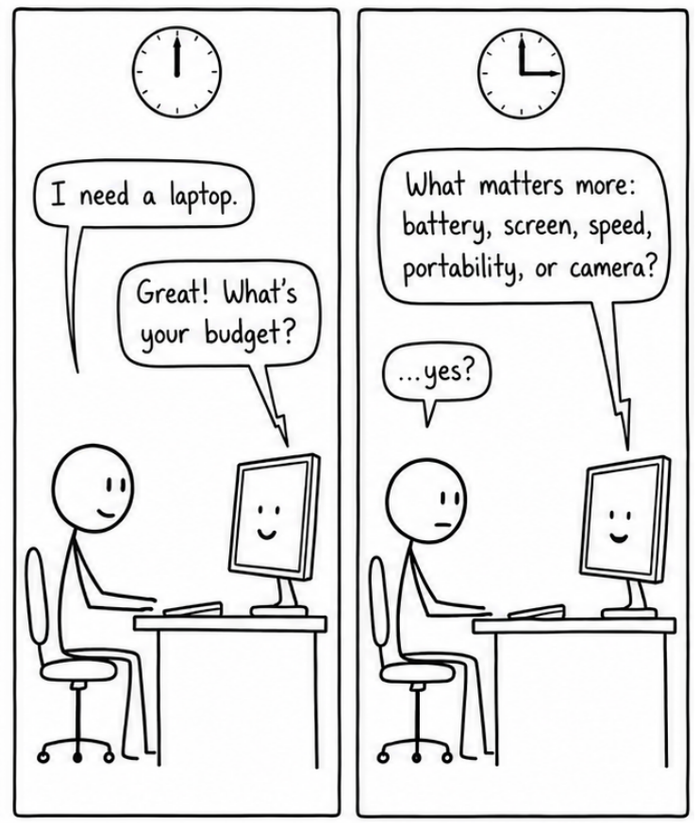
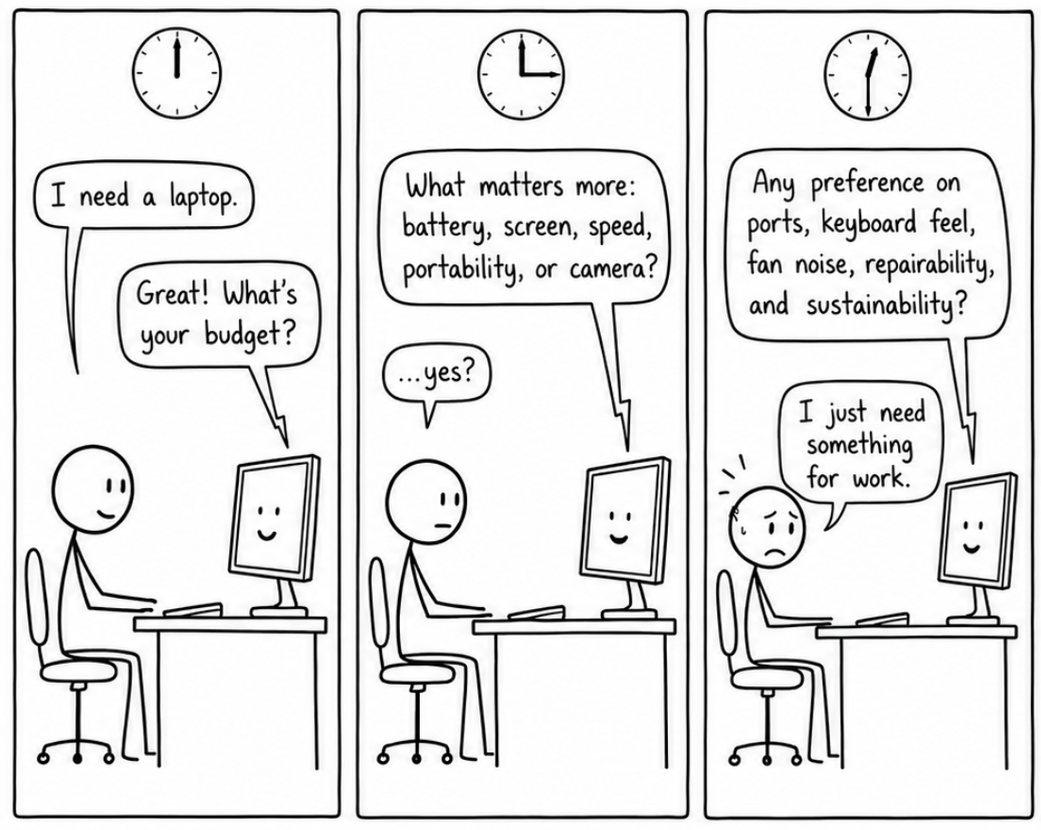
