

Building a Personalized Messaging System at Netflix

Grace Huang

July 31th, 2020

Data Council SF 2020



DOES IT SPARK JOY?




**When members find content
they love, they enjoy our
service more**




Our messages are designed to help them find content to enjoy

Some ways to reach out to our members

Email



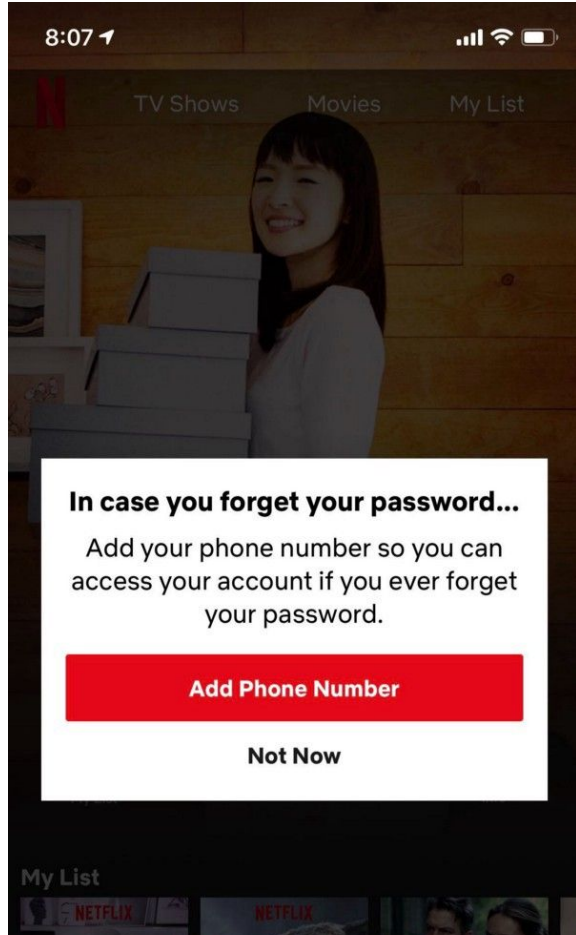

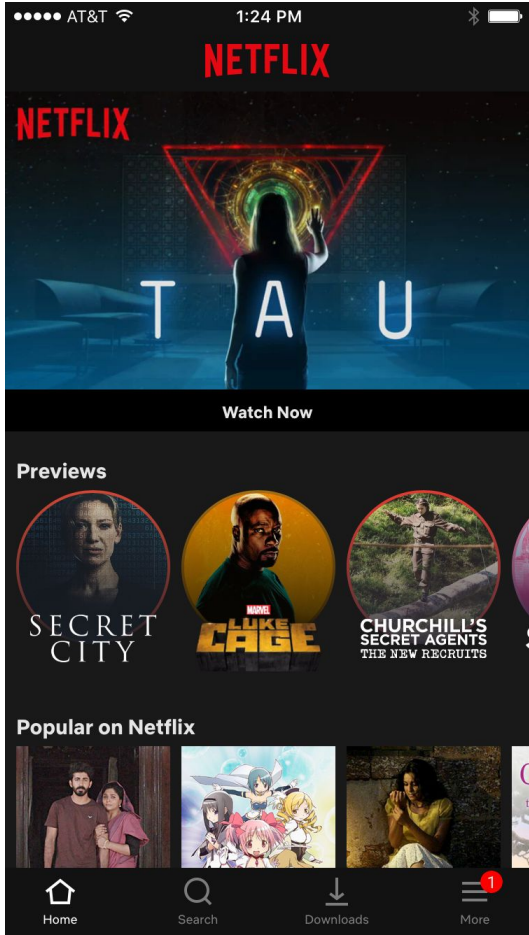
Push



SMS

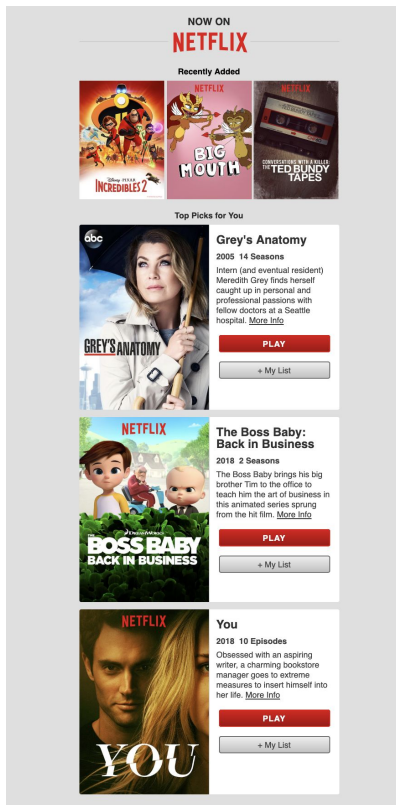


In-App: Notifications and Alerts



A Variety of Message Types

Recommendations



NOW ON NETFLIX

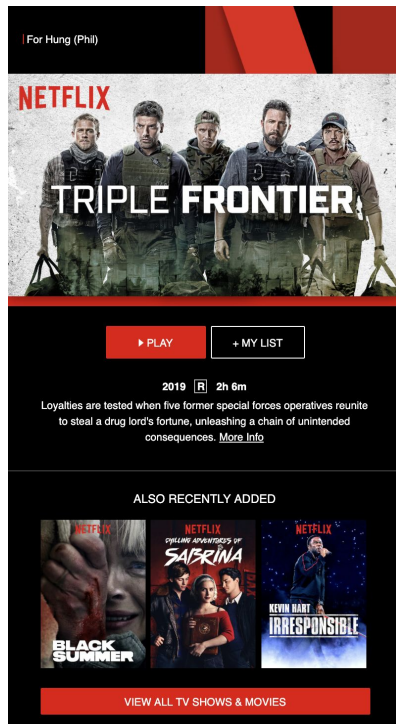
Recently Added

- THE INCREDIBLES 2**
- BIG MOUTH**
- THE UNDISCOVERED COUNTRY OF TED BLINDY TAPES**

Top Picks for You

- Grey's Anatomy**
2019 14 Seasons
Intern (and eventual resident) Meredith Grey finds herself caught up in personal and professional passions with fellow doctors at a Seattle hospital. [More Info](#)
- The Boss Baby: Back in Business**
2018 2 Seasons
The Boss Baby brings his big brother Tim to the office to teach him the art of business in this animated series springing from the hit film. [More Info](#)
