# The Fault in Our Recommendations
## On the Perils of Optimizing the Measurable

Omar Besbes, Yash Kanoria, and Akshit Kumar
Columbia University, Graduate School of Business

RecSys'24 | October 15, 2024

# Something is missing in today's recommendation systems

Graydon Snider · Follow

4 min read · May 21, 2021

Medium

Search

Write

Sign up

Sign in

# Something is missing in today's recommendation systems

Graydon Snider · Follow

4 min read · May 21, 2021

Medium

Search

Write

Sign up

Sign in

# Personalization and Popularity Bias

Is popularity bias just a convenient excuse?

Michael Roizner · Follow

3 min read · Dec 26, 2023

# Mismatch in Desiderata & Practice of RecSys

# Mismatch in Desiderata & Practice of RecSys

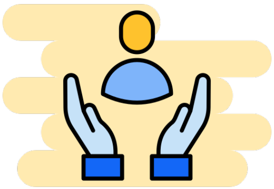

RecSys want to optimize user utility

# Mismatch in Desiderata & Practice of RecSys

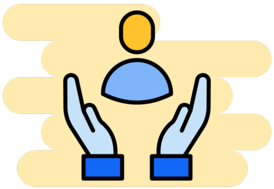

RecSys want to
optimize user utility

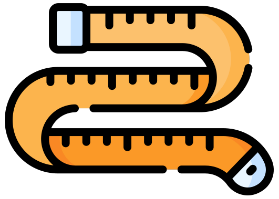

But measuring utility
is challenging

# Mismatch in Desiderata & Practice of RecSys

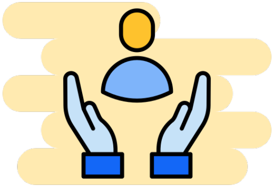

RecSys want to
optimize user utility

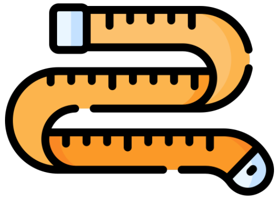

But measuring utility
is challenging

Optimize for proxies
such as engagement

# Mismatch in Desiderata & Practice of RecSys

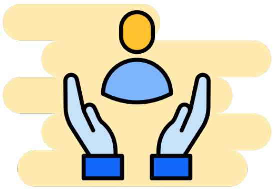

RecSys want to
optimize user utility

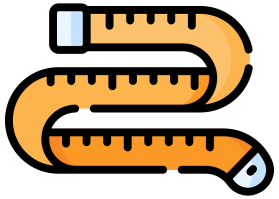

But measuring utility
is challenging

Optimize for proxies
such as engagement

By optimizing for measurable proxies, are recommendation systems at risk of significantly under-delivering on utility?

# Mismatch in Desiderata & Practice of RecSys

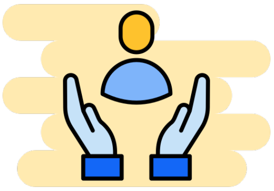

RecSys want to
optimize user utility

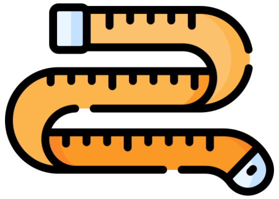

But measuring utility
is challenging

Optimize for proxies
such as engagement

- By optimizing for measurable proxies, are recommendation systems at risk of significantly under-delivering on utility?
- How can we optimize for utility despite not being able to measure it?

# In a Nutshell

# In a Nutshell

Study a stylized model of repeated user interaction

# In a Nutshell

Study a stylized model of repeated user interaction

Engagement optimization may lead to recommending only popular items, a.k.a, “popularity bias”

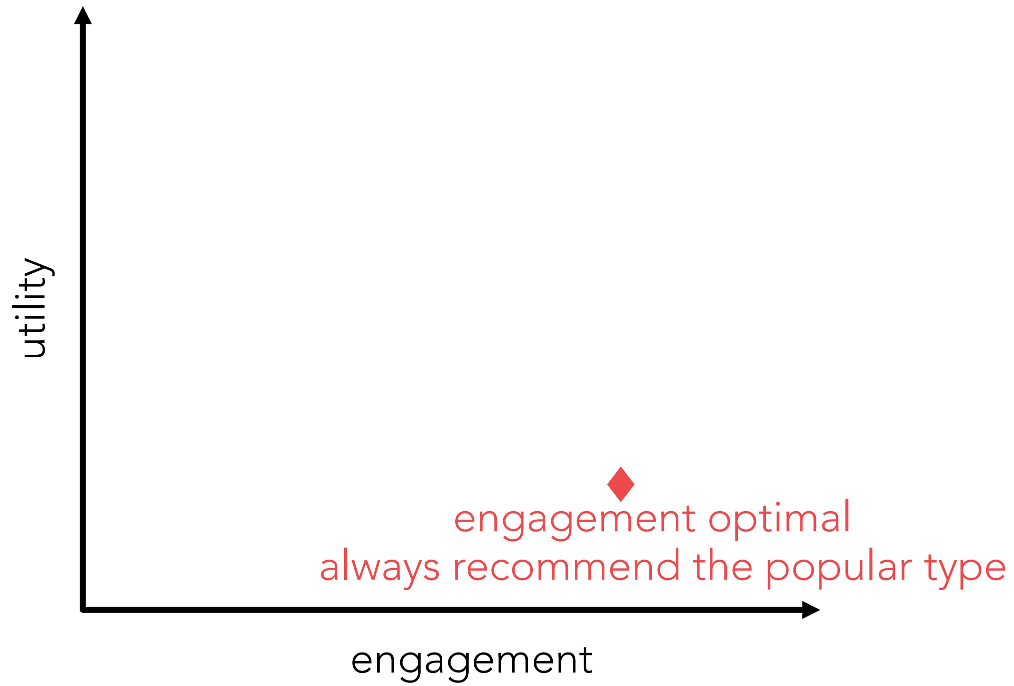

# In a Nutshell

Study a stylized model of repeated user interaction

Engagement optimization may lead to recommending only popular items, a.k.a, "popularity bias"

Utility optimization requires recommending a diverse set of options a.k.a "forced exploration"

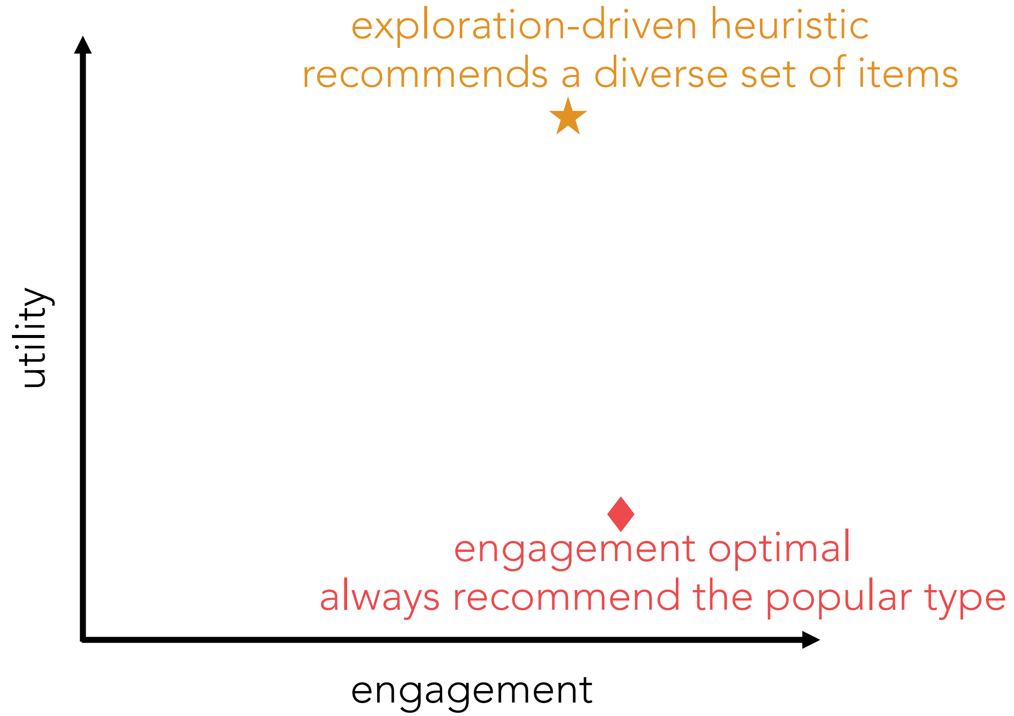

# In a Nutshell

Study a stylized model of repeated user interaction

Engagement optimization may lead to recommending only popular items, a.k.a, "popularity bias"

Utility optimization requires recommending a diverse set of options a.k.a "forced exploration"