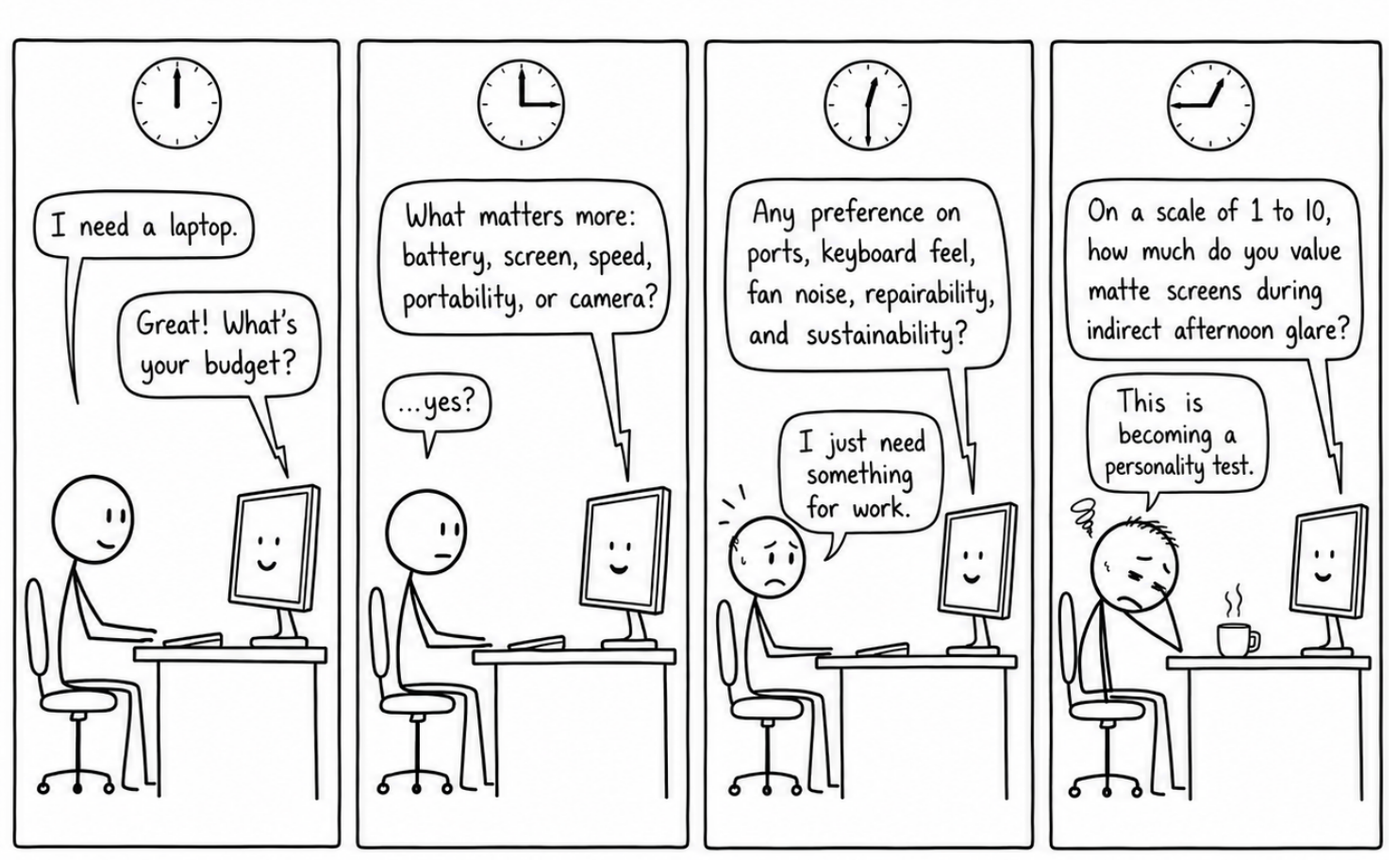
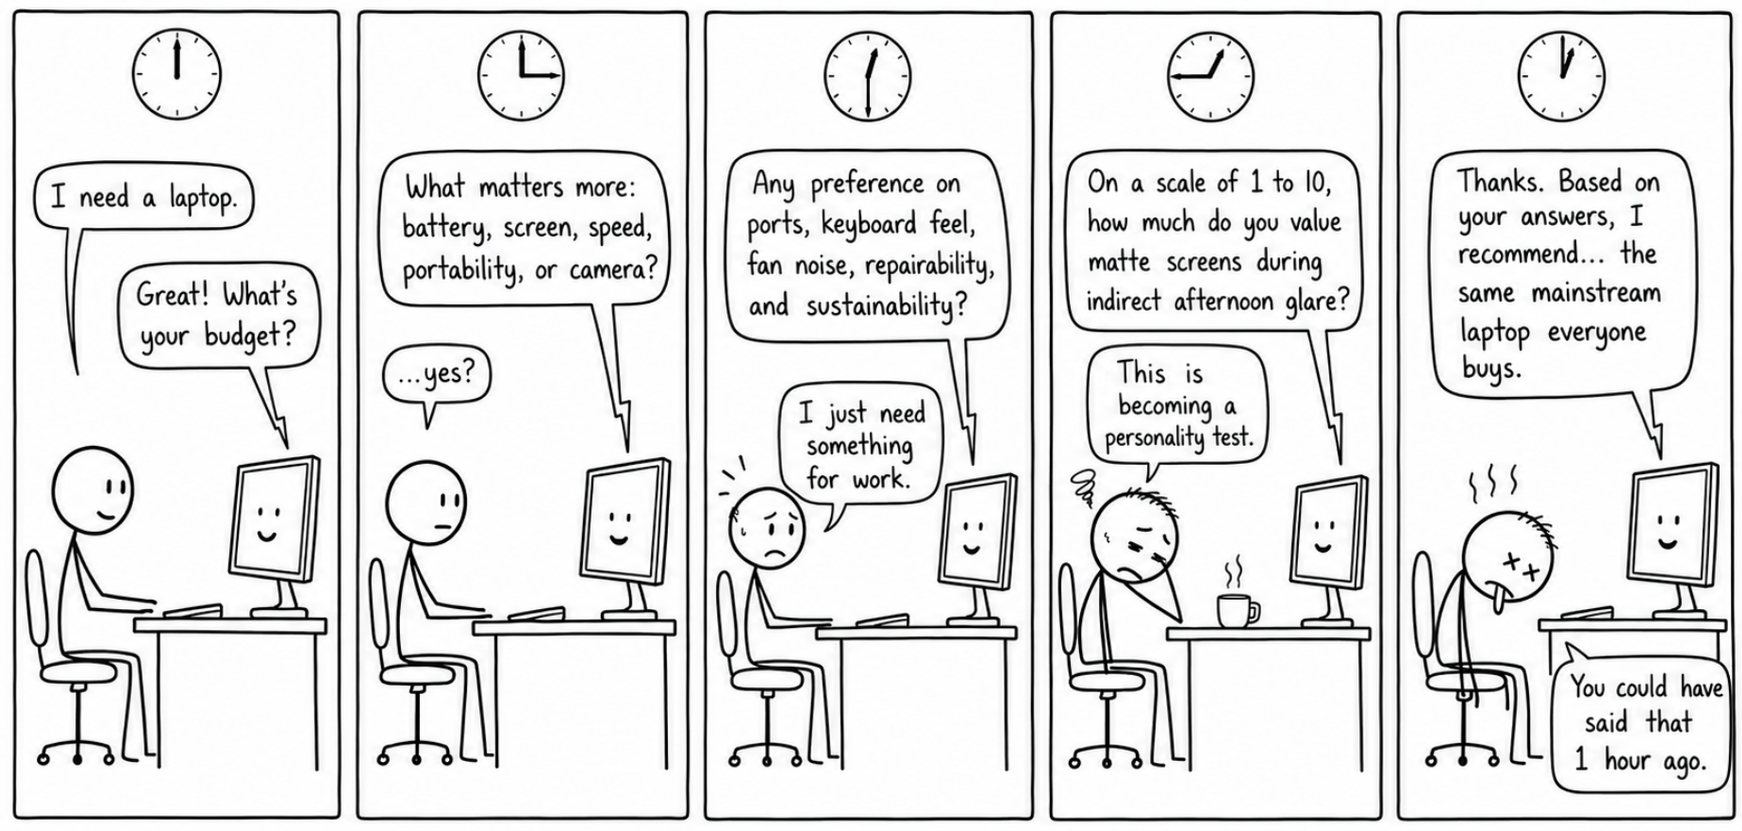