- You**
2018 10 Episodes
Obsessed with an aspiring writer, a charming bookstore manager goes to extreme measures to insert himself into her life. [More Info](#)

New Arrival



For Hung (Phil)

NETFLIX

TRIPLE FRONTIER

PLAY **+ MY LIST**

2019 **R** 2h 6m

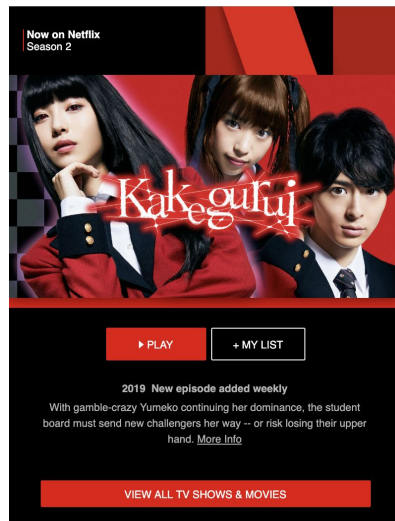
Loyalties are tested when five former special forces operatives reunite to steal a drug lord's fortune, unleashing a chain of unintended consequences. [More Info](#)

ALSO RECENTLY ADDED

- BLACK SUMMER**
- THE UNDISCOVERED COUNTRY OF TED BLINDY TAPES**
- KEVIN KURT IRRESPONSIBLE**

VIEW ALL TV SHOWS & MOVIES

New Season Alert



Now on Netflix Season 2

Kakegurui

PLAY **+ MY LIST**

2019 New episode added weekly

With gamble-crazy Yumeko continuing her dominance, the student board must send new challengers her way -- or risk losing their upper hand. [More Info](#)

VIEW ALL TV SHOWS & MOVIES

Coming Soon



AT&T 100%


9:41
Tuesday, September 13

Netflix now

Mirage, The Dirt, and more.