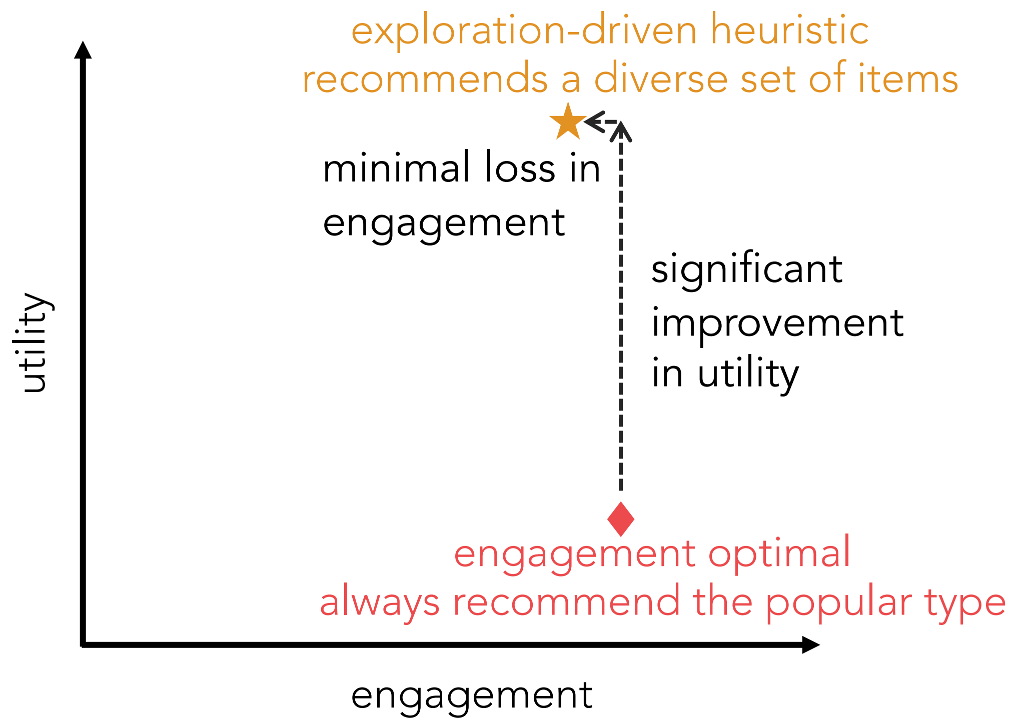


Given the ability to recommend multiple items, one can optimize for utility without directly measuring it, and without incurring substantial reduction in engagement

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

user with unknown preference

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)
  - $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)
  - $U_{\tau_1}$: utility of the recommended option 1
  - $U_{\tau_2}$: utility of the recommended option 2

$C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} \; U_i + \varepsilon_i$ $\qquad \varepsilon_i$: zero mean Gumbel noise

user with unknown preference

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

$C = \mathrm{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}}\ U_i + \varepsilon_i$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

user with unknown preference

popular

$U_P > 0$

niche

$U_N \sim F_N, \mathbb{E}[U_N] = 0$

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

$C = \mathrm{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}}\ U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

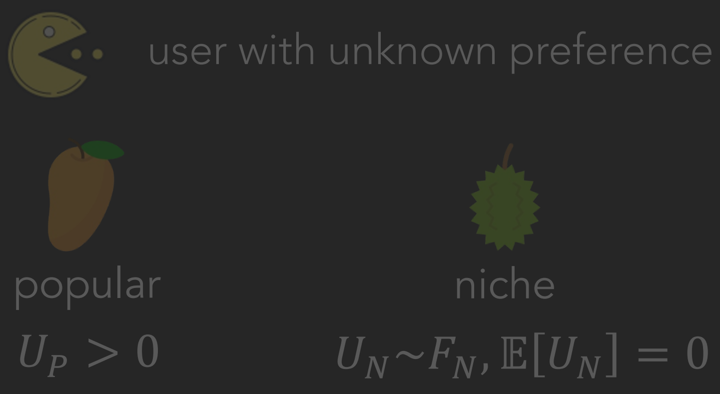

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_{\emptyset}$: utility of the outside option (assume that $U_{\emptyset} = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}}\ U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

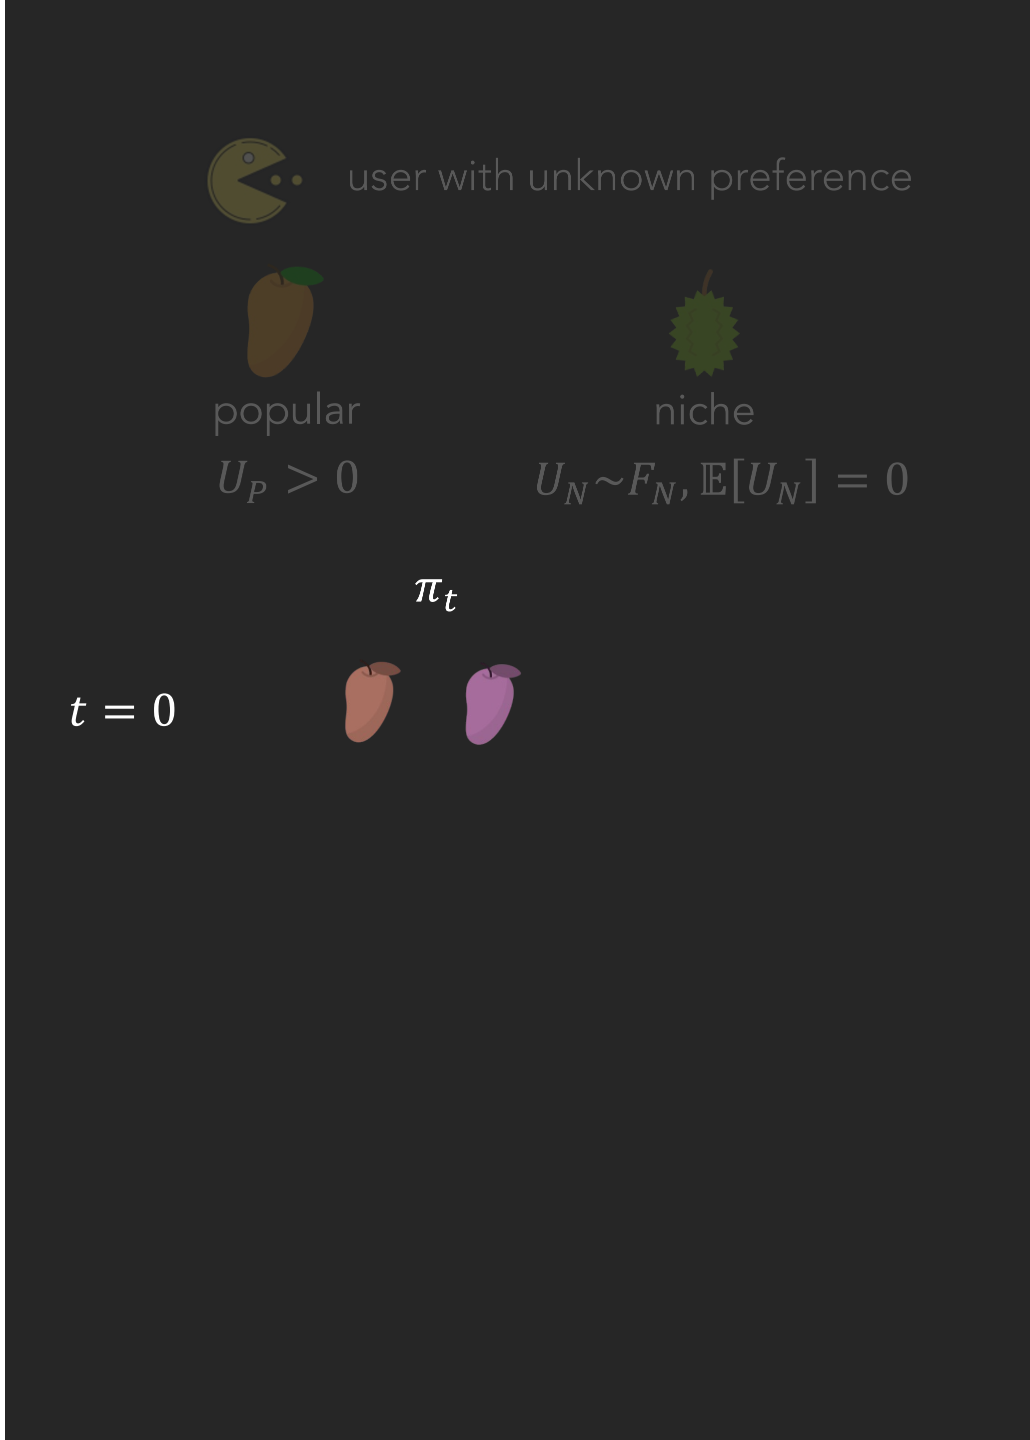

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} \; U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

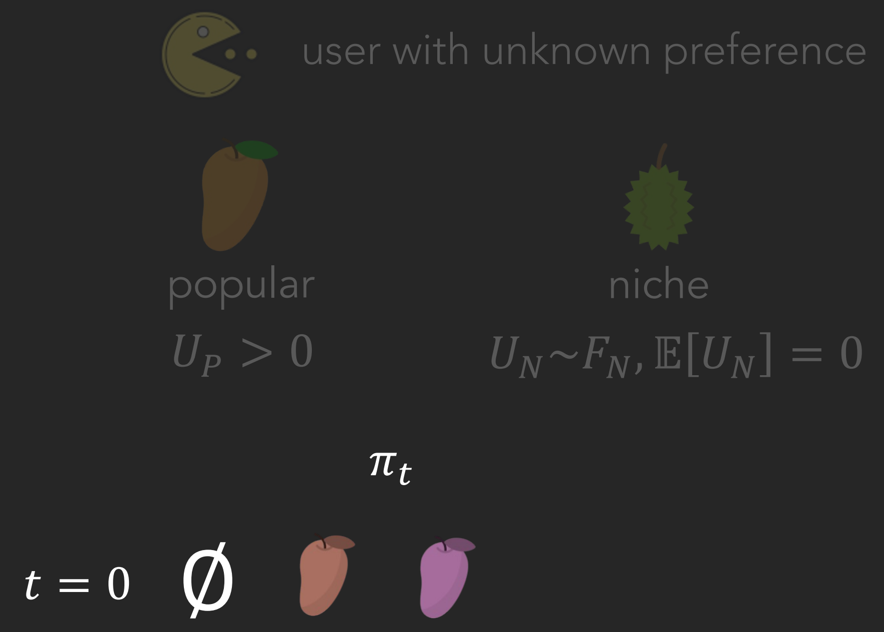

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \operatorname{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} U_i + \varepsilon_i$ $\quad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

