Measuring Strategization in Recommendation:

Users Adapt Their Behavior to Shape Future Content

Sarah H. Cen, **Andrew Ilyas**, Jennifer Allen, Hannah Li, and Aleksander Mądry

EC | July 9, 2024

Recommendation is data-driven

Platform provides (personalized) suggestions to each user

Recommendation is data-driven

Platform provides (personalized) suggestions to each user

Recommender systems are powered by **data-driven algorithms**

Is the training data reliable?

In theory, user behavior is **exogenous** (i.e., the user always responds to the same content in the same way, no matter how it's generated or shown)

Is the training data reliable?

In theory, user behavior is **exogenous** (i.e., the user always responds to the same content in the same way, no matter how it's generated or shown)

> Assuming exogeneity is convenient → it implies that differences in behavior <u>must</u> be due to differences in content

Is the training data reliable?

In theory, user behavior is **exogenous** (i.e., the user always responds to the same content in the same way, no matter how it's generated or shown)

> Assuming exogeneity is convenient → it implies that differences in behavior <u>must</u> be due to differences in content

But it's unclear if exogeneity really holds → users are increasingly "aware" of their recommendation algorithms

Is the training data reliable?

So, we hypothesized in Cen et al. (2023) that user behavior is <u>not</u> exogeneous (i.e., users **strategize**)

Is the training data reliable?

So, we hypothesized in Cen et al. (2023) that user behavior is <u>not</u> exogeneous (i.e., users **strategize**)

When users are strategic, the training data becomes unreliable
→ bad for platform learning

Is the training data reliable?

So, we hypothesized in Cen et al. (2023) that user behavior is not exogeneous (i.e., users **strategize**)

When users are strategic, the training data becomes unreliable
→ bad for platform learning

But do users strategize?
If they do, is the effect noticeable?

This work

This work

In this project, we systematically test for user strategization in a lab experiment & survey

This work

In this project, we systematically test for user strategization in a lab experiment & survey

Driving questions:

This work

In this project, we systematically test for user strategization in a lab experiment & survey

Driving questions:

Are users aware of their recommendation algorithms?

This work

In this project, we systematically test for user strategization in a lab experiment & survey

Driving questions:

Are users aware of their recommendation algorithms?

Do users behave strategically in response to algorithms?

This work

In this project, we systematically test for user strategization in a lab experiment & survey

Driving questions:

Are users aware of their recommendation algorithms?

Do users behave strategically in response to algorithms?

If so, how much and why?

This work

In this project, we systematically test for user strategization in a lab experiment & survey

Driving questions:

Are users aware of their recommendation algorithms?

Do users behave strategically in response to algorithms?

If so, how much and why?

Model & Hypotheses

Model (Cen, Ilyas & Mądry ’23)

Model (Cen, Ilyas & Mądry ’23)

User’s true preferences: Utility function U

Model (Cen, Ilyas & Mądry ’23)

User’s true preferences: Utility function U

U (video , click) = payoff user receives if they click on the video

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function U

U (video , click) = payoff user receives if they click on the video

User's revealed preferences: How they behave (what platform sees)

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function *U*

U (video , click) = payoff user receives if they click on the video

User's revealed preferences: How they behave (what platform sees)

Platform cannot observe *U*. Platform observes revealed preferences

Model (Cen, Ilyas & Mądry ’23)

User’s true preferences: Utility function U

U (video , click) = payoff user receives if they click on the video

User’s revealed preferences: How they behave (what platform sees)

Platform cannot observe U. Platform observes revealed preferences

Exogeneity: Revealed preferences depend only on content

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function U

U (video , click) = payoff user receives if they click on the video

User's revealed preferences: How they behave (what platform sees)

Platform cannot observe U. Platform observes revealed preferences

Exogeneity: Revealed preferences depend only on content

Whether user decides to click on video depends on U but <u>not</u> on algorithm

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function *U*

U (video , click) = payoff user receives if they click on the video

User's revealed preferences: How they behave (what platform sees)

Platform cannot observe *U*. Platform observes revealed preferences

Strategic user: Chooses action to optimize long-run payoffs

1. Users are aware that *current* actions affect *future* recommendations
2. Based on knowledge of algorithm, users balance current & future payoffs

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function *U*

U (video , click) = payoff user receives if they click on the video

User's revealed preferences: How they behave (what platform sees)

Platform cannot observe *U*. Platform observes revealed preferences

Strategic user: Chooses action to optimize long-run payoffs

1. Users are aware that *current* actions affect *future* recommendations
2. Based on knowledge of algorithm, users balance current & future payoffs