Netflix
Coming Soon: 44 New Titles
Mirage, The Dirt, and more.

Netflix
Coming Soon: 44 New Titles
We think you're going to like these.

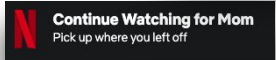


LOOK AHEAD





A Heuristics Driven Paradigm

Candidate Messages



Decision Engine



Email



Push Notifications



In-App Alerts

Heuristics based:

- If a user watched Stranger Things Season 1, then send a message about Season 2 arrival
- Do not send if a user has had a similar recommendation in the past x days
-

A Heuristics Driven Paradigm

Candidate Messages


N Now on Netflix
Season 3

N Continue Watching for Mom
Pick up where you left off


Top 10
United States
As of 7/31/2020

Top Picks for You

ML-Driven
Decision Engine



How many to send?
What to send?



Email



Push Notifications



In-App Alerts

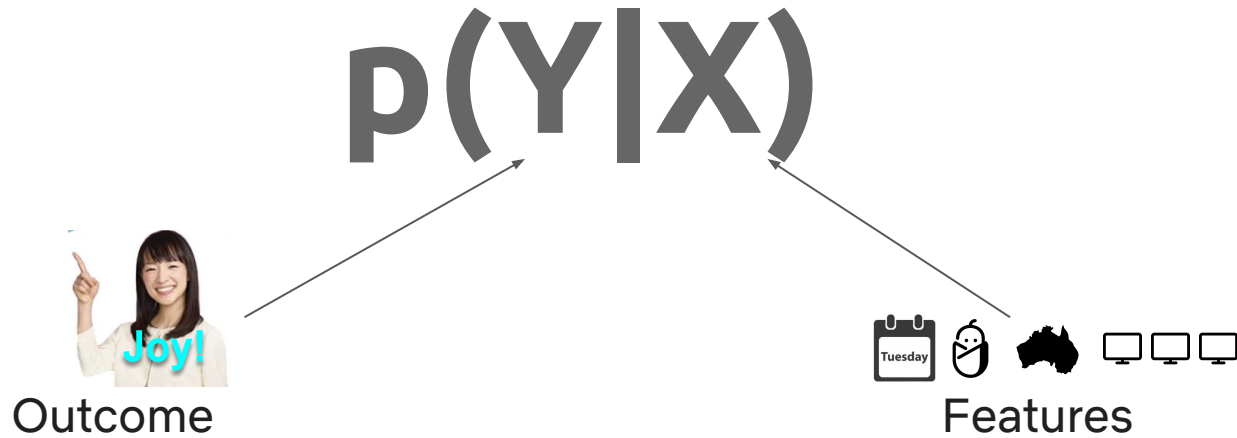
Key Considerations for the System

- Making a **personalized, timely** decision for every Netflix subscriber
- Removing **bias** from the system
- Maximizing **causal** impact
- Balancing **reward** against **cost**

Personalizing the messaging decision



A Personalized Messaging Decision



Can be estimated using a variety of classification (or regression) techniques -

Linear (or Logistic) Regression, GBDT, Neural Network...etc

But...


How to obtain data with the full range of messaging frequency and message type variations?