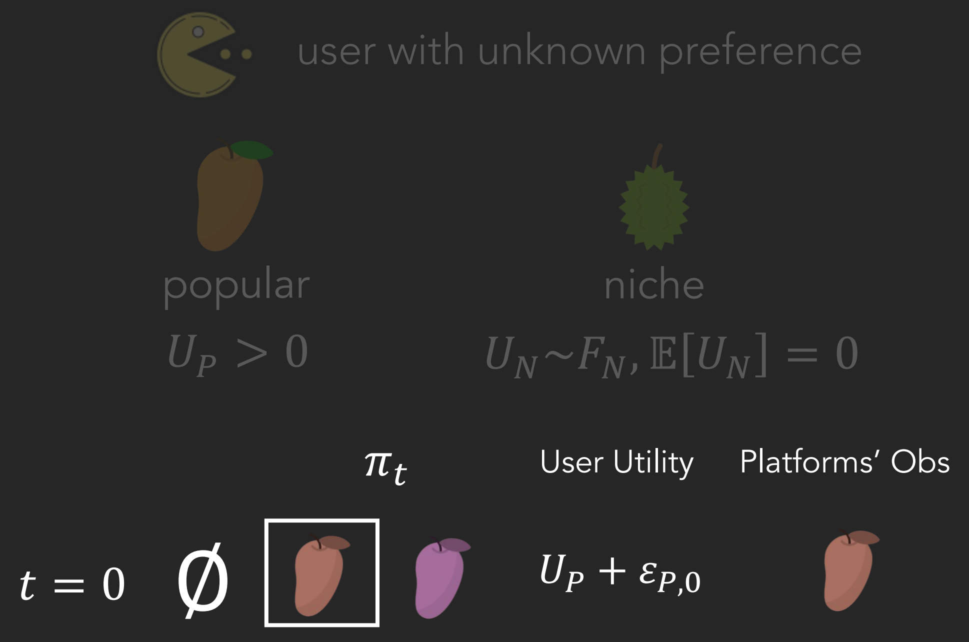

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} U_i + \varepsilon_i$ $\quad\quad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

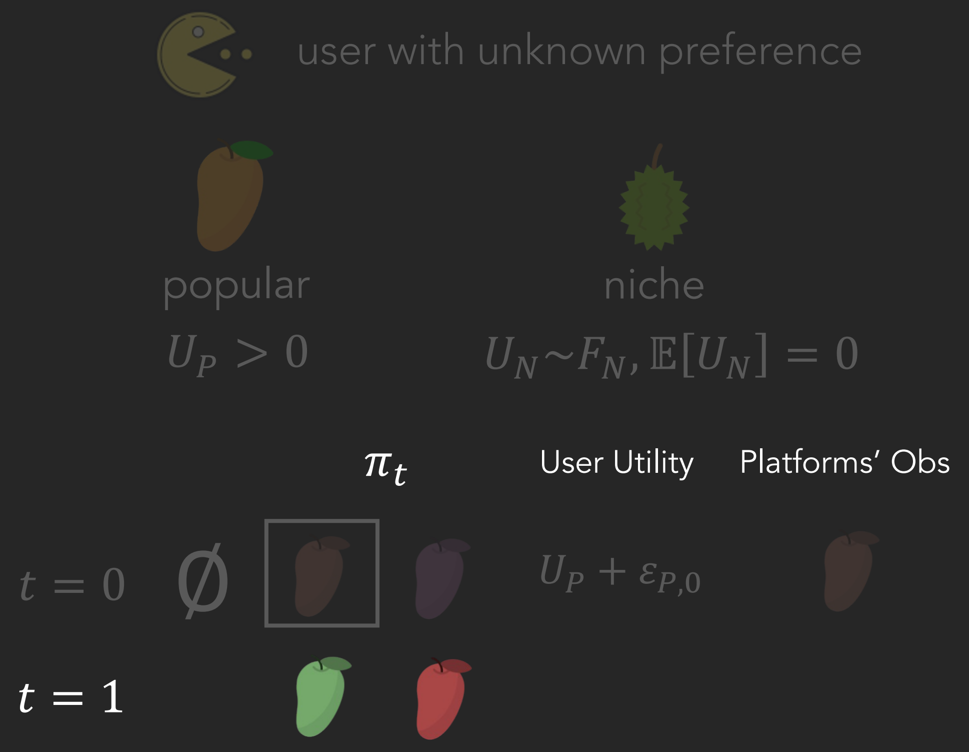

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \mathrm{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} \; U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

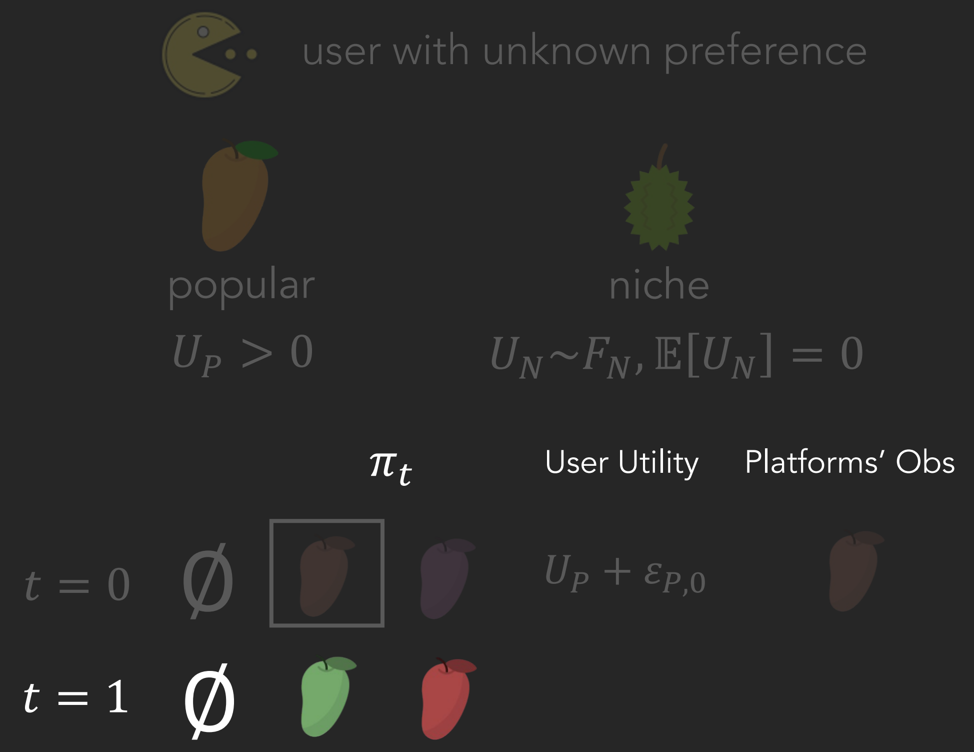

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

  $C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} \; U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

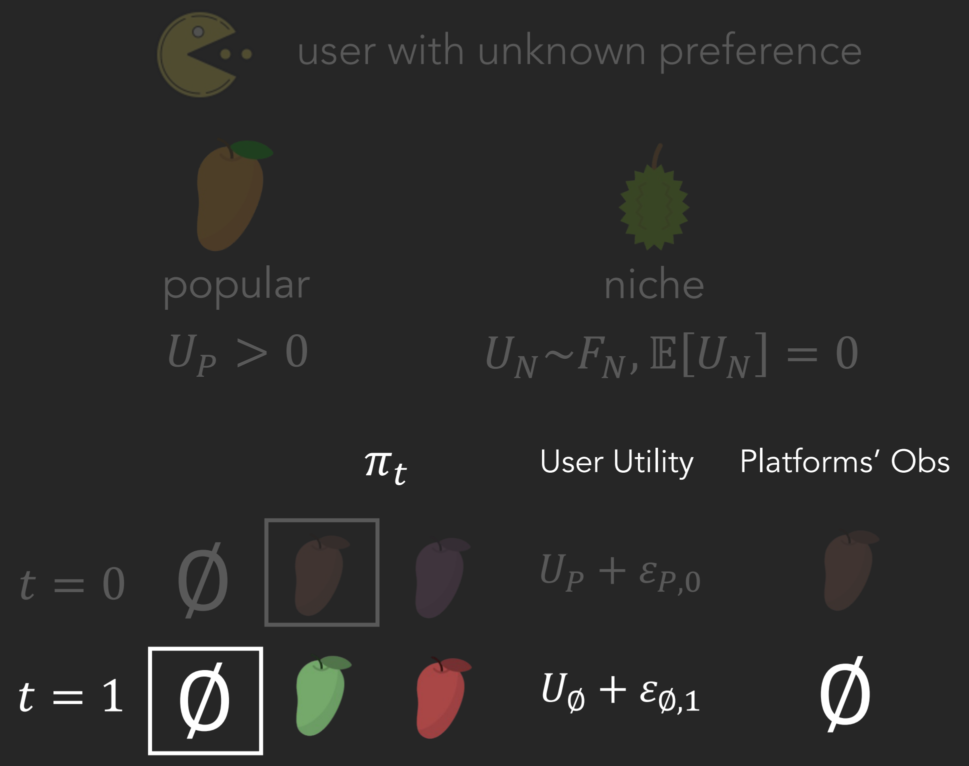

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)
  - $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)
  - $U_{\tau_1}$: utility of the recommended option 1
  - $U_{\tau_2}$: utility of the recommended option 2