I need a laptop.
Great! What's your budget?

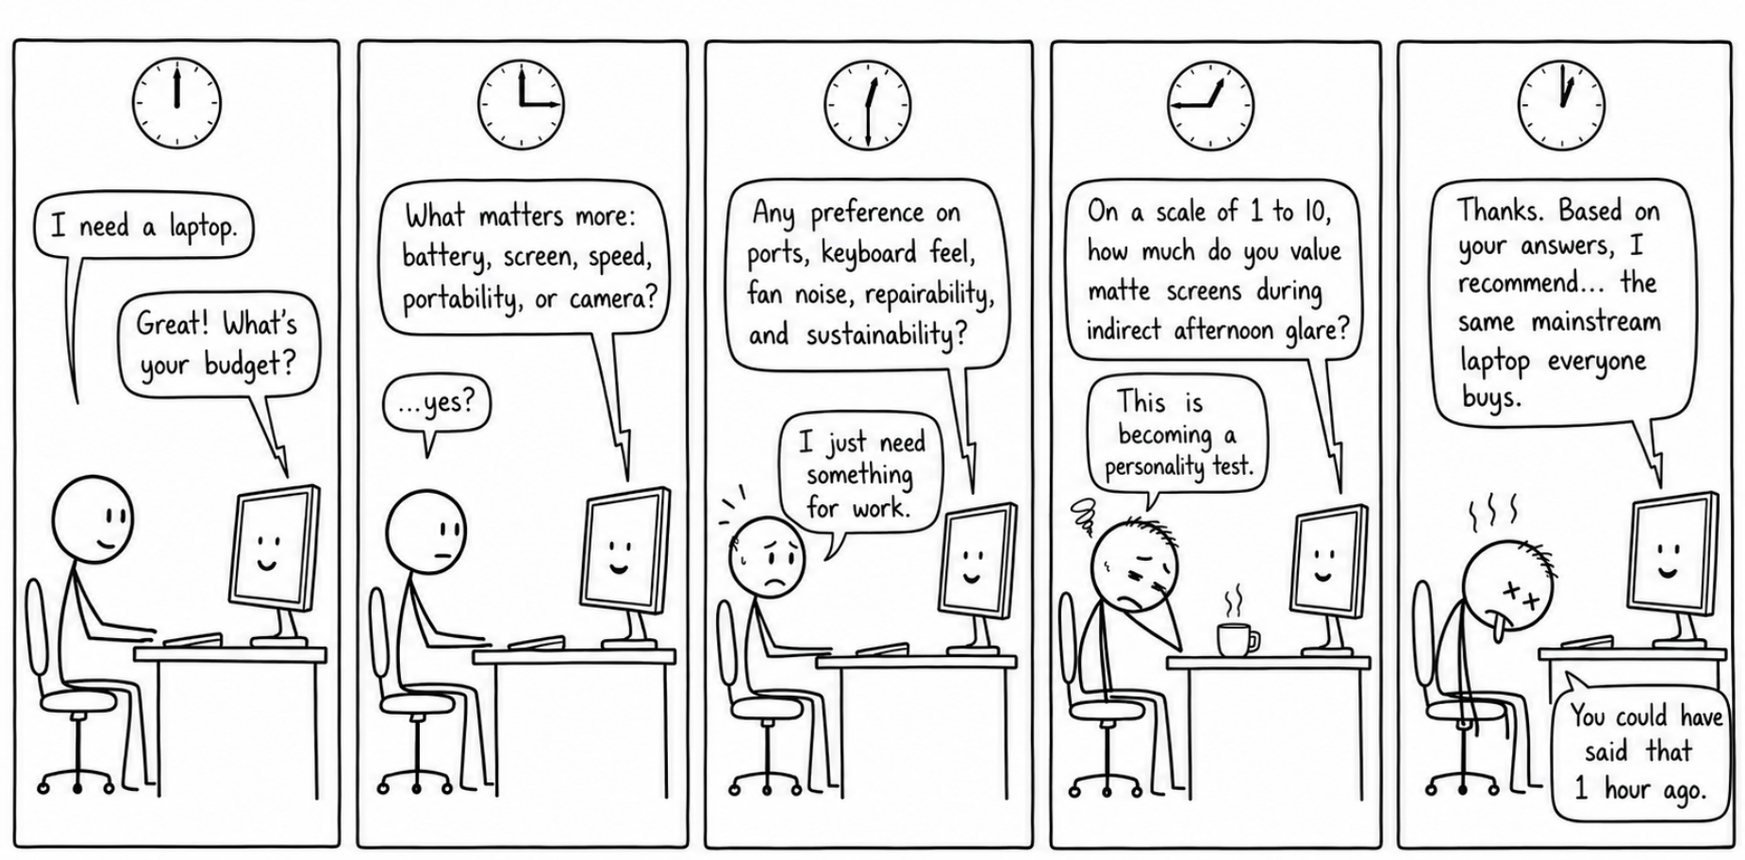

I need a laptop.
Great! What's your budget?
What matters more: battery, screen, speed, portability, or camera?
...yes?


I need a laptop.
Great! What's your budget?
What matters more: battery, screen, speed, portability, or camera?
...yes?
Any preference on ports, keyboard feel, fan noise, repairability, and sustainability?
I just need something for work.