Model (Cen, Ilyas & Mądry '23)

User's true preferences: Utility function U

U(video , click) = payoff user receives if they click on the video

It's difficult to test directly for strategization
(Every user has a different <u>unknown</u> utility U)

Instead, we test for the *effects* of strategization

1. Users are <u>aware</u> that *current* actions affect *future* recommendations
2. <u>Based on knowledge of algorithm</u>, users balance current & future payoffs

We test for two effects of strategization

We test for two effects of strategization

Hypothesis #1: Different algorithms induce different behaviors

We test for two effects of strategization

Hypothesis #1: Different algorithms induce different behaviors

This would imply that users adapt their behavior to algorithms (specifically, to how they believe the algorithm learns preferences)

We test for two effects of strategization

Hypothesis #1: Different algorithms induce different behaviors

This would imply that users adapt their behavior to algorithms (specifically, to how they believe the algorithm learns preferences)

Hypothesis #2: Telling users that they will receive personalized recommendations causes different behaviors

We test for two effects of strategization

Hypothesis #1: Different algorithms induce different behaviors

This would imply that users adapt their behavior to algorithms (specifically, to how they believe the algorithm learns preferences)

Hypothesis #2: Telling users that they will receive personalized recommendations causes different behaviors

This would imply that users strategize w.r.t. their algorithm <u>because</u> they believe their current actions impact future recommendations

We test for two effe

We will argue our results aren't explained by experimenter demand (i.e., participants aren't just subconsciously responding to experiment cues)

Hypothesis #1: Different algorithms induce different behaviors

This would imply that users adapt their behavior to algorithms (specifically, to how they believe the algorithm learns preferences)

Hypothesis #2: Telling users that they will receive personalized recommendations causes different behaviors

This would imply that users strategize w.r.t. their algorithm <u>because</u> they believe their current actions impact future recommendations

Experiment

Lab experiment

Session 1 (Time remaining: 04:53)

I'll Play The Blues For You (Live At Wattstax / 1972)

Albert King

0:04 / 5:39

Lab experiment

Session 1 (Time remaining: 04:53)

I'll Play The Blues For You (Live At Wattstax / 1972)

Albert King

0:04 / 5:39

Experimental methodology

635 participants on CloudResearch

Each participant underwent two 5-minute listening sessions, then a post-experiment survey (demographics & open-ended questions)

We randomly assign users to **incentive** and **information** **groups**

Two incentive conditions

Tests Hypothesis 2

Incentive condition: Will I receive recommendations?

Control: Told their behaviors are used to learn population preferences

Treatment: Told they will be given personalized recommendations

Study Description

Info Control | 5-minute listening warmup

Info Condition | 5-minute listening session

Post-experiment survey

Incentive Control

Learning People's Music Preferences

(Please temporarily disable ad blockers, and do not press the back button in your browser.)

Re-read the study description here.

Study Description

In this study, we are gathering information on what music the general population likes. During this study, we will also observe how people interact with songs (like which songs people "thumbs-up").

There are two stages to this study:

1. **Stage 1**: We'll show you songs and observe how you interact with them. There will be three listening sessions. All the songs are chosen randomly.
2. **Stage 2**: Your behavior from Stage 1 will be used in our study, and you'll be asked to perform a brief survey.

Incentive Treatment

Testing a Music Recommendation Algorithm

(Please temporarily disable ad blockers, and do not press the back button in your browser.)

Re-read the study description here.

Stage 1

Let's start Stage 1. In this stage, we will show you a series of songs over three listening sessions. **During all three sessions, the songs are generated randomly (not using an algorithm)**.

Each session of the three sessions will last <u>5 minutes</u>. A few things to know:

1. You cannot go back to songs once they have ended or been skipped.
2. You can interact with the songs (see the buttons below). Interacting with the songs is **optional**.

As a reminder, our algorithm will recommend three music artists for you at the end of this study based on how you interact with songs during this stage.

Three information conditions

Tests Hypothesis 1

Information condition: How are my preferences learned?

Control: Given no information about learning

Likes/dislikes: Told preferences are mostly learned from likes/dislikes

Dwell time: Told preferences mostly learned from dwell time

Study Description

Info Control | 5-minute listening warmup

Info Condition | 5-minute listening session

Post-experiment survey

Example: Information Condition

Stage 1: Warm-Up Session

This session will last 5 minutes. We will show you a random selection of songs and log how you interact.

Remember: You do not have to interact with the songs. Feel free to skip songs.

During this session, we want to get a baseline for what songs you like. We ask that you interact as you would with a song recommender like Spotify, Pandora, or YouTube.