$C = \text{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} \; U_i + \varepsilon_i$ $\quad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

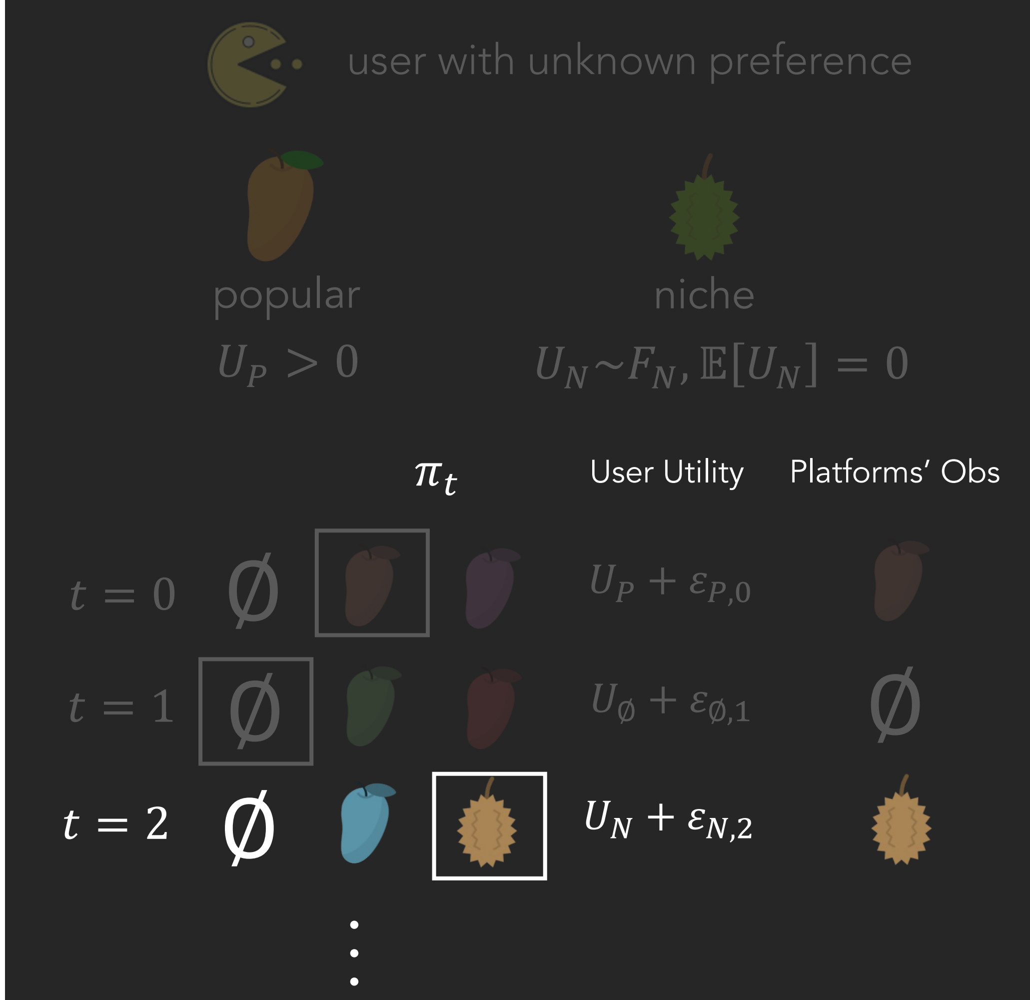

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)

  $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)

  $U_{\tau_1}$: utility of the recommended option 1

  $U_{\tau_2}$: utility of the recommended option 2

$C = \operatorname{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} U_i + \varepsilon_i$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

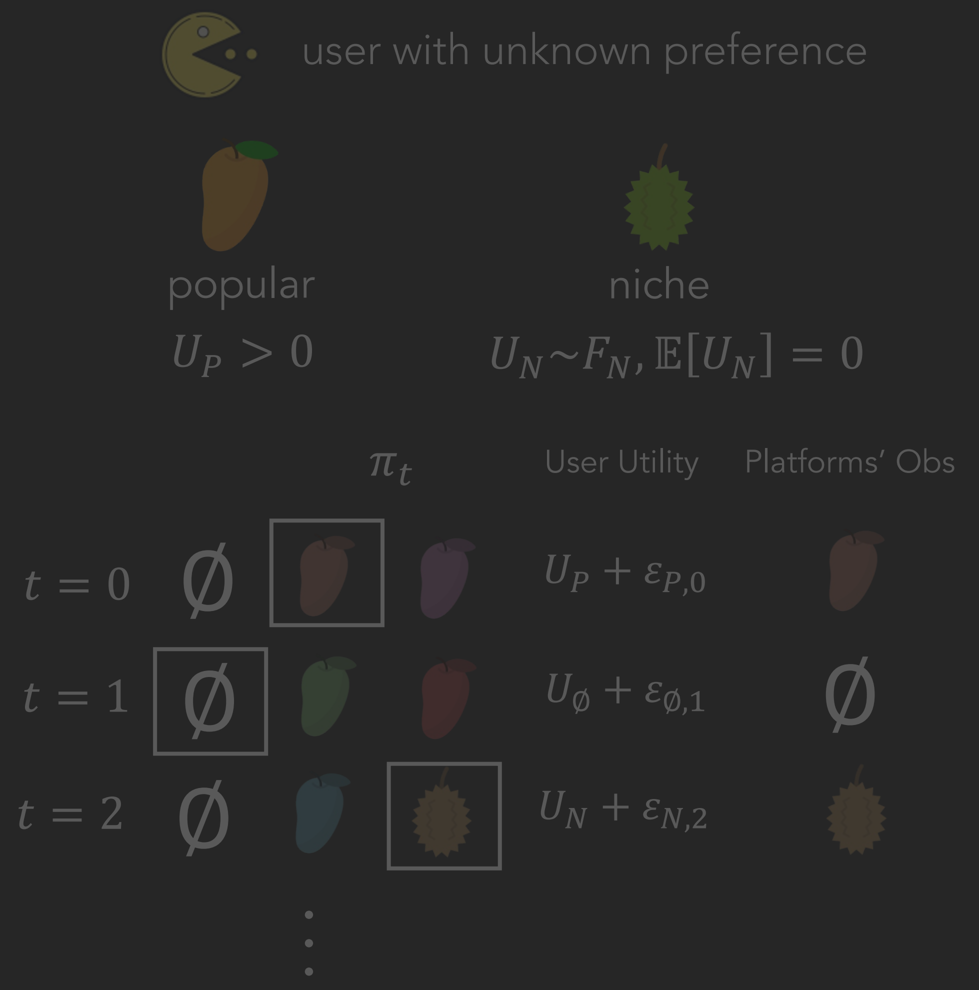


Engagement $= \mathbb{E}[\sum_{t=0}^{t=\infty} \delta^t 1(C_t \neq \emptyset)]$

# Model

- Users' Decision: Select between an outside option and the best of the recommended options (assume only two options are recommended)
  - $U_\emptyset$: utility of the outside option (assume that $U_\emptyset = 0$)
  - $U_{\tau_1}$: utility of the recommended option 1
  - $U_{\tau_2}$: utility of the recommended option 2

  $C = \operatorname{argmax}_{i \in \{\emptyset, \tau_1, \tau_2\}} U_i + \varepsilon_i$ $\qquad$ $\varepsilon_i$: zero mean Gumbel noise

- There are two types of items: popular and niche
  - popular: positive mean utility, low variance
  - niche: zero mean utility, high variance

- A RecSys repeatedly recommends an assortment of items $\pi_t$ to a user whose utility is unknown.
  - We assume an infinite horizon setting with discount factor $\delta \in [0,1]$, which captures how forward-looking the RecSys is.