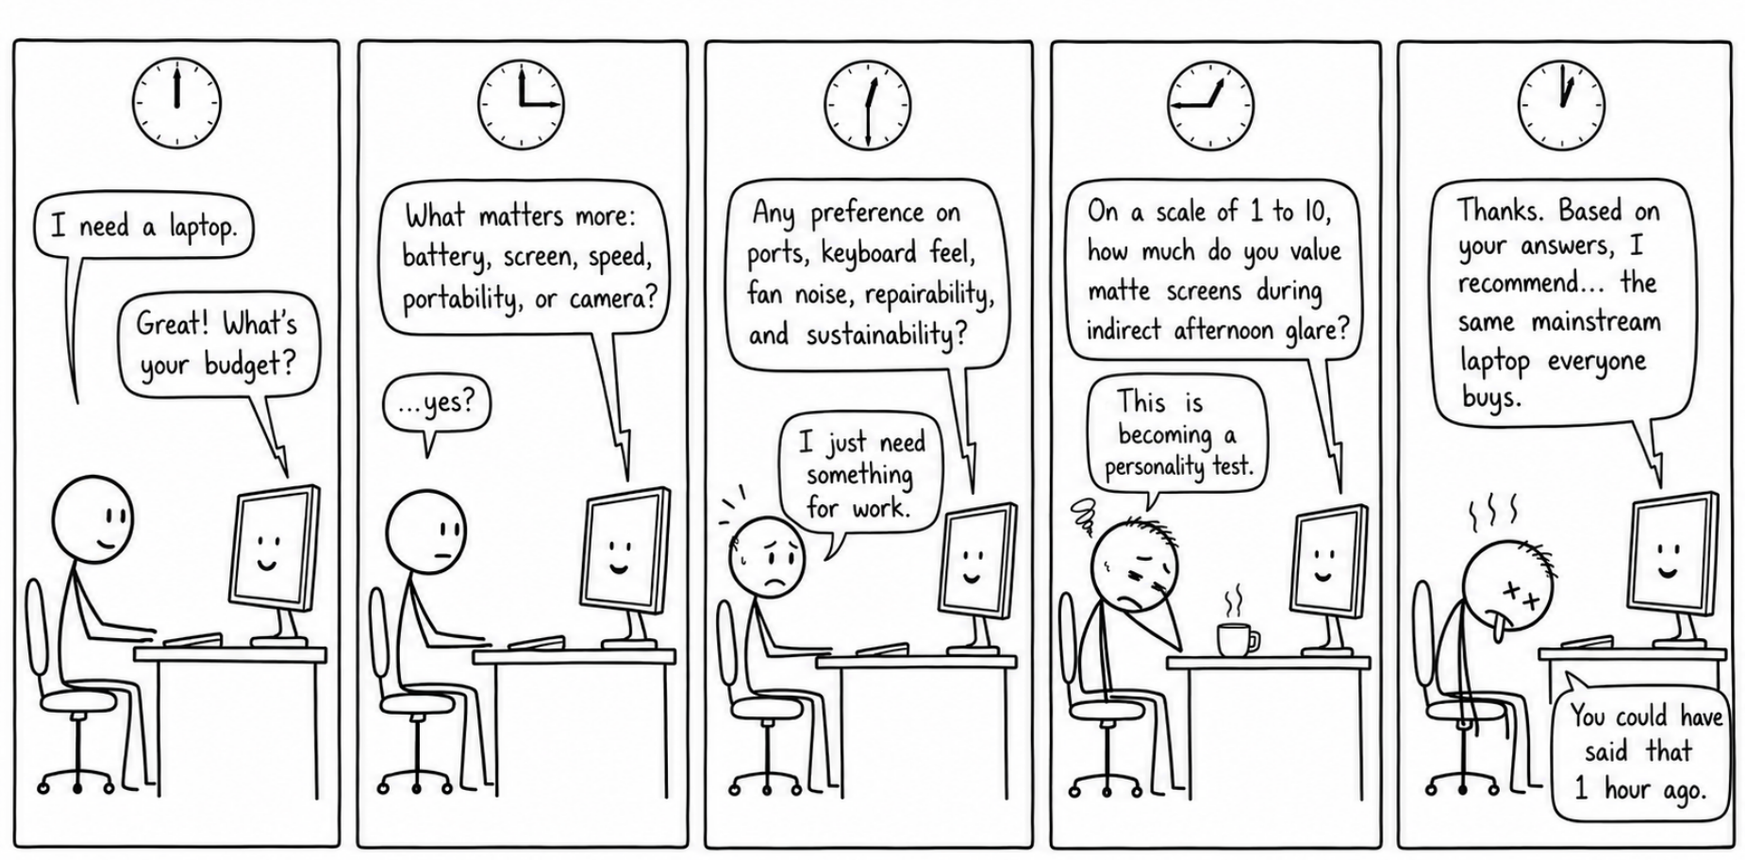

I need a laptop.
Great! What's your budget?
What matters more: battery, screen, speed, portability, or camera?
...yes?
Any preference on ports, keyboard feel, fan noise, repairability, and sustainability?
I just need something for work.
On a scale of 1 to 10, how much do you value matte screens during indirect afternoon glare?
This is becoming a personality test.


I need a laptop.
Great! What's your budget?
What matters more: battery, screen, speed, portability, or camera?
...yes?
Any preference on ports, keyboard feel, fan noise, repairability, and sustainability?
I just need something for work.
On a scale of 1 to 10, how much do you value matte screens during indirect afternoon glare?
This is becoming a personality test.
Thanks. Based on your answers, I recommend... the same mainstream laptop everyone buys.
You could have said that 1 hour ago.

How much should a conversational recommendation system converse?



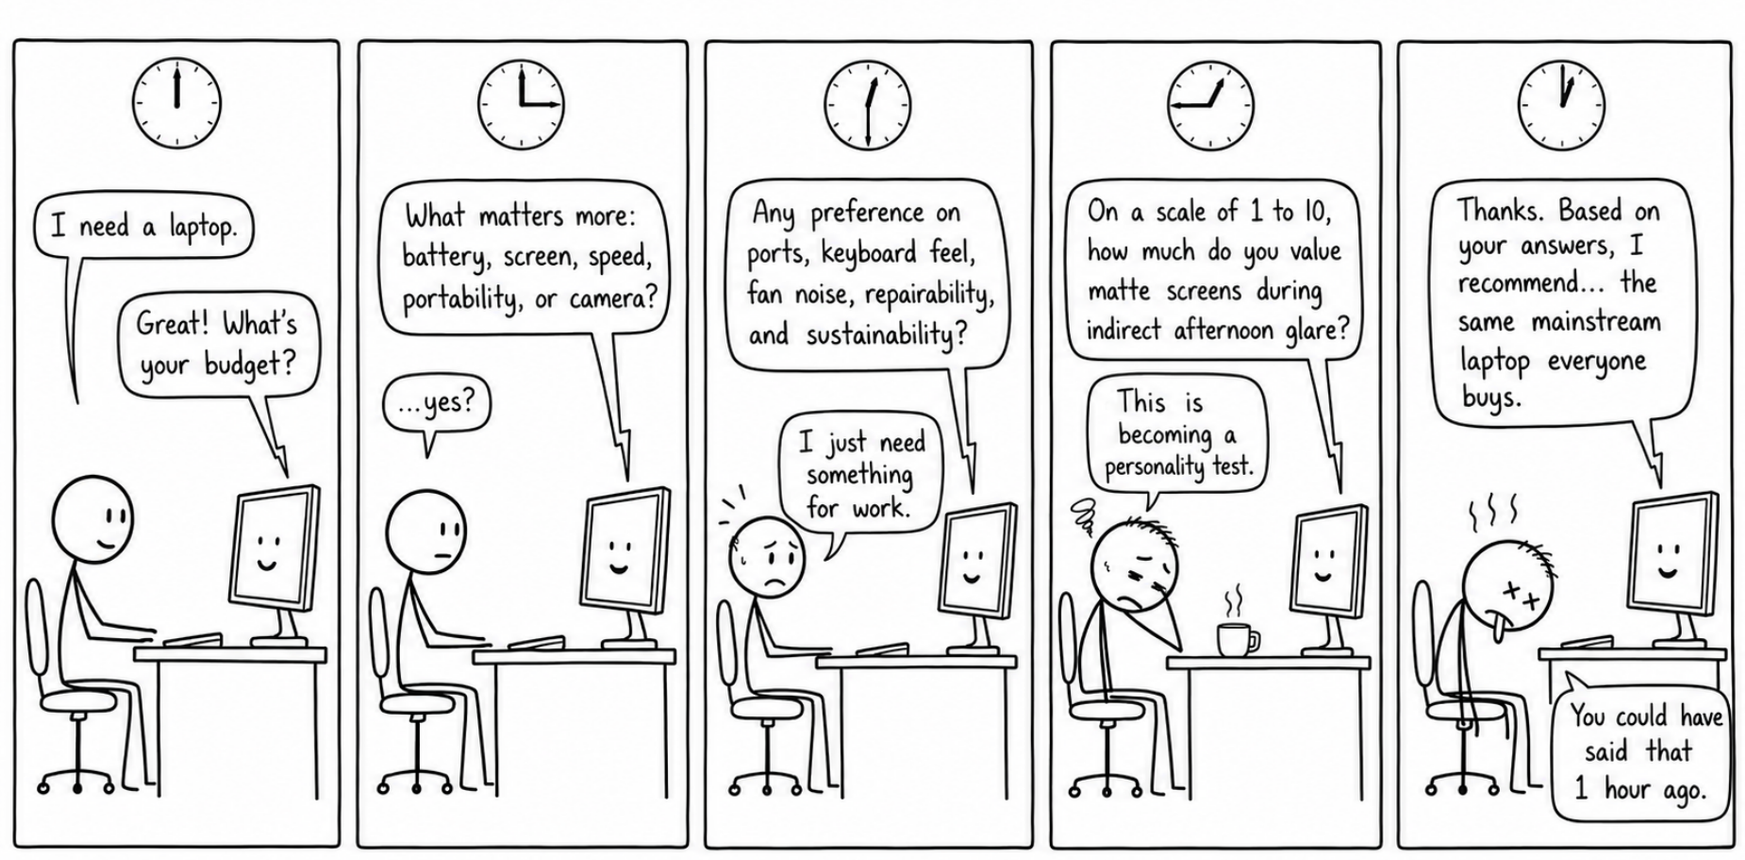
How much should a conversational recommendation system converse?