Session 1 (Time remaining: 04:53)

I'll Play The Blues For You (Live At Wattstax / 1972)

Albert King

0:04 / 5:39

Example: Information Condition

Stage 1: Warm-Up Session

This session will last 5 minutes. We will show you a random selection of songs and log how you interact.

Remember: You do not have to interact with the songs. Feel free to skip songs.

During this session, we want to get a baseline for what songs you like. We ask that you interact as you would with a song recommender like Spotify, Pandora, or YouTube.

Session 1 (Time remaining: 04:53)

I'll Play Th
1972)

Albert King

0:04 / 5:39

Stage 1: Training the Algorithm (Session 2)

As before, we will show you a random selection of songs for 5 minutes and log how you interact.

Remember: You do not have to interact with the songs. Feel free to skip songs.

In this second session, **the algorithm will pay more attention to what you thumbs-up and thumbs-down than the algorithm did in the warm-up** in order to figure out what types of songs you enjoy. As a reminder, the songs during this session are chosen randomly.

2x3 Factorial Experiment

Information condition (Hypothesis 1)

Incentive condition (Hypothesis 2)

Incentive Control Info Control	Incentive Control Likes/Dislikes Info	Incentive Control Dwell Time Info
Incentive Treatment Info Control	Incentive Treatment Dwell Time Info	Incentive Treatment Dwell Time Info

Results

Analysis

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Model: Fixed-effects model (with main effects $\beta_1$and β_2, interaction effect β_3)

$$Y \sim \beta_0 + \beta_1 D_{\text{Incentive}} + \beta_2 D_{\text{Info}} + \beta_3 (D_{\text{Incentive}} \times D_{\text{Info}}) + \beta_4 Y_{pre} + \varepsilon$$

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Model: Fixed-effects model (with main effects $\beta_1$and β_2, interaction effect β_3)

$$Y \sim \beta_0 + \beta_1 D_{\text{Incentive}} + \beta_2 D_{\text{Info}} + \beta_3 (D_{\text{Incentive}} \times D_{\text{Info}}) + \beta_4 Y_{pre} + \varepsilon$$

Findings: Strong evidence supporting both hypotheses

Interaction effects (between hypotheses) are suggestive but not significant

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Model: Fixed-effects model (with main effects $\beta_1$and β_2, interaction effect β_3)

$$Y \sim \beta_0 + \beta_1 D_{\text{Incentive}} + \beta_2 D_{\text{Info}} + \beta_3 (D_{\text{Incentive}} \times D_{\text{Info}}) + \beta_4 Y_{pre} + \varepsilon$$

Findings: Strong evidence supporting both hypotheses

Interaction effects (between hypotheses) are suggestive but not significant

Potential limitations:

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Model: Fixed-effects model (with main effects β_1 and β_2, interaction effect β_3)

$$Y \sim \beta_0 + \beta_1 D_{\text{Incentive}} + \beta_2 D_{\text{Info}} + \beta_3 (D_{\text{Incentive}} \times D_{\text{Info}}) + \beta_4 Y_{pre} + \varepsilon$$

Findings: Strong evidence supporting both hypotheses

Interaction effects (between hypotheses) are suggestive but not significant

Potential limitations:

- Experimenter demand → cannot explain away our findings

Analysis

Measured outcomes: (dis)likes, skips, total clicks, average & variance of dwell time

Model: Fixed-effects model (with main effects $\beta_1$and β_2, interaction effect β_3)

$$Y \sim \beta_0 + \beta_1 D_{\text{Incentive}} + \beta_2 D_{\text{Info}} + \beta_3 (D_{\text{Incentive}} \times D_{\text{Info}}) + \beta_4 Y_{pre} + \varepsilon$$

Findings: Strong evidence supporting both hypotheses

Interaction effects (between hypotheses) are suggestive but not significant

Potential limitations:

- Experimenter demand → cannot explain away our findings
- Our experiment design can only surface <u>average</u> treatment effects