Removing Bias from the System

2


The Obvious Candidate: Explore/Exploit

ϵ -greedy




- Explore with probability ϵ
- Otherwise, choose arm with best action

UCB



- Pull arm with the highest upper confidence bound

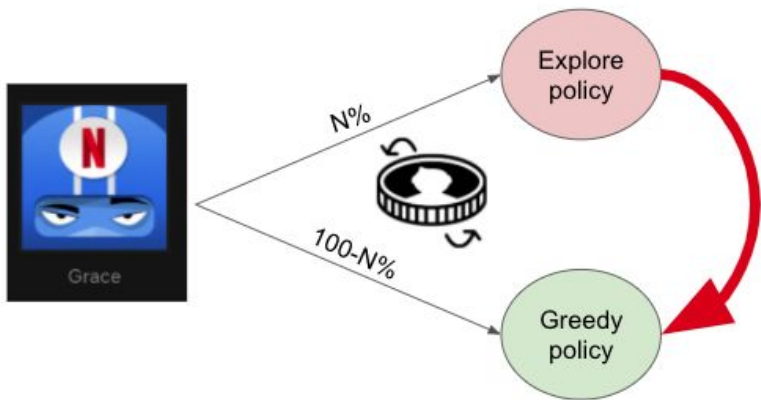
Thompson Sampling



[chart source](#)

- Take a random sample from each arm's PDF
- Choose the arm with the highest sampled value.

An Example Approach: Personalized Messaging using Contextual Bandit



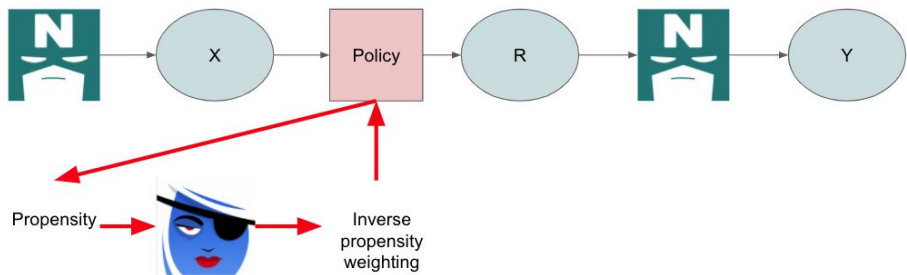
Choose a random message:

Top Picks for You


Context	Messages Chosen
Tuesday   	Top Picks for You Top 10 United States As of 7/31/2020
Sunday   	Continue Watching for Mom Pick up where you left off
Saturday   	Top Picks for You N Now on Netflix Season 3

Other Examples of Debiasing Techniques: Propensity Correction

- Propensity Correction (e.g. IPS)



- Instrument Variables..



But...

Did a subscriber visit Netflix and watch a movie because the message we sent was truly relevant and helpful?


Would they have watched a show even if we did not reach out?

Maximizing Causal Impact

A large, bold, white number '3' is centered vertically on a red background. The red background has a subtle, diagonal, wavy pattern. The number '3' is cut off on the right side by the edge of the frame.

A Causal Personalization System

- $p(Y|X)$ only captures correlation
- Every observation is influenced by past actions from our messaging system




<https://xkcd.com/552/>



Recall that we built a *Correlational Personalization Model* for Messaging...

$$P(Y|X)$$



Explicitly Model Past Actions

$$p(Y|X, \text{do}(R))$$

Personalized Decision Making

$R \in \{\emptyset, [\text{Top Picks for You}], [\text{Continue Watching for Mom}], \dots\}$


N Now on Netflix
Season 3

Send a Message when...