Akshit Kumar

University of Toronto

How much should a conversational recommendation system converse?



Akshit Kumar
University of Toronto

Vahideh Manshadi
Yale University

Akhilesh Tumu
Yale University

Conversational RecSys are changing how we search



Home > TV > News

May 7, 2025 5:15am PT

Netflix Overhauls TV Homepage With 'More Intuitive' Design, Adds OpenAI-Powered Chatbot

By Jennifer Maas

Conversational RecSys are changing how we search



Got a Tip? Newsletters U.S. Edition

VARIETY

Home > TV > News

May 7, 2025 5:

Netflix Overhauls TV Homepage Wit 'More Intuitive' Design, Adds OpenAI Powered Chatbot

By Jennifer Maas

November 24, 2025 Product Company

Introducing shopping research in ChatGPT

A new shopping experience that helps you find the right products for you.

Try in ChatGPT ↗

Conversational RecSys are changing how we search

Got a Tip? Newsletters U.S. Edition

VARIETY

Home > TV > News

May 7, 2025 5:

Netflix Overhauls TV Homepage Wit 'More Intuitive' Design, Adds OpenAI Powered Chatbot

By Jennifer Maas

November 24, 2025 Product Company

Introducing shopping research in ChatGPT

A new shopping experience that helps you find the right products for you.

Try in ChatGPT ↗

About News & Features Careers Sustainability & Governance Investors Press Shop

GUEST EXPERIENCE

Meet Your New Shopping Buddy: The Target App in ChatGPT

Nov 25, 2025 | 3-minute read

Copy link to share

Conversational RecSys are changing how we search

Got a Tip? Newsletters U.S. Edition

VARIETY

Home > TV > News

May 7, 2025 5:

Netflix Overhauls TV Homepage Wit 'More Intuitive' Design, Adds OpenAI Powered Chatbot

By Jennifer Maas

November 24, 2025 Product Company

Introducing shopping research in ChatGPT

A new shopping experience that helps you find the right products for you.

About News & Features Careers

GUEST EXPERIENCE

Meet Your New S in ChatGPT

Nov 25, 2025 | 3-minute read

Copy link to share

CNBC

Search quotes, news & videos WATCHLIST

MARKETS BUSINESS INVESTING TECH POLITICS VIDEO INVESTING CLUB JOIN PRO JOIN LIVESTREAM

RETAIL

Walmart teams up with Google's Gemini to make it easier for shoppers to find and buy products

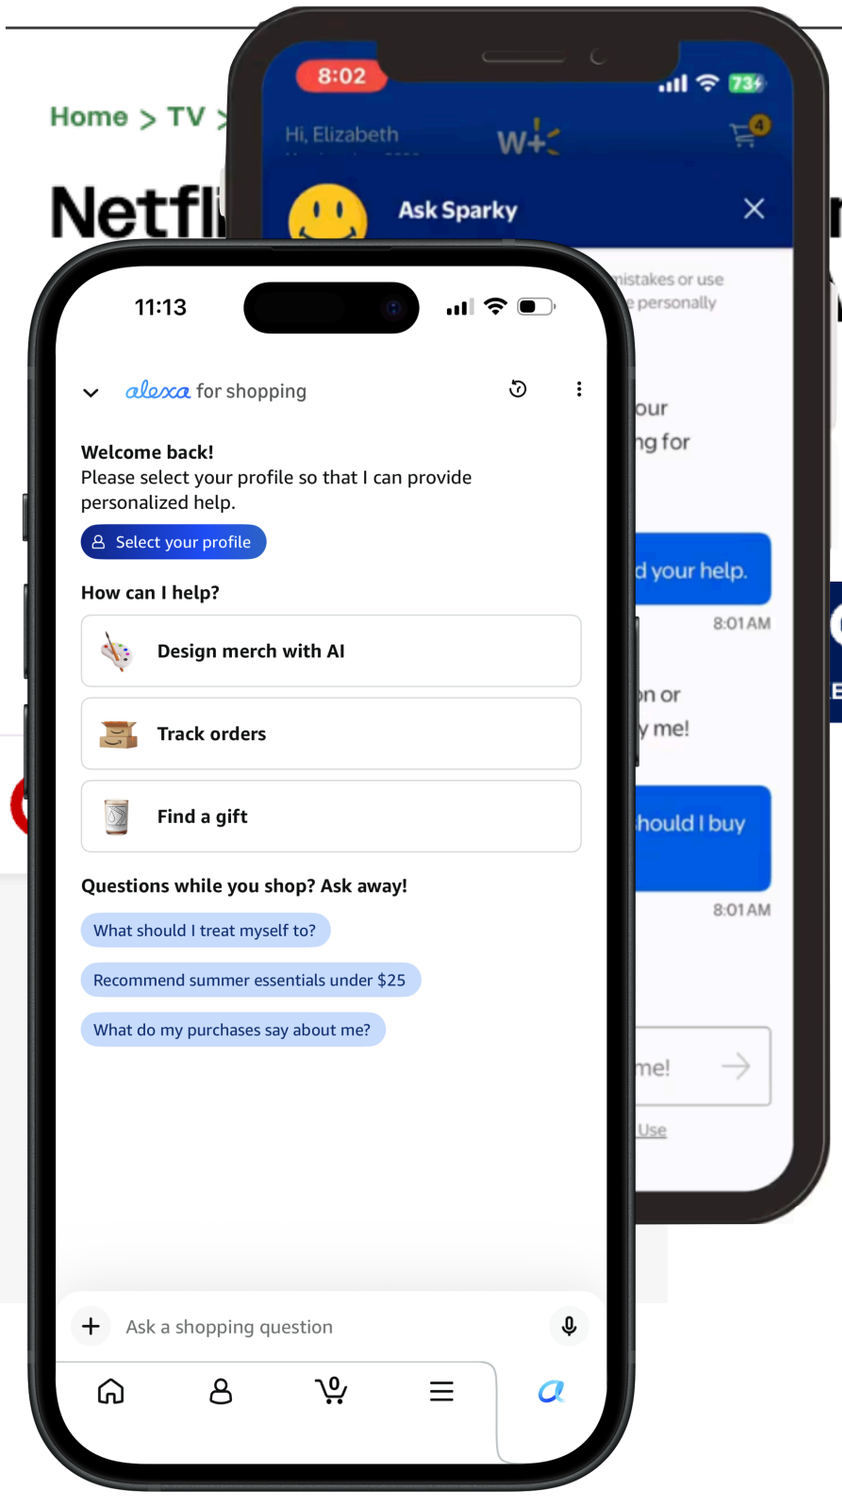
PUBLISHED SUN, JAN 11 2026•10:00 AM EST | UPDATED SUN, JAN 11 2026•10:42 AM EST



