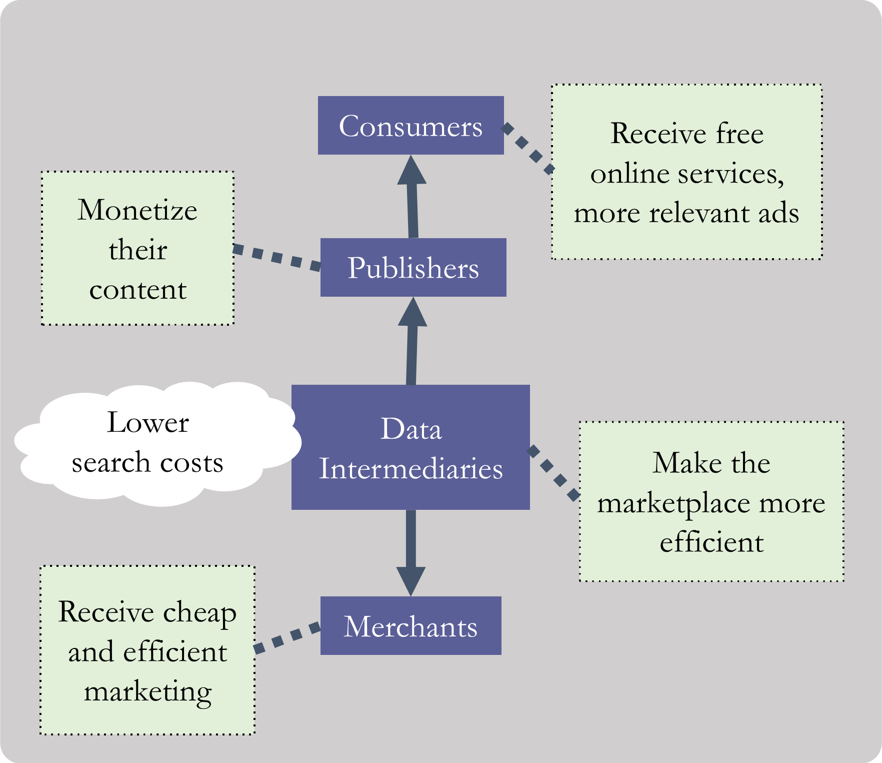

Consumers

Receive free online services, more relevant ads

Monetize their content

Publishers

Lower search costs

Data Intermediaries

Make the marketplace more efficient

Receive cheap and efficient marketing

Merchants

Theory 2: Surplus-Extraction Framing

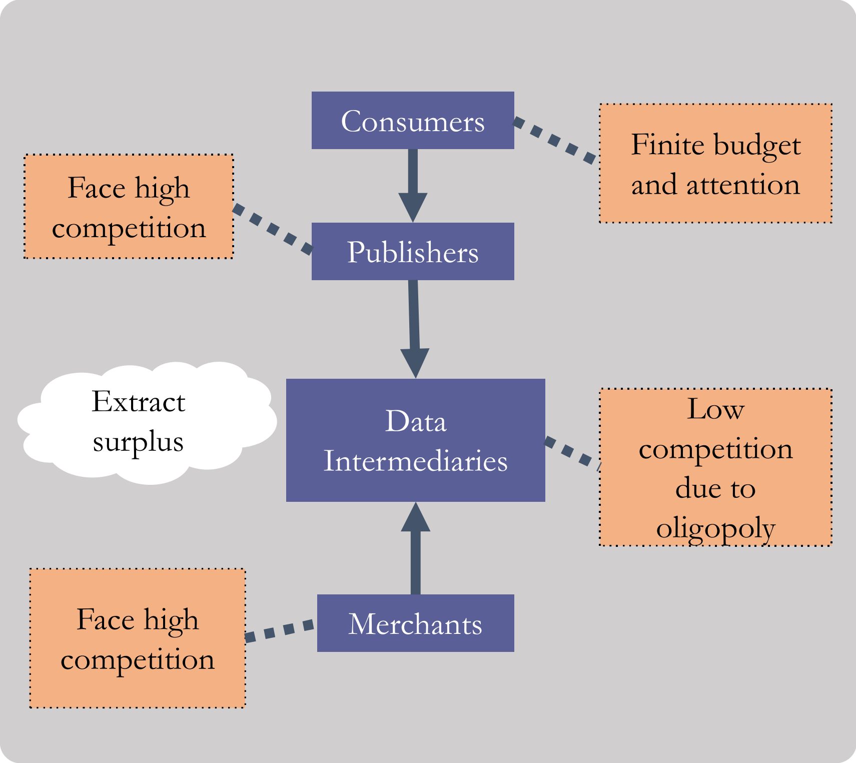

Consumers

Finite budget and attention

Face high competition

Publishers

Extract surplus

Data Intermediaries

Low competition due to oligopoly

Face high competition

Merchants

Source: Moradi, Cheyre, and Acquisti (2025); Acquisti (2024)

Empirical Evidence: Framework 1 or 2?

- Behavioral targeting improves **ad performance metrics** (CTR, conversions…)
- But better metrics do **not establish merchant welfare gains**. Why not?
  - They may be **priced in**
  - Or **competed away** in advertiser rivalry for the same users

- Targeted impressions often command a **price premium**
- But higher prices per impression **do not establish publisher gains**. Why not?
  - Publisher outcomes depend on **price and volume**
  - Gains may be **competed away** across publishers and against other media
  - **Shifting benchmark**: As targeting diffuses, more data is needed to sustain prices

- Evidence for **consumers** is mixed: some studies find **search-cost reduction**, others find **welfare harms** (prices, quality, fraud). Evidence not conclusive:
  - Measure **different welfare margins**
  - Gains can coexist with downstream harms

Toward a Better Evidence Base for Policy

- **Cross-stakeholder** Incidence Studies
  - Measure outcomes across the ecosystem (users, merchants, publishers, and intermediaries)
  - Estimate the distribution of gains and losses
- **Equilibrium and Competition** Studies
  - Analyze how results shift once firms adjust pricing, bidding, and strategies
  - Distinguish between true efficiency gains and zero-sum competition
- **System-Wide** Interventions
  - Distinguish between true efficiency gains and zero-sum competition
  - Use field experiments to study system-wide tradeoffs and effects

CONSUMER INJURIES AND BENEFITS

IN THE DATA-DRIVEN ECONOMY

Panel Five: Measuring Consumer Preferences, Beliefs, and Decisions