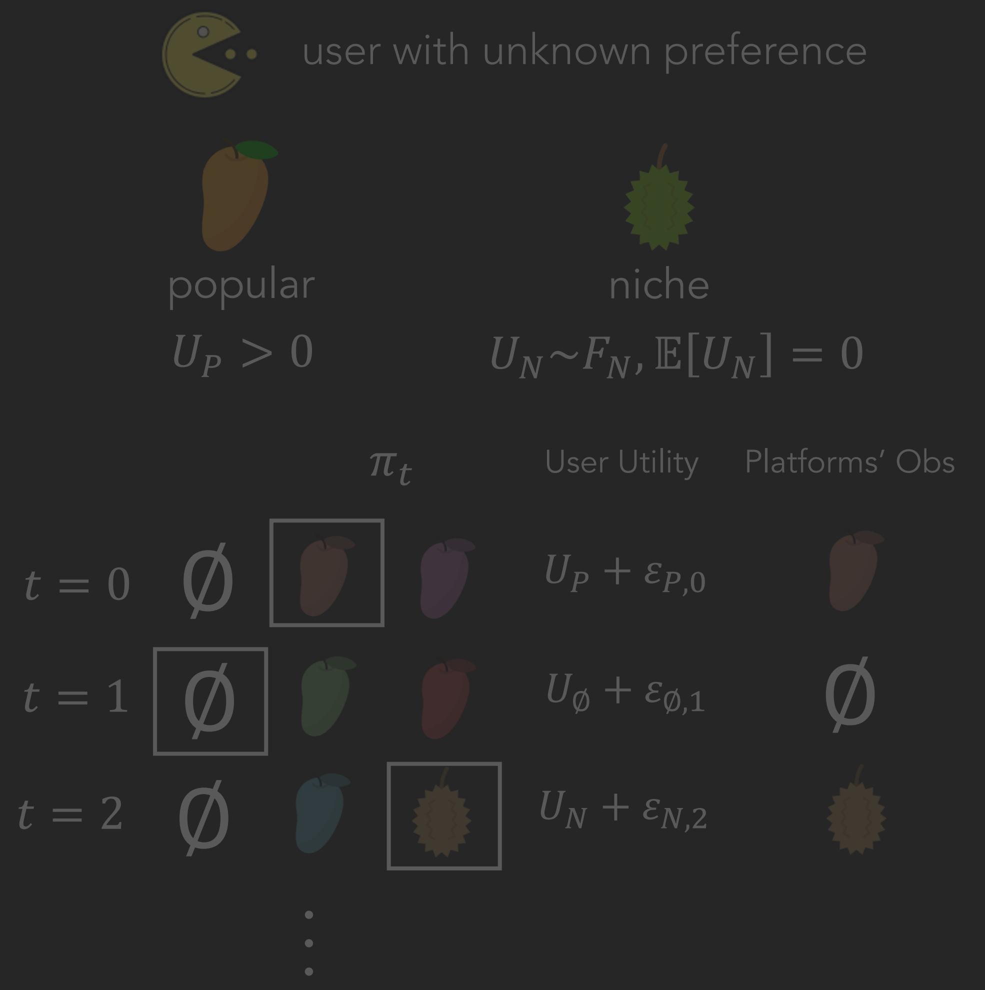


Engagement = $\mathbb{E}[\sum_{t=0}^{t=\infty} \delta^t 1(C_t \neq \emptyset)]$

Utility = $\mathbb{E}[\sum_{t=0}^{t=\infty} \delta^t (U_{C_t} + \varepsilon_{C_t,t})]$

# Toy Example

# Toy Example

- popular product with utility $U_P > 0$

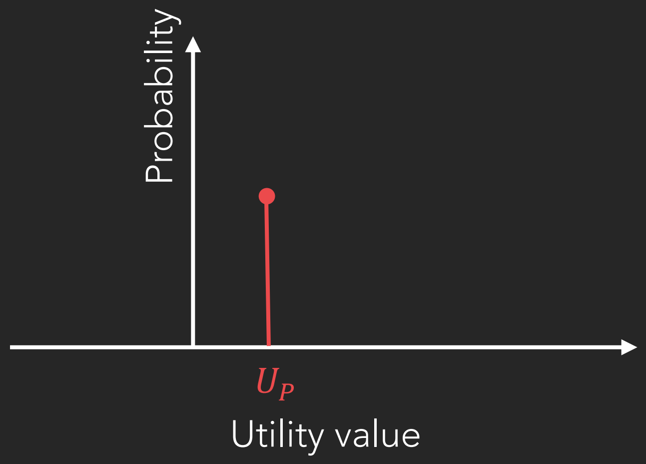

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1-p$

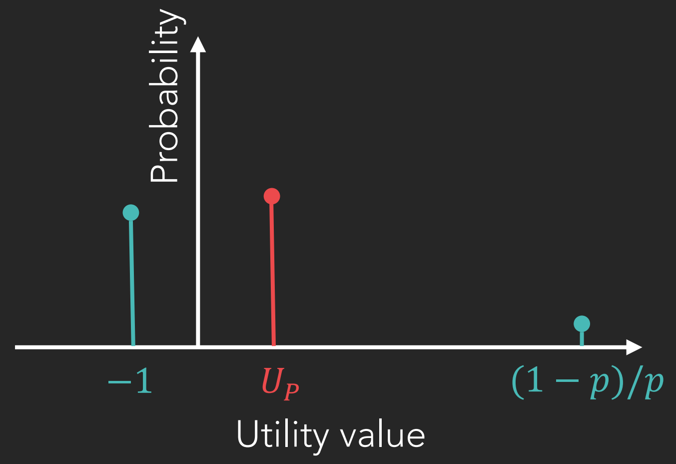

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1-p$

APP: Always Popular Policy

Similar set of recommendations

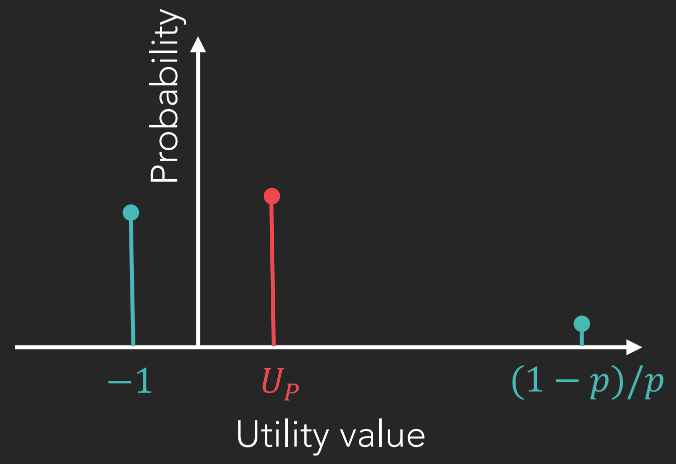

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1 - p$

APP: Always Popular Policy

Similar set of recommendations

**Theorem (APP is engagement optimal)**

Fix any discount factor $\delta \in (0,1)$, there exists a small enough $p > 0$ such that APP is the uniquely engagement optimal policy.

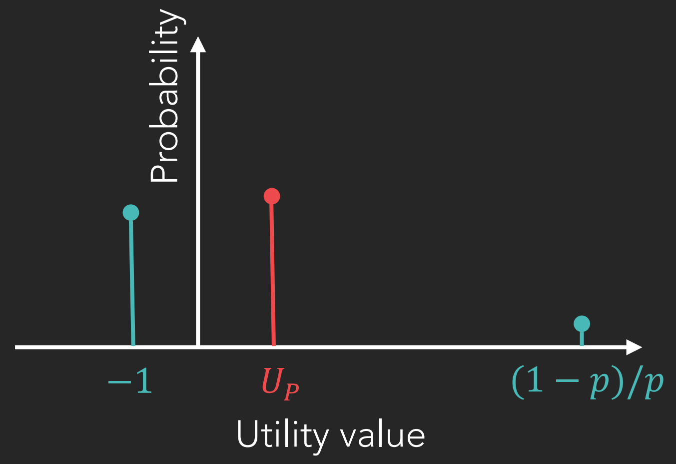

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1 - p$

APP: Always Popular Policy

Similar set of recommendations

**Theorem (APP is engagement optimal)**

Fix any discount factor $\delta \in (0,1)$, there exists a small enough $p > 0$ such that APP is the uniquely engagement optimal policy.