Melissa Repko
@IN/MELISSA-REPKO
@MELISSA_REPKO

SHARE

Conversational RecSys are changing how we search



Got a Tip? Newsletters U.S. Edition

VARIETY

Home > TV >

May 7, 2025 5:

Netfli mepage Wit
dds OpenAI

November 24, 2025 Product Company

Introducing shopping research in ChatGPT

A new shopping experience that helps you find the right products for you.

Search quotes, news & videos WATCHLIST

ETS BUSINESS INVESTING TECH POLITICS VIDEO INVESTING CLUB JOIN PRO JOIN LIVESTREAM

RETAIL

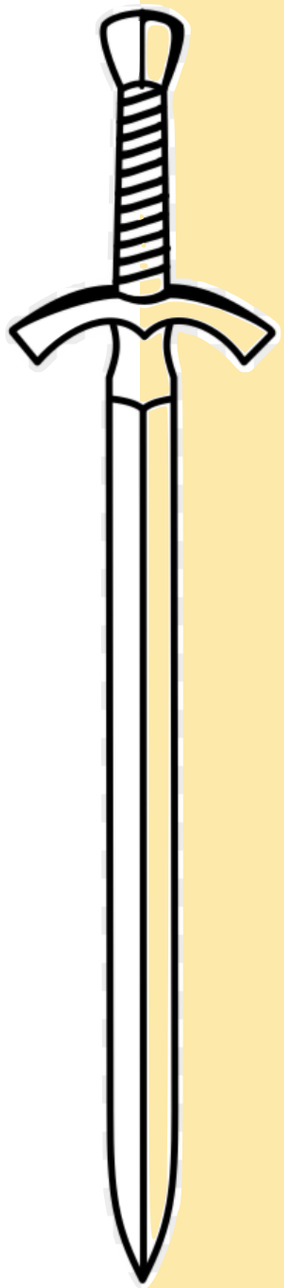
Walmart teams up with Google's Gemini to make it easier for shoppers to find and buy products

PUBLISHED SUN, JAN 11 2026·10:00 AM EST | UPDATED SUN, JAN 11 2026·10:42 AM EST

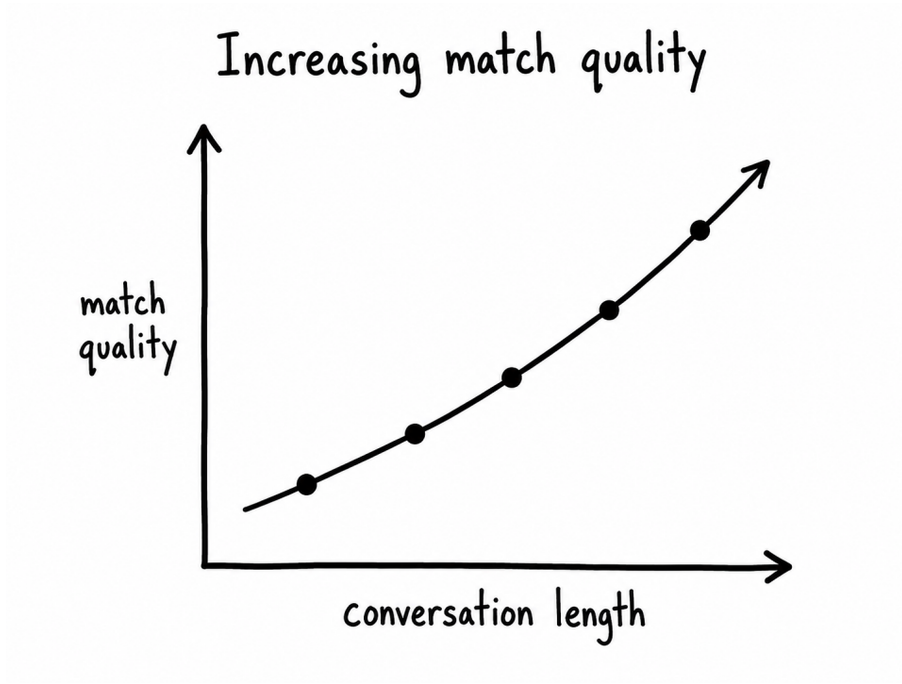
Melissa Repko
@IN/MELISSA-REPKO
@MELISSA_REPKO

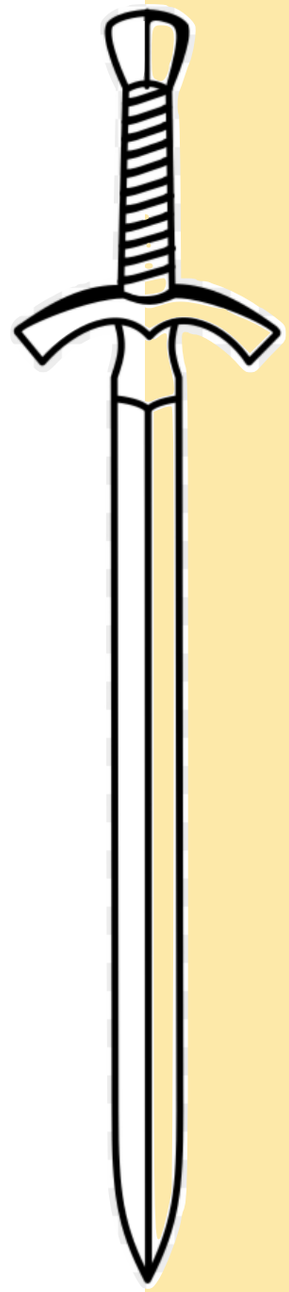
SHARE

Conversational Elicitation is a double-edged sword



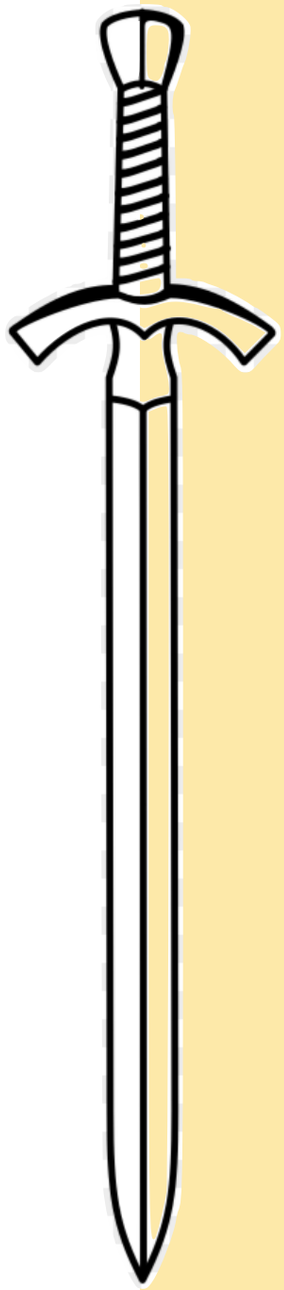
Conversational Elicitation is a double-edged sword