Means and 95% confidence intervals

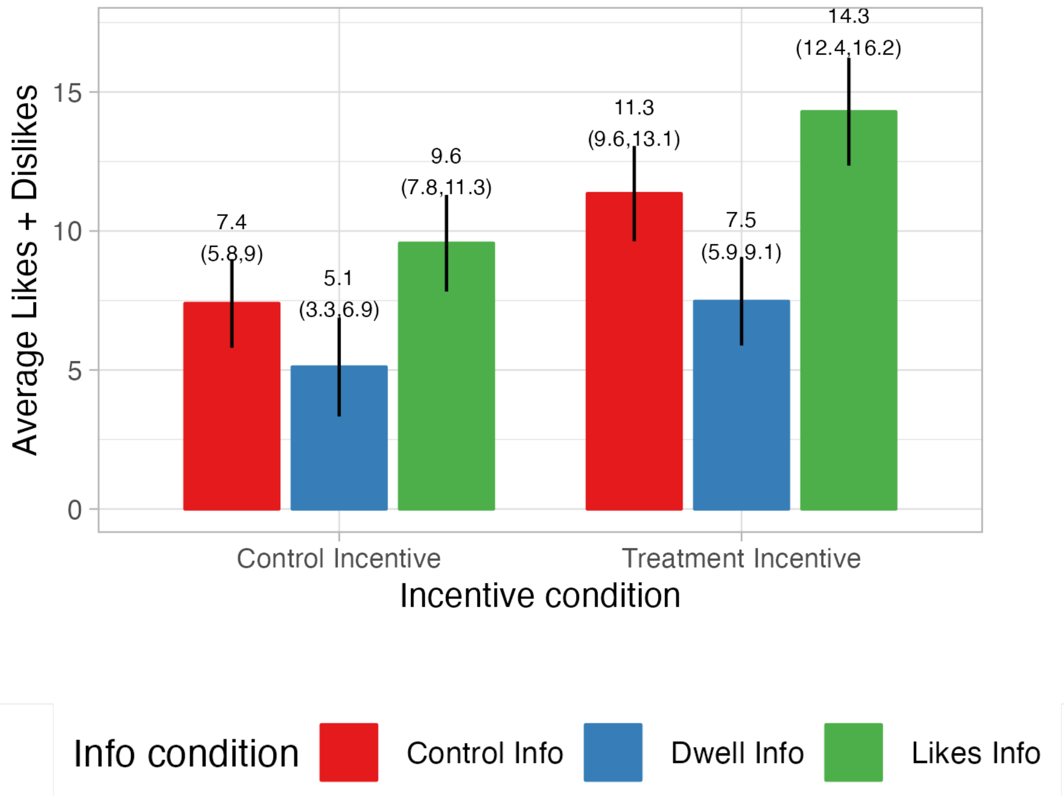

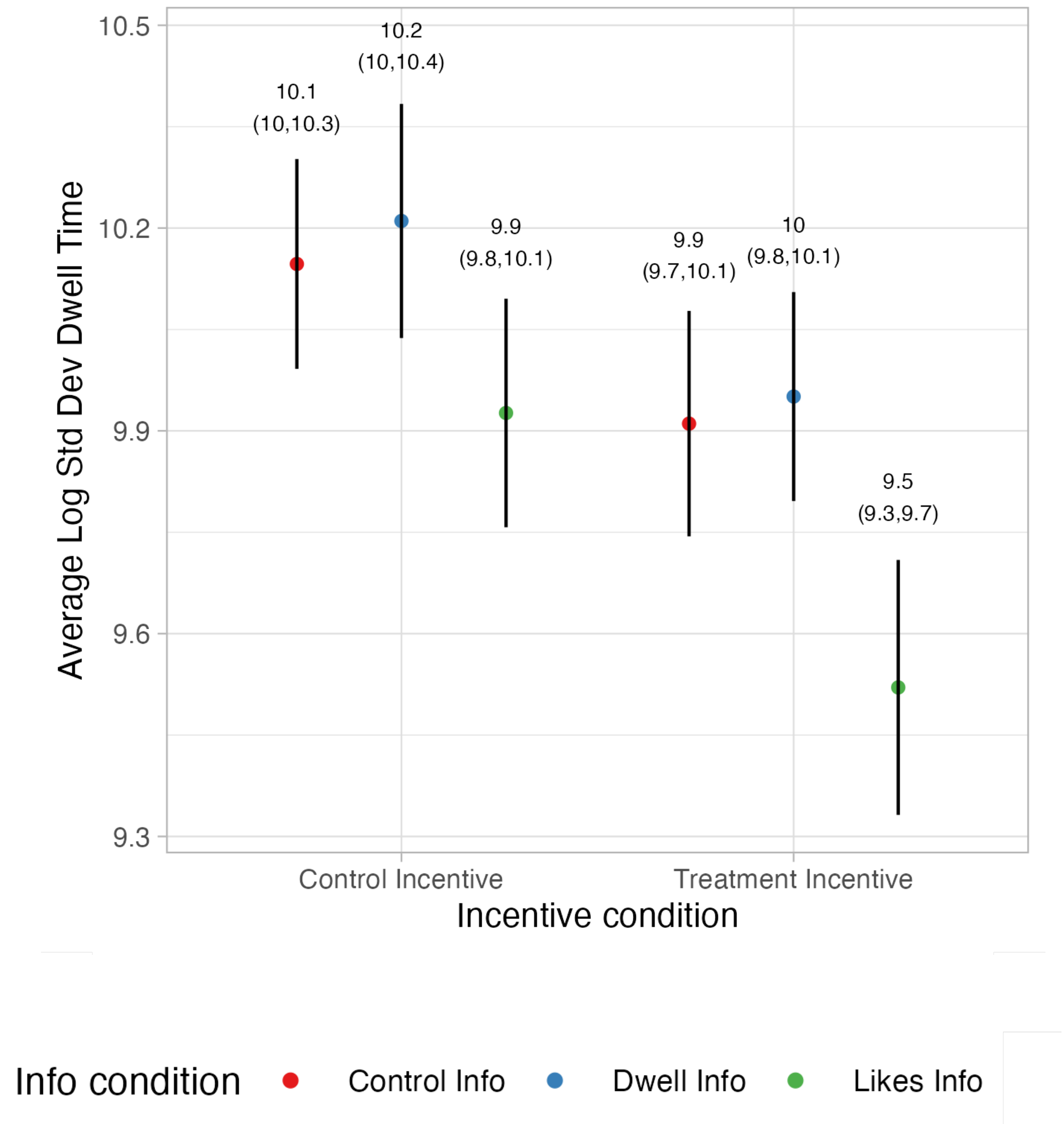

Marginal effects (Information)

Effect of the Likes and Dwell Information conditions, compared to the Control. OLS regression (left) and quasi-Poisson regression (right) with controls for behavior in the Warm-up session.