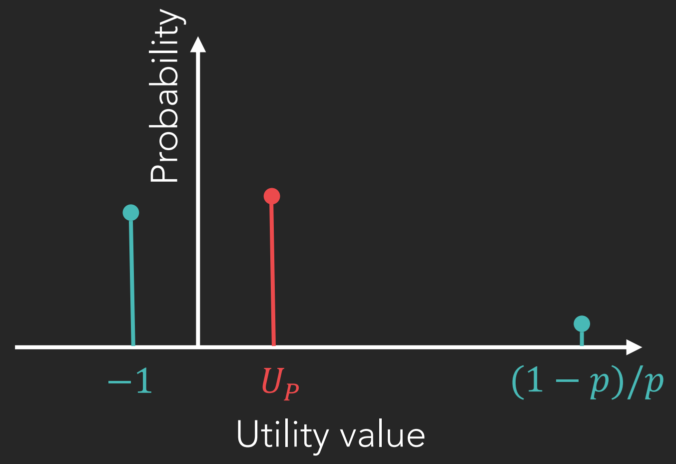


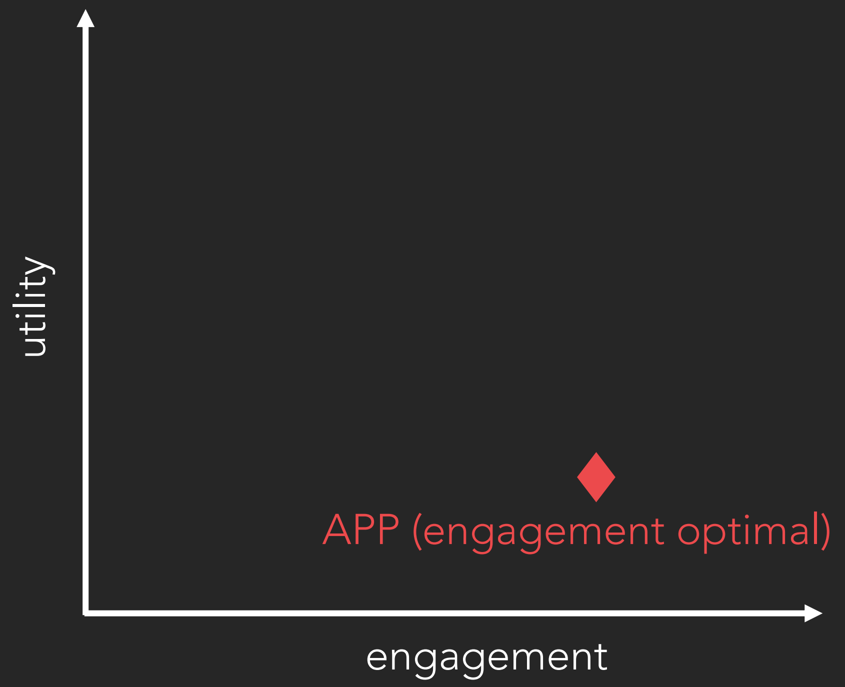

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1 - p$

APP: Always Popular Policy

Similar set of recommendations

**Theorem (APP is engagement optimal)**

Fix any discount factor $\delta \in (0,1)$, there exists a small enough $p > 0$ such that APP is the uniquely engagement optimal policy.

PEAR: Posterior-based Exploration-driven Adaptive Recos

Diverse set of recommendations

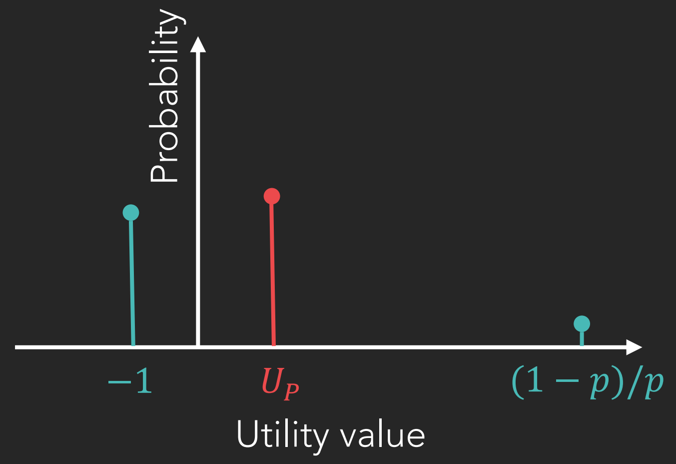


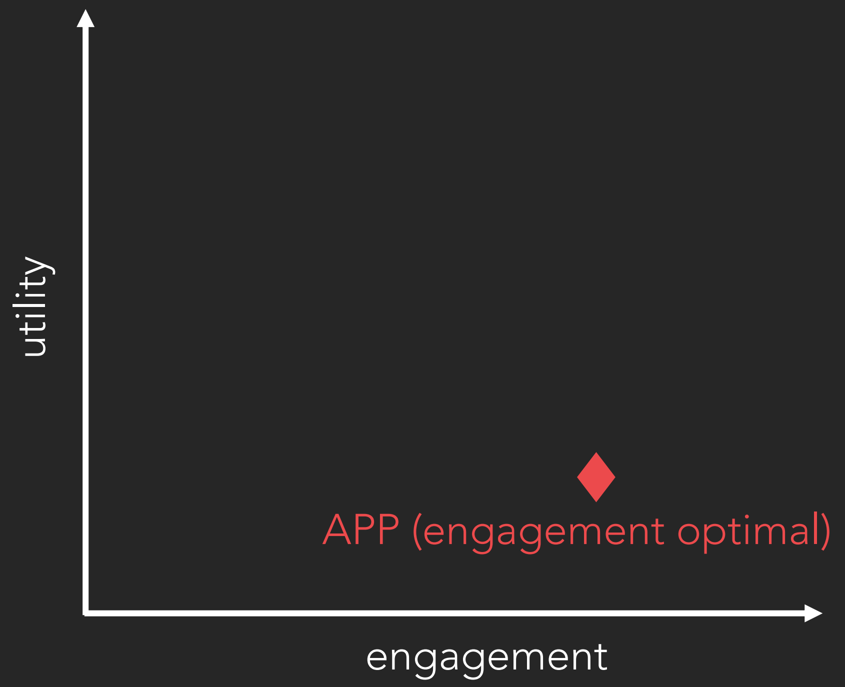

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1 - p$

APP: Always Popular Policy

Similar set of recommendations

**Theorem (APP is engagement optimal)**

Fix any discount factor $\delta \in (0,1)$, there exists a small enough $p > 0$ such that APP is the uniquely engagement optimal policy.

PEAR: Posterior-based Exploration-driven Adaptive Recos

Diverse set of recommendations

**Theorem (PEAR is near utility optimal)**

For sufficiently large discount factor $\delta \in (0,1)$, PEAR is near optimal both in terms of engagement as well as utility.

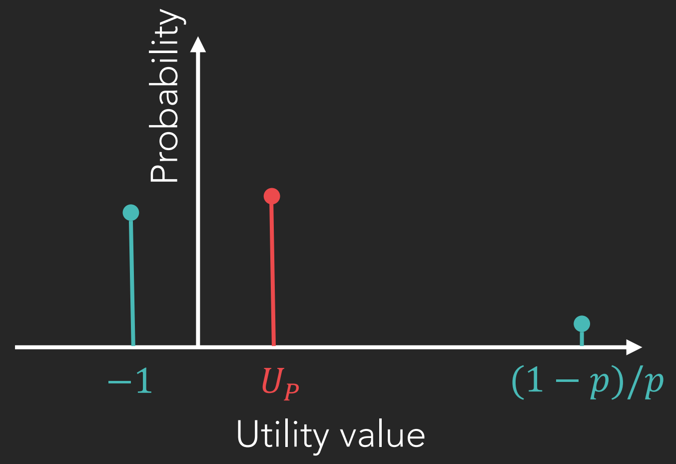


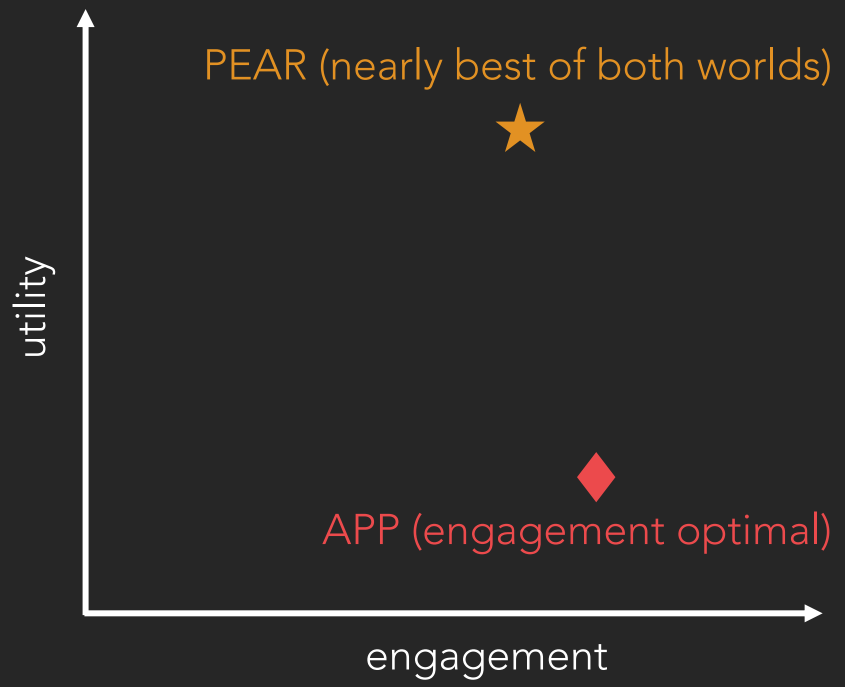

# Toy Example

- popular product with utility $U_P > 0$
- niche product has a two-point utility distribution
  - $\mathbb{P}(U_N = (1-p)/p) = p$ and $\mathbb{P}(U_N = -1) = 1 - p$

APP: Always Popular Policy

Similar set of recommendations

| |
|---|
| Theorem (APP is engagement optimal)<br>Fix any discount factor $\delta \in (0,1)$, there exists a small enough $p > 0$ such that APP is the uniquely engagement optimal policy. |

PEAR: Posterior-based Exploration-driven Adaptive Recos

Diverse set of recommendations

| |
|---|
| Theorem (PEAR is near utility optimal)<br>For sufficiently large discount factor $\delta \in (0,1)$, PEAR is near optimal both in terms of engagement as well as utility. |