Conversational Elicitation is a double-edged sword

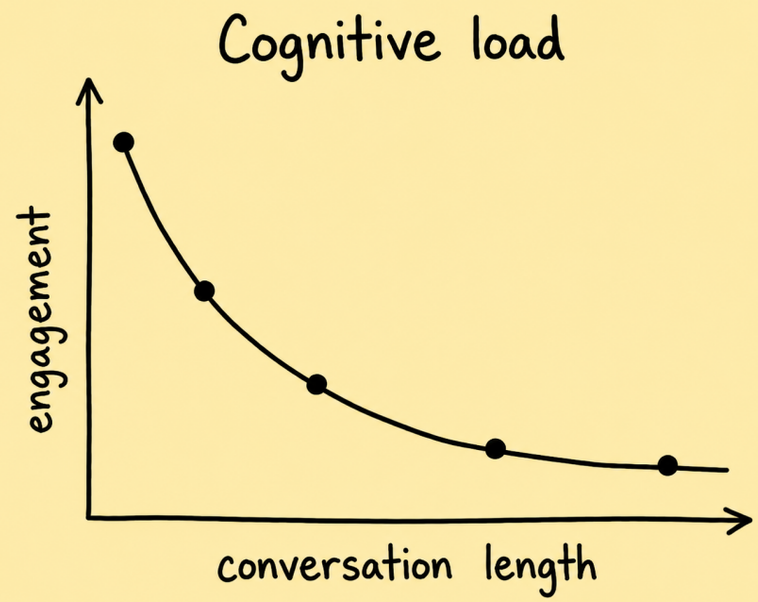




Conversational Elicitation is a double-edged sword







Conversational Elicitation is a double-edged sword





Different Elicitation Strategies in Practice



Different Elicitation Strategies in Practice



Quick Recommendations

Different Elicitation Strategies in Practice



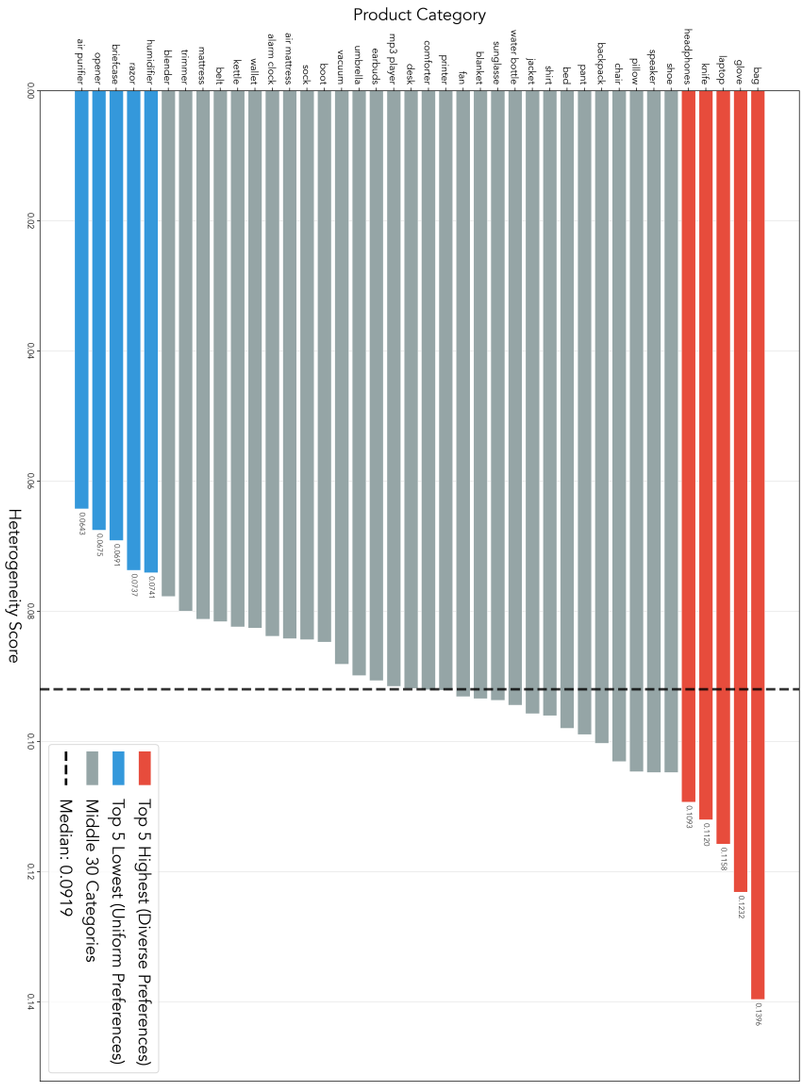
Quick Recommendations



Preference Elicitation

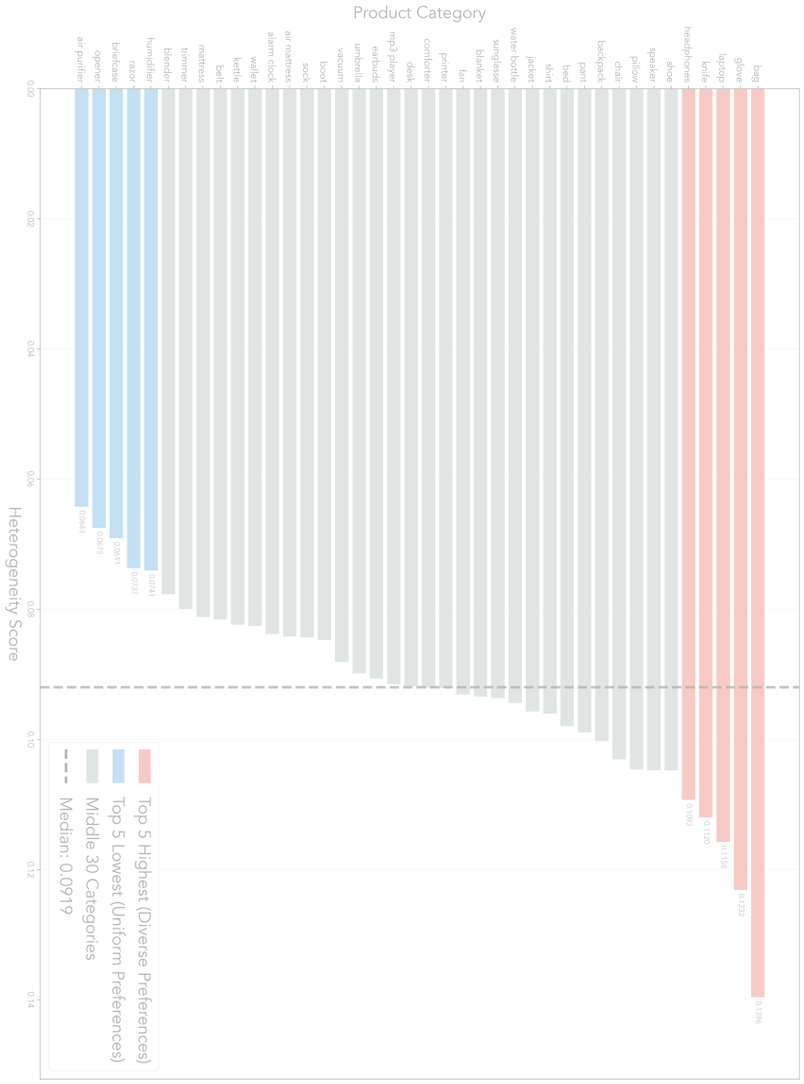
What forces drive differing strategies?