Member satisfaction with message

$$p(Y|X, \text{do}([\text{STRANGER THINGS}]))$$

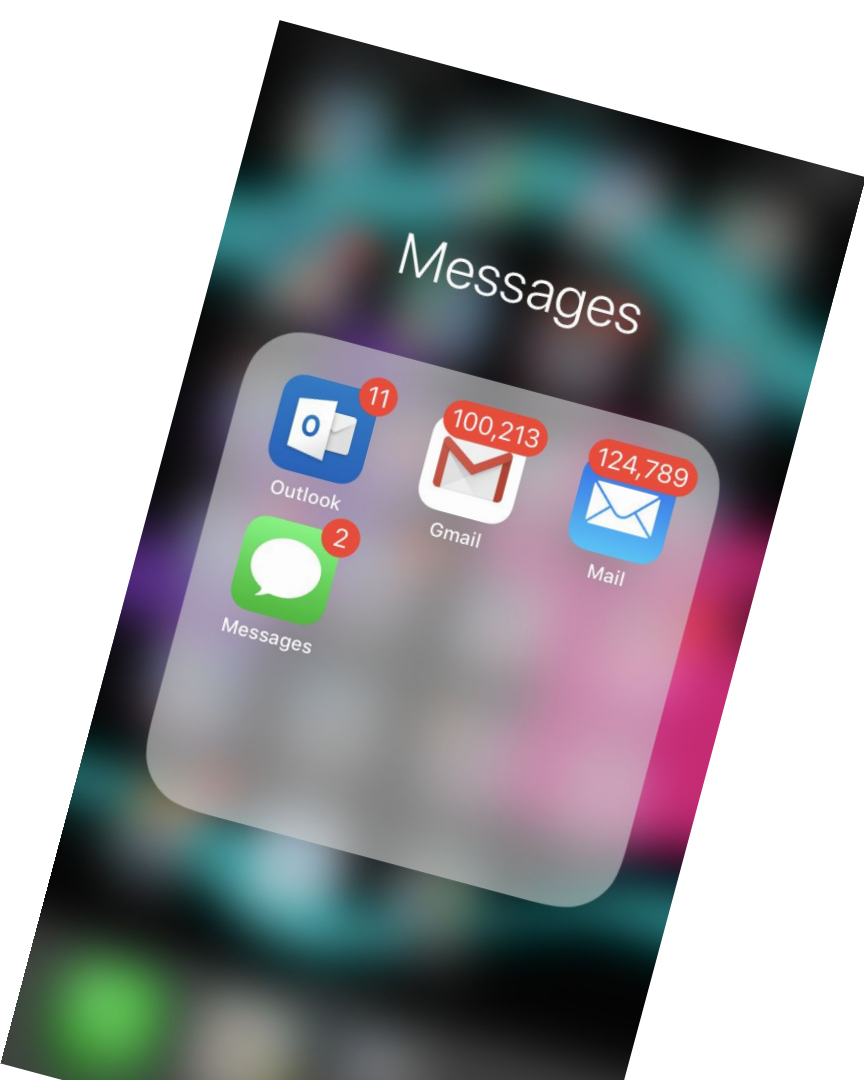


Member satisfaction with no message

$$p(Y|X, \text{do}(\emptyset))$$

Balancing Reward against Cost

4




More \neq Better

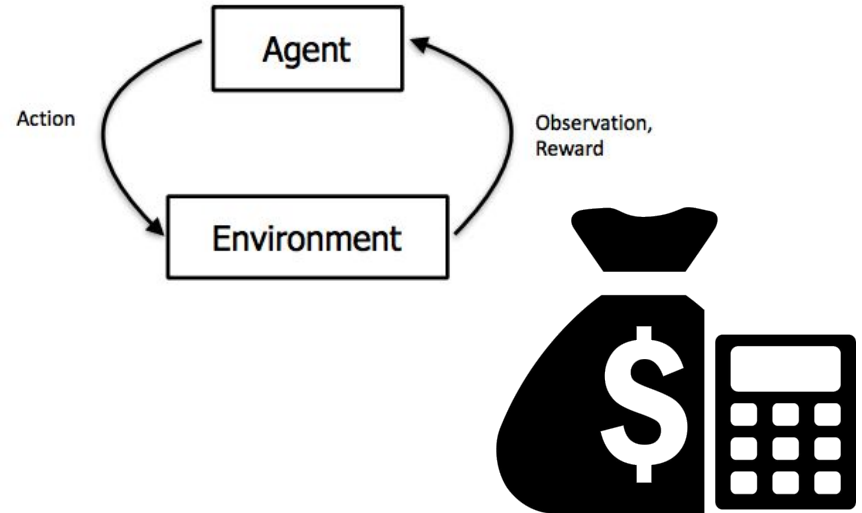
How to impose a volume constraint?