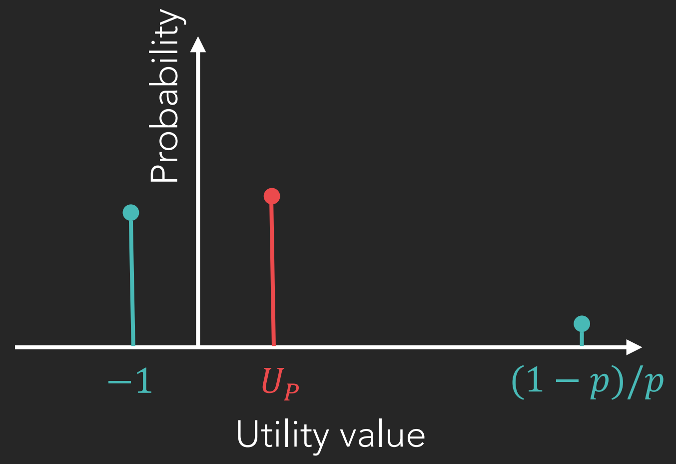


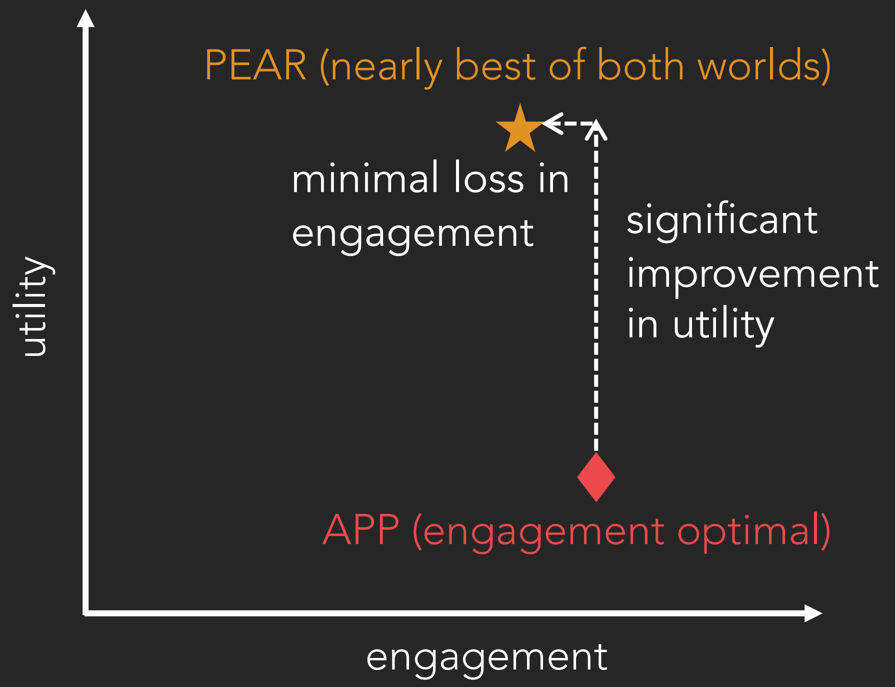

# Beyond the Toy Example

# Beyond the Toy Example

- Platform does not know the utility values

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type
- DICE: DIverse-then-CustomizEd recommendations
  - Set an exploration window of length $\tau > 0$.
  - Recommend a mix of both popular and niche types of item for the first $\tau$ periods.
  - Based on the user choices, decide the best kind of product for them.

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type
- DICE: DIverse-then-CustomizEd recommendations
  - Set an exploration window of length $\tau > 0$.
  - Recommend a mix of both popular and niche types of item for the first $\tau$ periods.
  - Based on the user choices, decide the best kind of product for them.

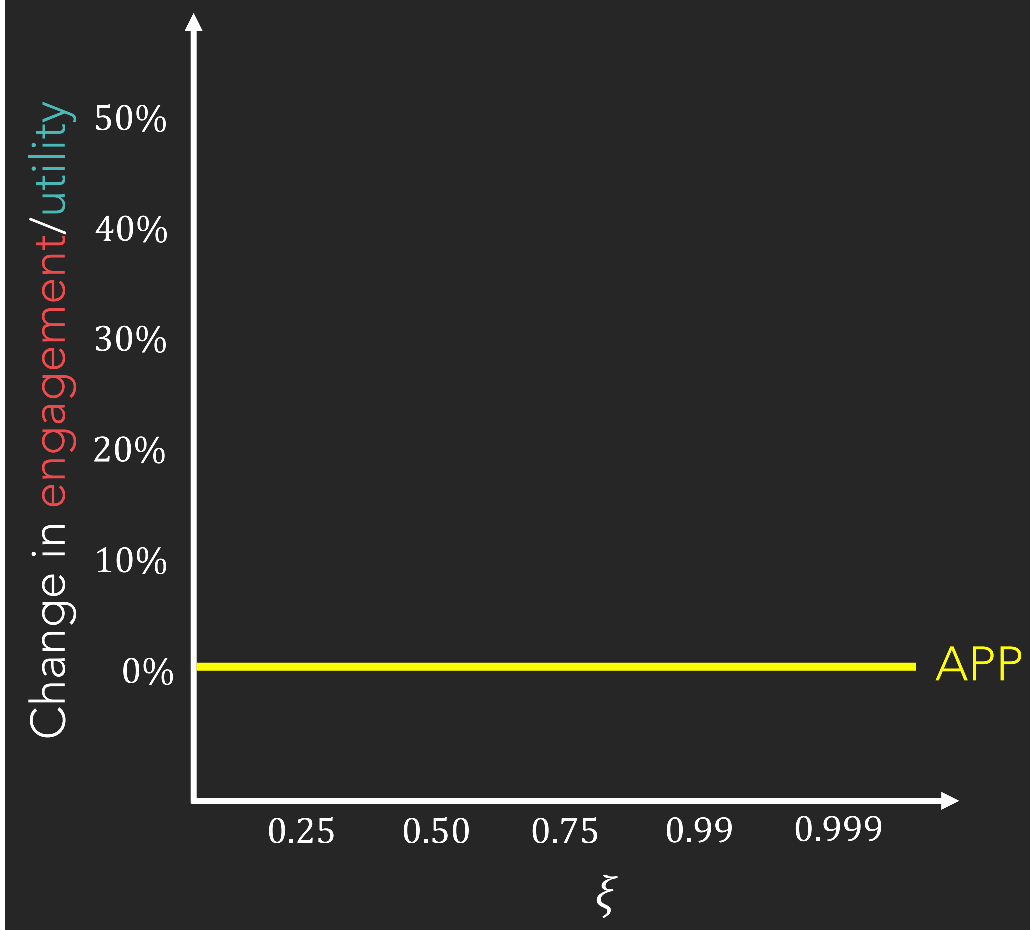

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type
- DICE: DIverse-then-CustomizEd recommendations
  - Set an exploration window of length $\tau > 0$.
  - Recommend a mix of both popular and niche types of item for the first $\tau$ periods.
  - Based on the user choices, decide the best kind of product for them.

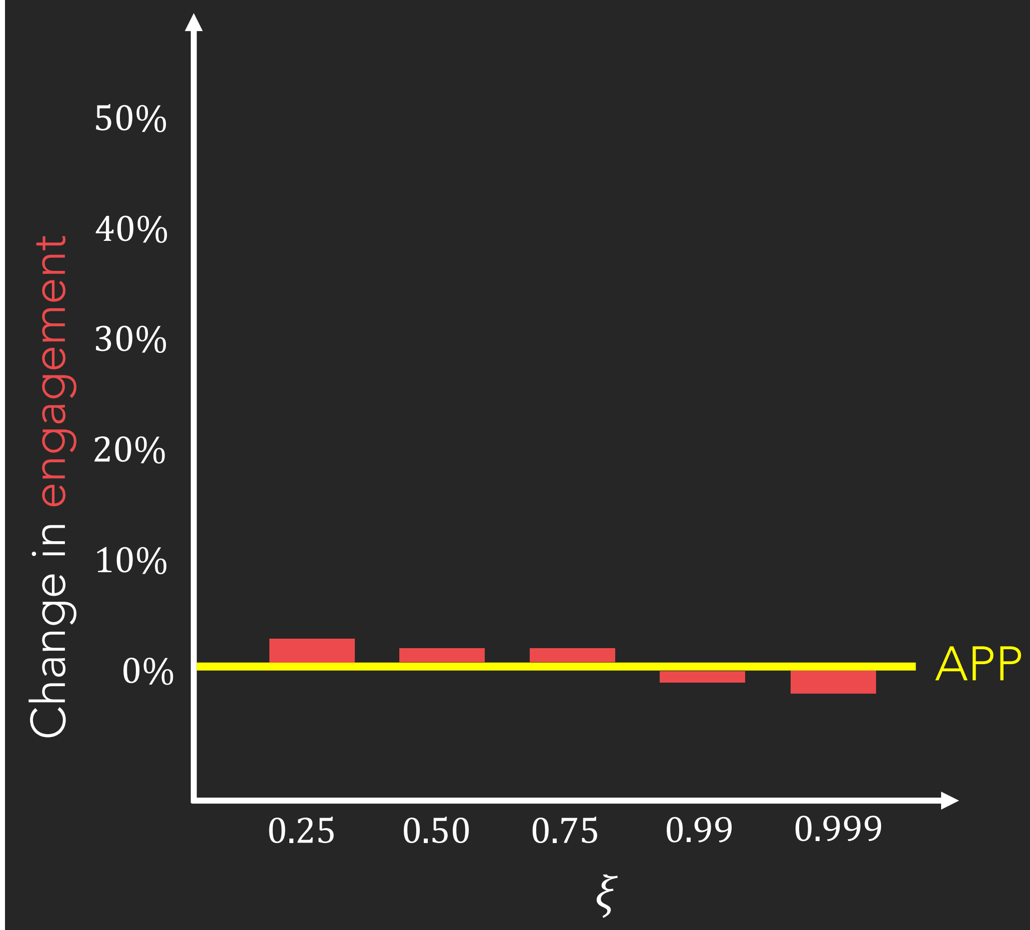

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type
- DICE: DIverse-then-CustomizEd recommendations
  - Set an exploration window of length $\tau > 0$.
  - Recommend a mix of both popular and niche types of item for the first $\tau$ periods.
  - Based on the user choices, decide the best kind of product for them.