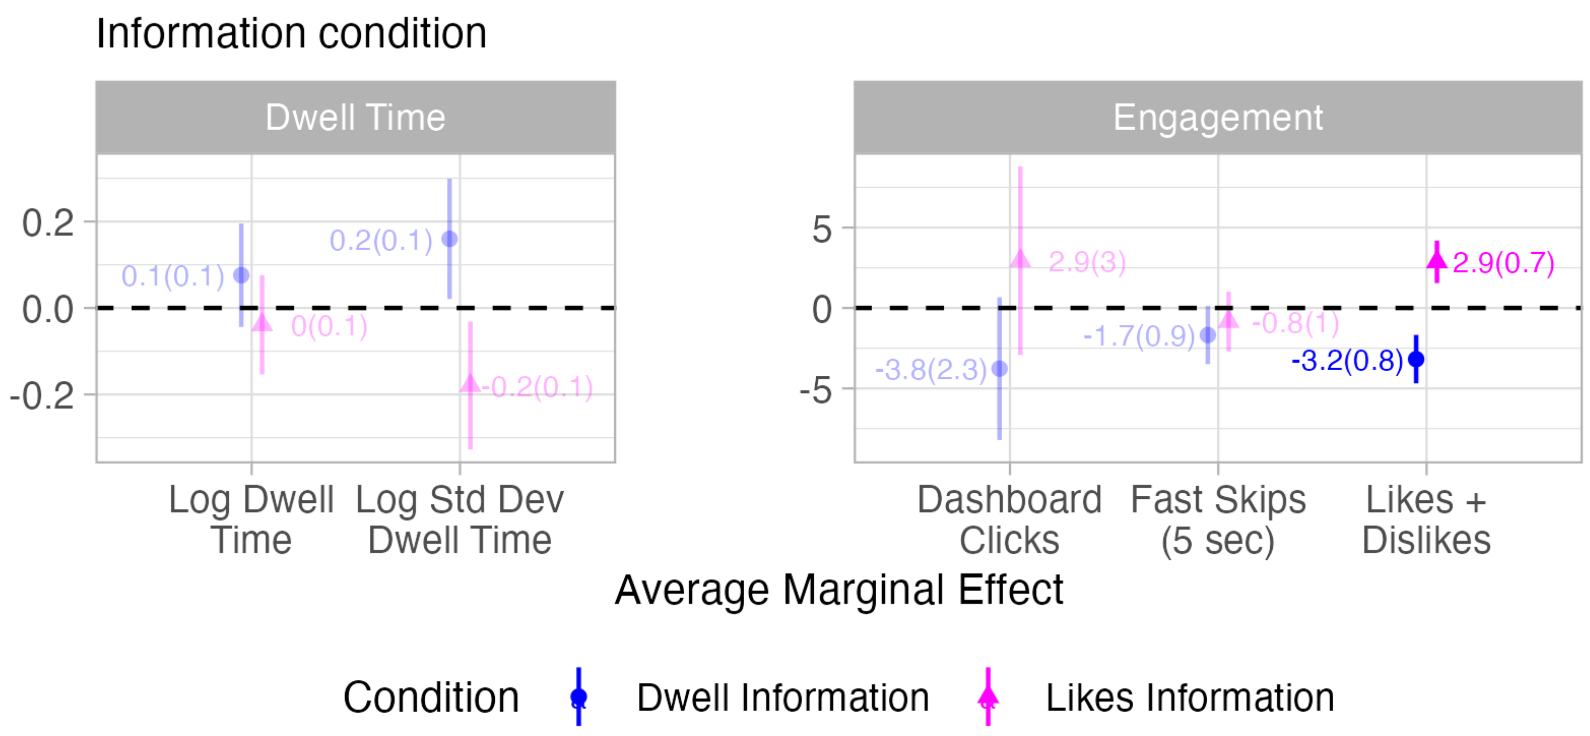

Marginal effects (Incentive)

Effect of the Treatment Incentive condition compared to the Control Incentive condition. OLS regression (left) and quasi-Poisson regression (right).