How to Impose a Volume Constraint?

- A simple approach:
Set an Incrementality Threshold




- Other flavors of
Reinforcement Learning




An example of how to put this together....

A Causal Bandit

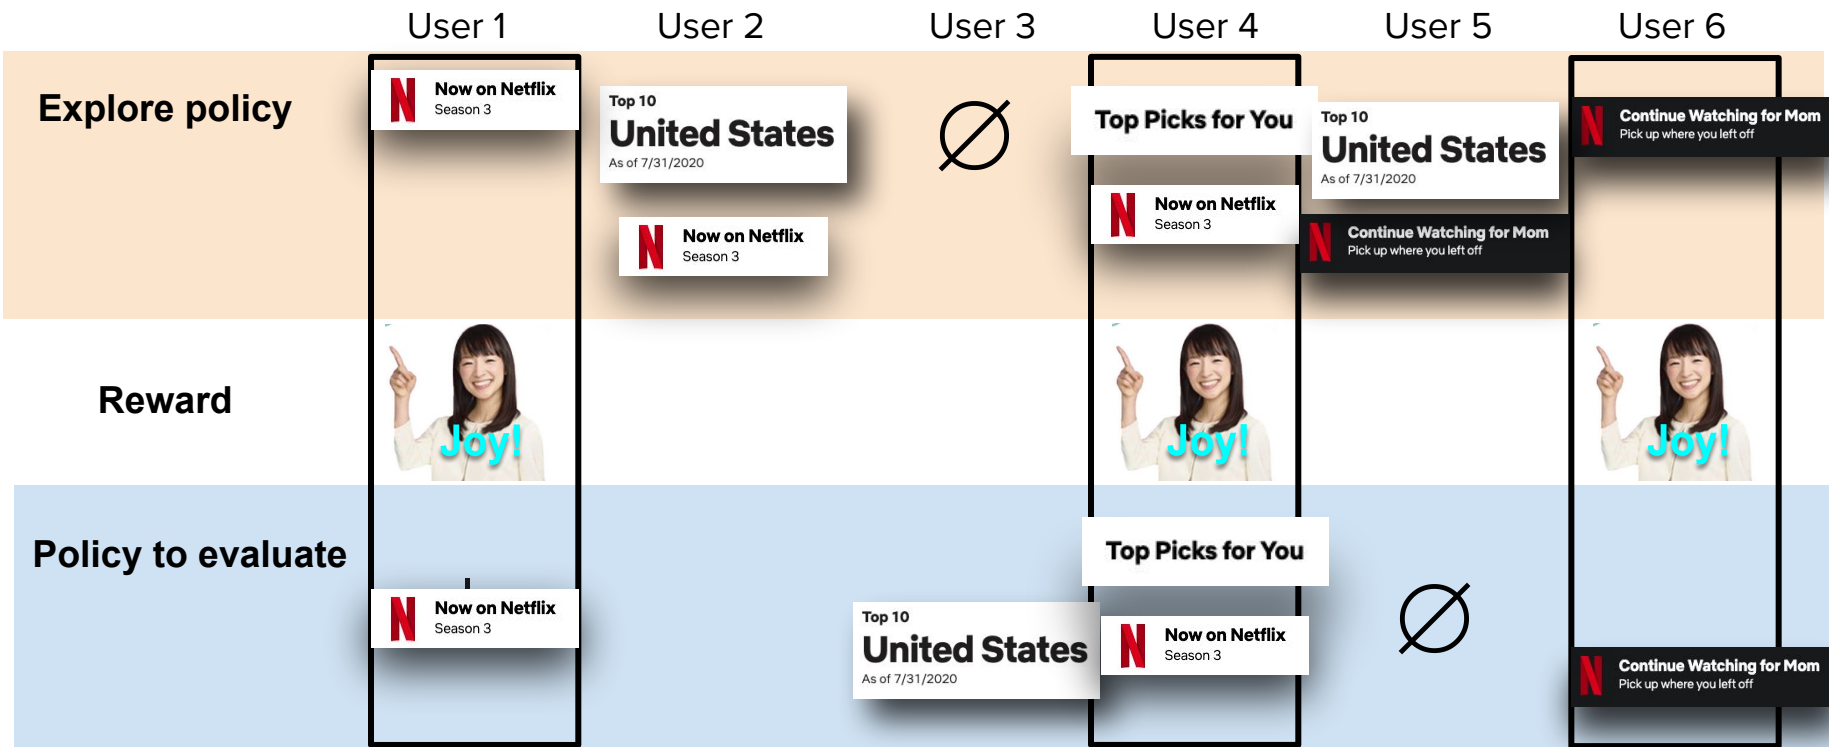


- Randomly sample from frequency and candidate message distributions

- Estimate incrementality of a message
- Greedily assemble the set of messages to send, subject to an incrementality threshold



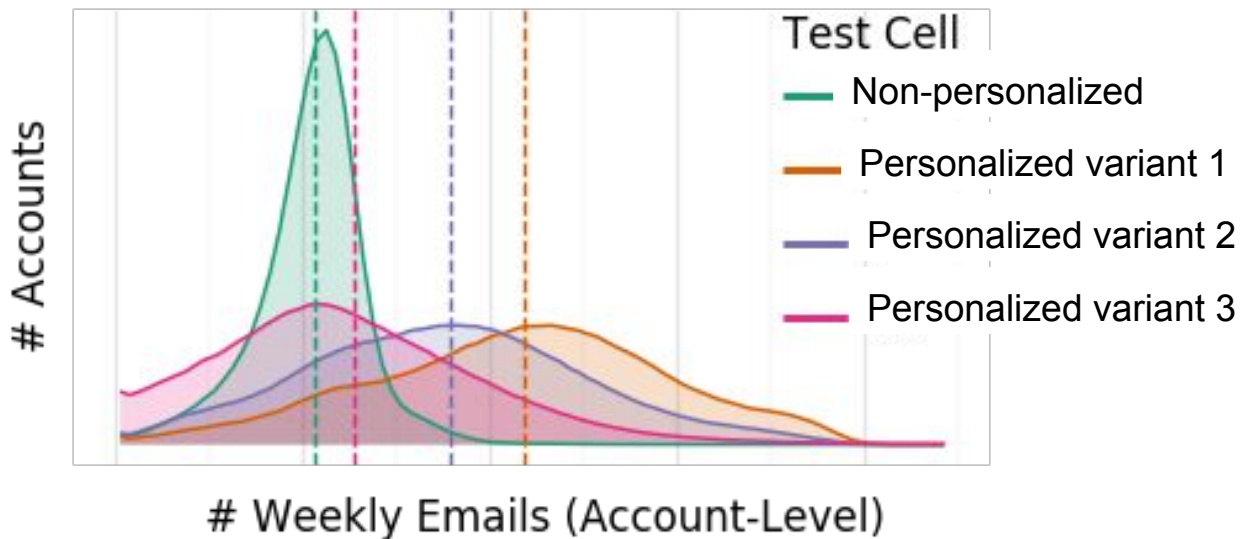
Offline Evaluation



Simulate online metrics offline!

Evaluating Online

- An A/B test still allows us to evaluate the **long term** behavior of a given policy



Recap

- Making a **personalized, timely** decision for every Netflix subscriber
- Removing **bias** from the system
- Maximizing **causal** impact
- Balancing **reward** against **cost**
- Lots of exciting future work.....

Check out more in [Recap: Designing a more Efficient Estimator for Off-policy Evaluation in Bandits with Large Action Spaces](#)

Thank You.

N

Grace Huang

ghuang@netflix.com