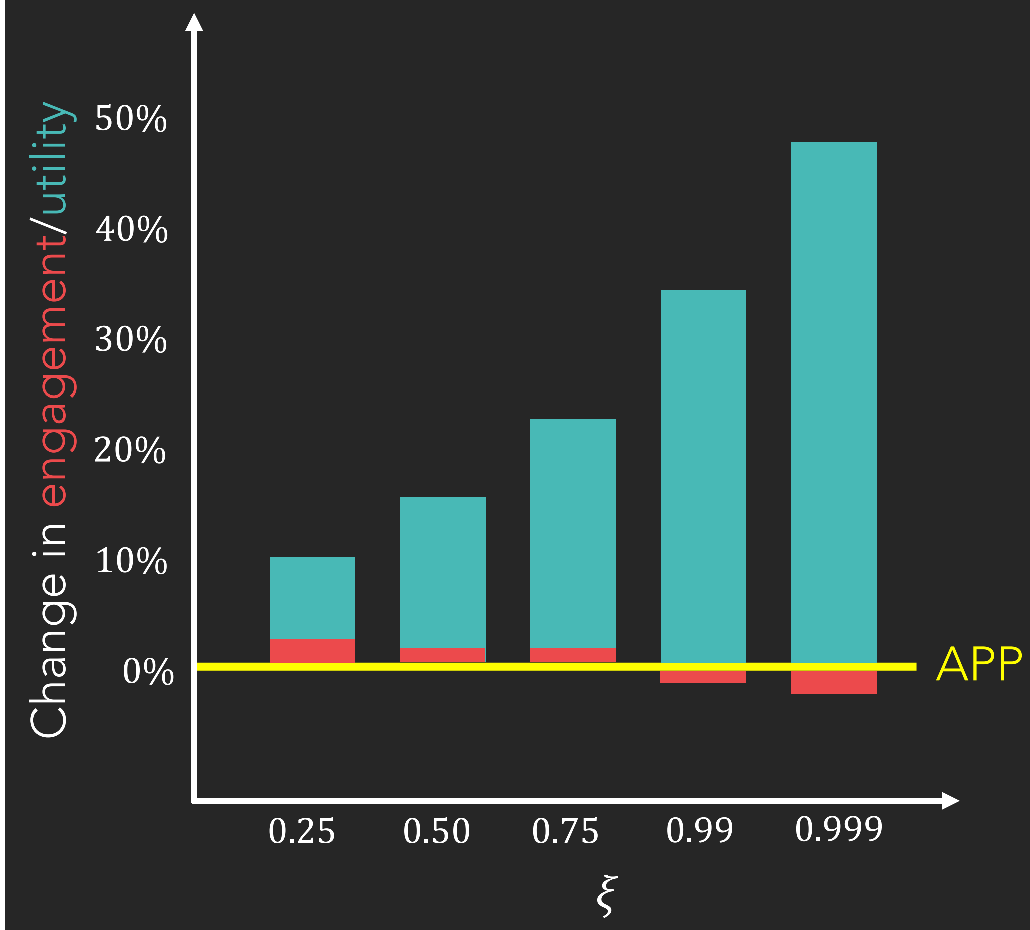

# Beyond the Toy Example

- Platform does not know the utility values
- popular product with utility $U_P > 0$
- niche product has utility $U_N$ distributed as a Generalized Pareto distribution with scale parameter $\xi$
- APP: Always Popular Policy
  - Always recommend items of the popular type
- DICE: DIverse-then-CustomizEd recommendations
  - Set an exploration window of length $\tau > 0$.
  - Recommend a mix of both popular and niche types of item for the first $\tau$ periods.
  - Based on the user choices, decide the best kind of product for them.

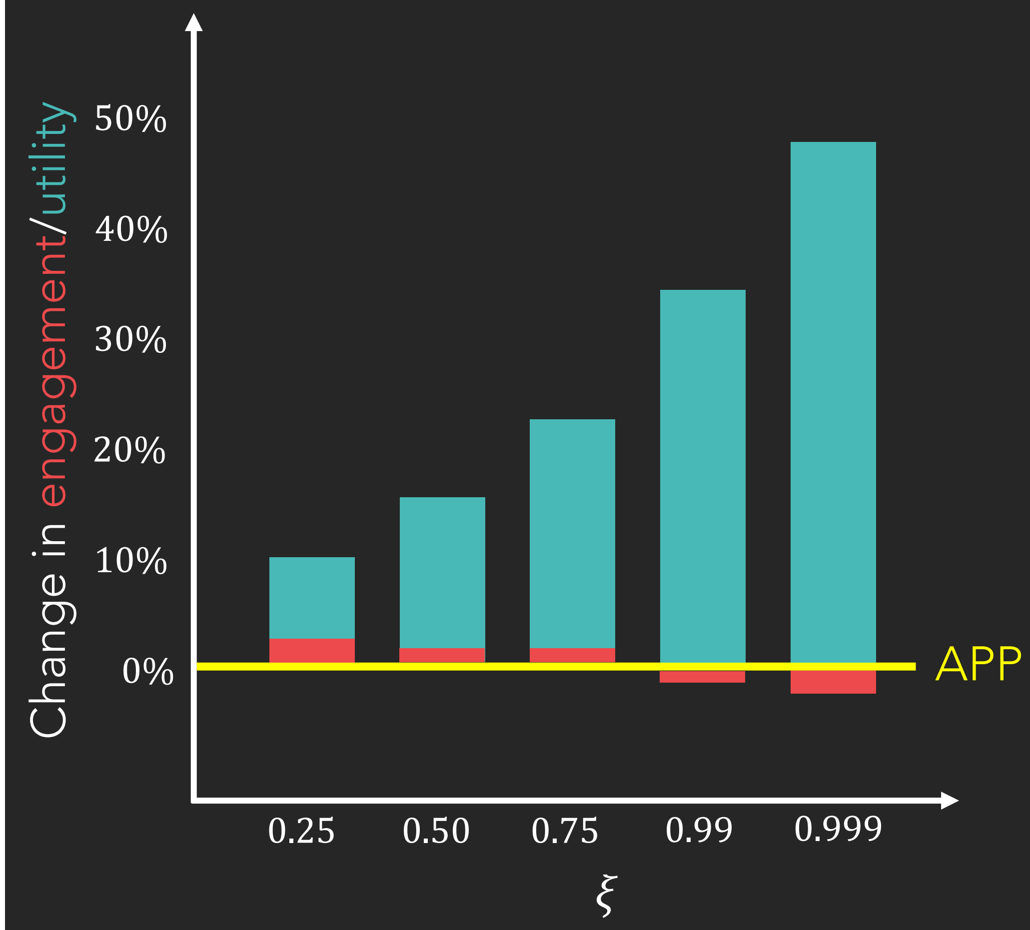


Given the ability to recommend multiple items, one can optimize for utility without directly measuring it, and without incurring substantial reduction in engagement

# Conclusion and Future Work

Icon credit: Icons made by Freepik, Taimoor D, VectorSlab, juicy_fish, Oz Mora, Kise1ki from www.flaticon.com

# Conclusion and Future Work

- Summary
  - Study a stylized model of repeated user interaction
  - Identify and formalize misalignment  engagement and utility maximization
  - Platforms can optimize for utility by providing diverse/exploratory recommendations

# Conclusion and Future Work

- Summary
  - Study a stylized model of repeated user interaction
  - Identify and formalize misalignment  engagement and utility maximization
  - Platforms can optimize for utility by providing diverse/exploratory recommendations

- Future Work
  - Tractably incorporate many different facets of the recommendation ecosystems: multiple product types, inconsistent preferences of users, etc.
  - Better design of recommendation systems to allow increased user exploration/autonomy

Icon credit: Icons made by Freepik, Taimoor D, VectorSlab, juicy_fish, Oz Mora, Kise1ki from www.flaticon.com

# Conclusion and Future Work

- Summary
  - Study a stylized model of repeated user interaction
  - Identify and formalize misalignment  engagement and utility maximization
  - Platforms can optimize for utility by providing diverse/exploratory recommendations

- Future Work
  - Tractably incorporate many different facets of the recommendation ecosystems: multiple product types, inconsistent preferences of users, etc.
  - Better design of recommendation systems to allow increased user exploration/autonomy

# Thank You!

Icon credit: Icons made by Freepik, Taimoor D, VectorSlab, juicy_fish, Oz Mora, Kise1ki from www.flaticon.com