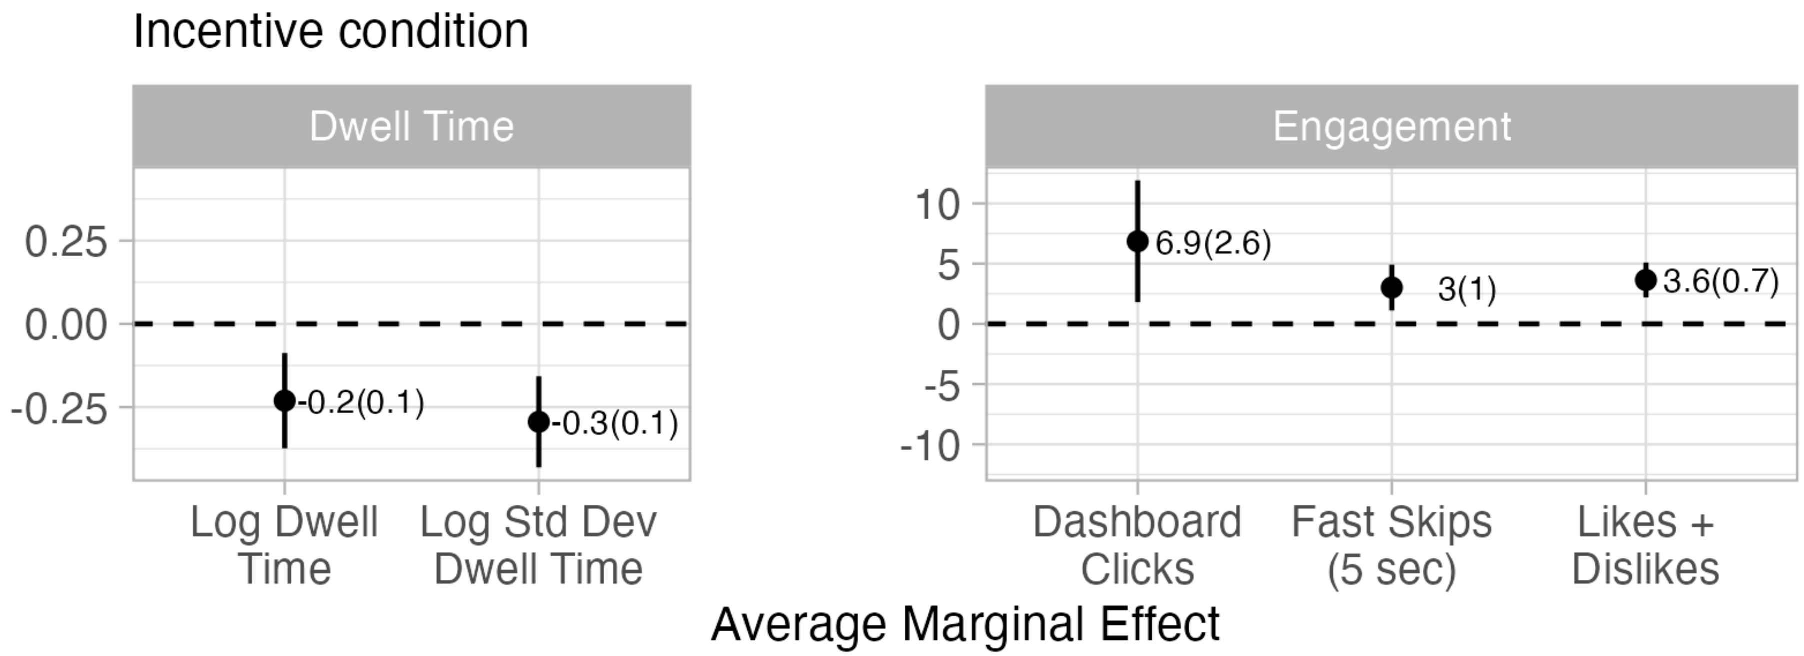

Some behaviors differ across age/platform use

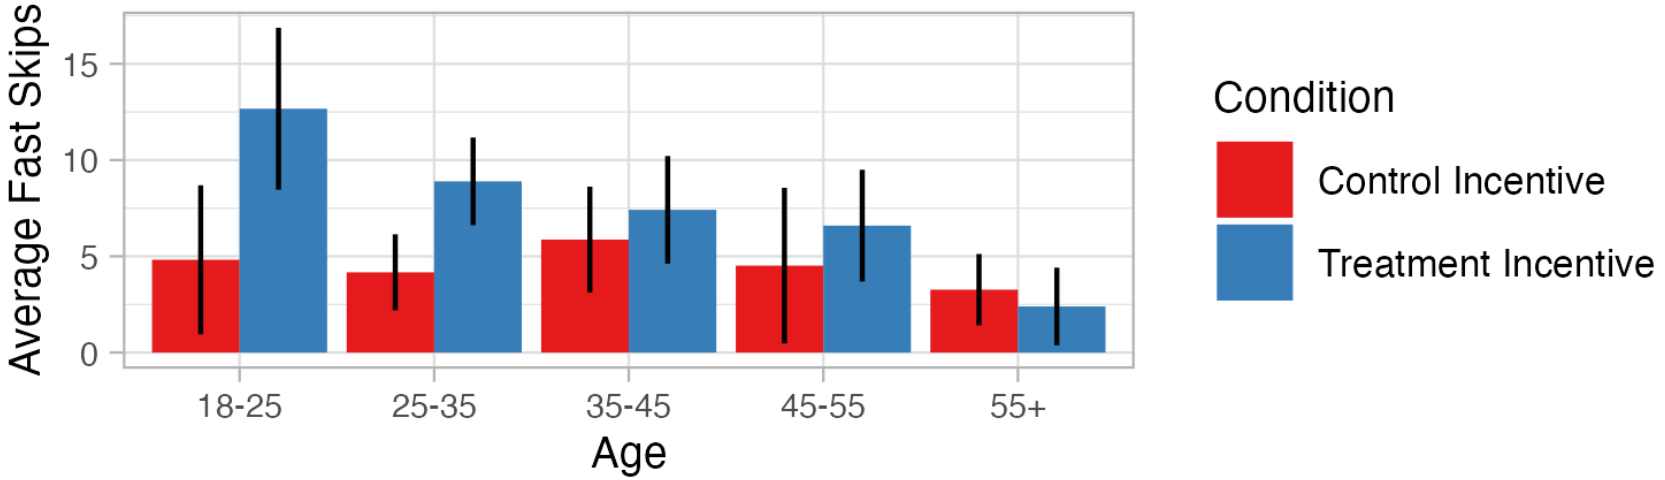

Some behaviors differ across age/platform use

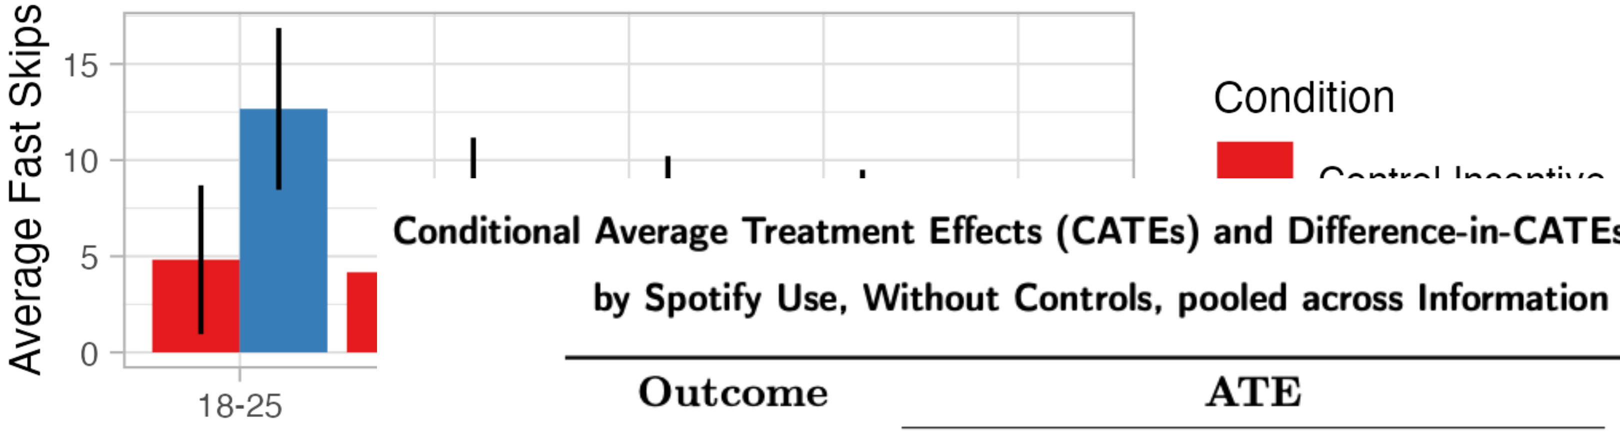

Conditional Average Treatment Effects (CATEs) and Difference-in-CATEs (DIC) of the Incentive condition, by Spotify Use, Without Controls, pooled across Information Conditions

Outcome	ATE		DIC
	Spotify Use=Often	Spotify Use=Rare	
Likes + Dislikes	4.67***(0.94)	1.54(1.13)	3.13*(1.47)
Fast Skips (5 sec)	4.18**(1.27)	0.87(1.46)	3.31†(1.93)
Dashboard Clicks	9.83**(3.32)	0.13(4.35)	9.7†(5.47)

[a] Heteroskedasticity Robust Standard Errors in Parentheses.
[b] Signif. Codes: ***: .001, **: .01, *: .05, †: .1

Post-experiment survey

We asked users if they strategize in the wild. We find that users strategize:

- To avoid seeing undesirable content (or advertisements) in future
- To preserve privacy (e.g., with private browsing)
- To avoid feedback loops (being pigeonholed)
- To help the algorithm
- Some don’t strategize at all

Post-experiment survey

We asked users if they strategize in the wild. We find that users strategize:

- To avoid seeing undesirable content (or advertisements) in future
- To preserve privacy (e.g., with private browsing)
- To avoid feedback loops (being pigeonholed)
- To help the algorithm
- Some don't strategize at all

"Sometimes I may like a song but not thumbs-up the song because I don't want my feed filled with similar artists/videos"

Post-experiment survey

We asked users if they strategize in the wild. We find that users strategize:

- To avoid seeing undesirable content (or advertisements) in future
- To preserve privacy (e.g., with private browsing)
- To avoid feedback loops (being pigeonholed)
- To help the algorithm
- Some don't strategize at all

"I avoid reading certain news stories on Google news because I know I will be bombarded with similar articles. Instead I switch to an untracked browser to read the story."

"Sometimes I may like a song but not thumbs-up the song because I don't want my feed filled with similar artists/videos"

Post-experiment survey

We asked users if they strategize in the wild. We find that users strategize:

- To avoid seeing undesirable content (or advertisements) in future
- To preserve privacy (e.g., with private browsing)
- To avoid feedback loops (being pigeonholed)
- To help the algorithm
- Some don't strategize at all

"I avoid reading certain news stories on Google news because I know I will be bombarded with similar articles. Instead I switch to an untracked browser to read the story."

"Sometimes I may like a song but not thumbs-up the song because I don't want my feed filled with similar artists/videos"

"I have many YouTube accounts so my algorithm does not pick up [on] a YouTube link a friend sends me to watch"

Takeaways

Takeaways

Users are aware of their recommendation algorithms

Takeaways

Users are aware of their recommendation algorithms

We find evidence that users *do* strategize

- Users change behavior based on perception of algorithm
- Users engage differently based on how current actions will affect them downstream

Takeaways

Users are aware of their recommendation algorithms

We find evidence that users *do* strategize

- Users change behavior based on perception of algorithm
- Users engage differently based on how current actions will affect them downstream

In post-experiment survey, we find that many users very consciously strategize to both help & hide from the algorithm!

Takeaways

Users are aware of their recommendation algorithms

We find evidence that users *do* strategize

- Users change behavior based on perception of algorithm
- Users engage differently based on how current actions will affect them downstream

In post-experiment survey, we find that many users very consciously strategize to both help & hide from the algorithm!

Paper: https://arxiv.org/abs/2405.05596

Email: ailyas@mit.edu

Web: https://andrewilyas.com