What forces drive differing strategies?



Different products have different preference heterogeneity

What forces drive differing strategies?



Different platforms have different monetization models



MEDIA

Netflix's advertising strategy shift is starting to pay off

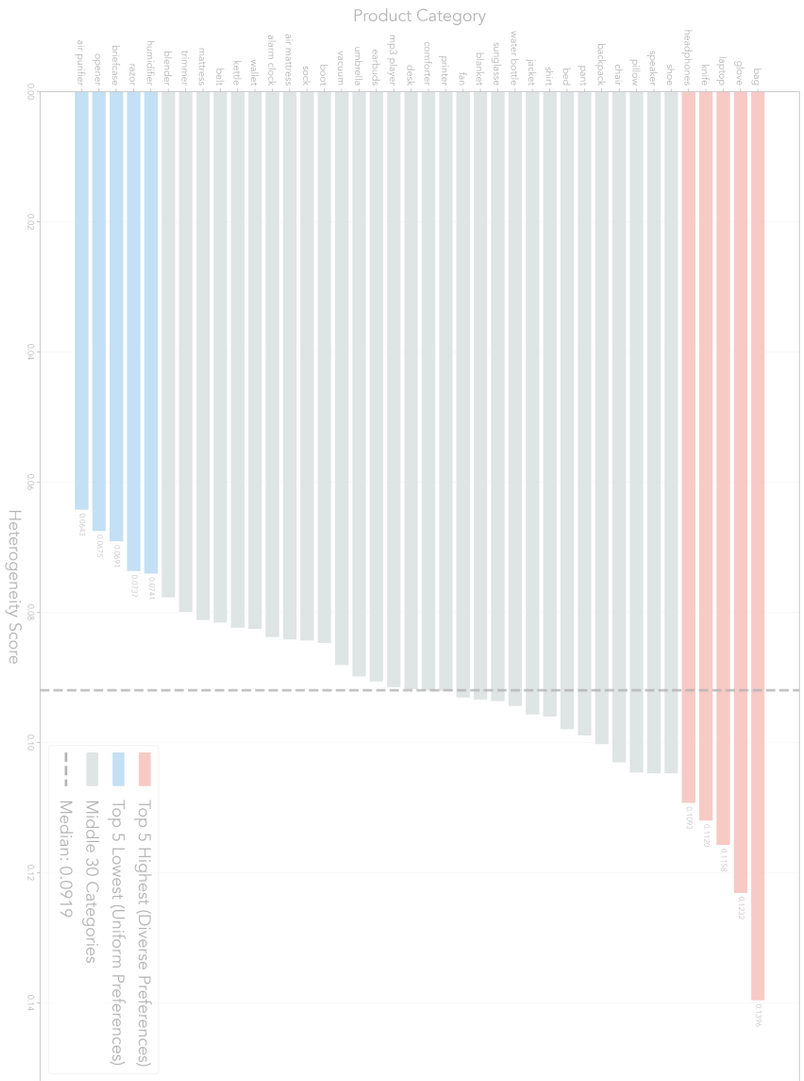
PUBLISHED WED, JAN 21 2026•1:39 PM EST | UPDATED WED, JAN 21 2026•4:43 PM EST



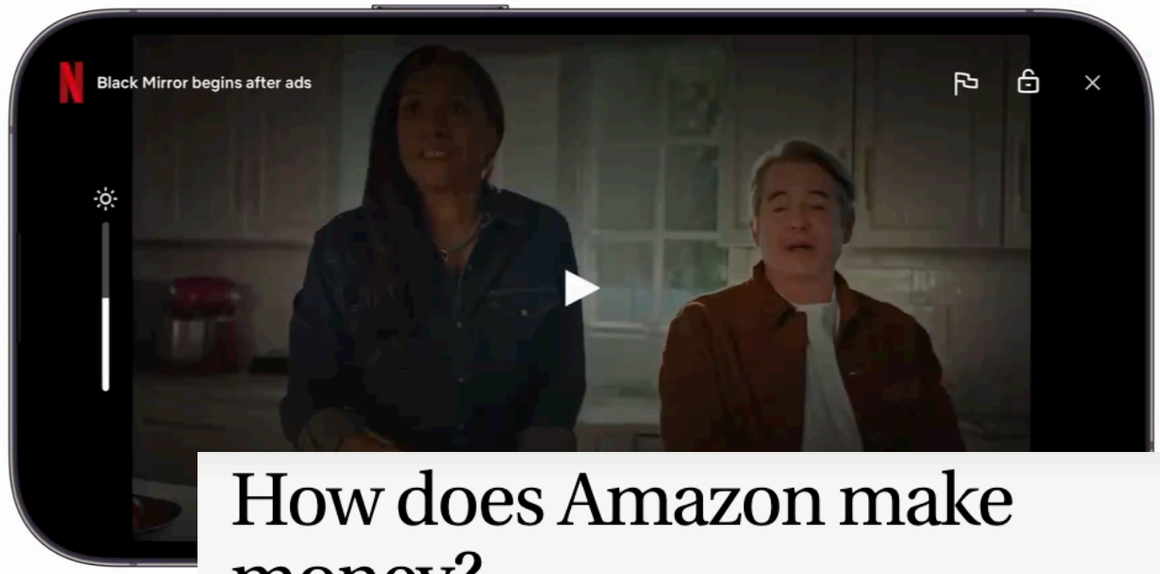
Lillian Rizzo
@LILLIANNNN

SHARE

What forces drive differing strategies?



Different platforms have different monetization models



MEDIA

Netflix's starting

PUBLISHED WED, JAN 21 202

Lillian Rizzo
@LILLIANNNN

How does Amazon make money?

Amazon has built a remarkably effective revenue model from which platform founders can learn. The company earns money at multiple points in its ecosystem across retail, infrastructure, digital media, and services. Let's examine each major revenue stream.

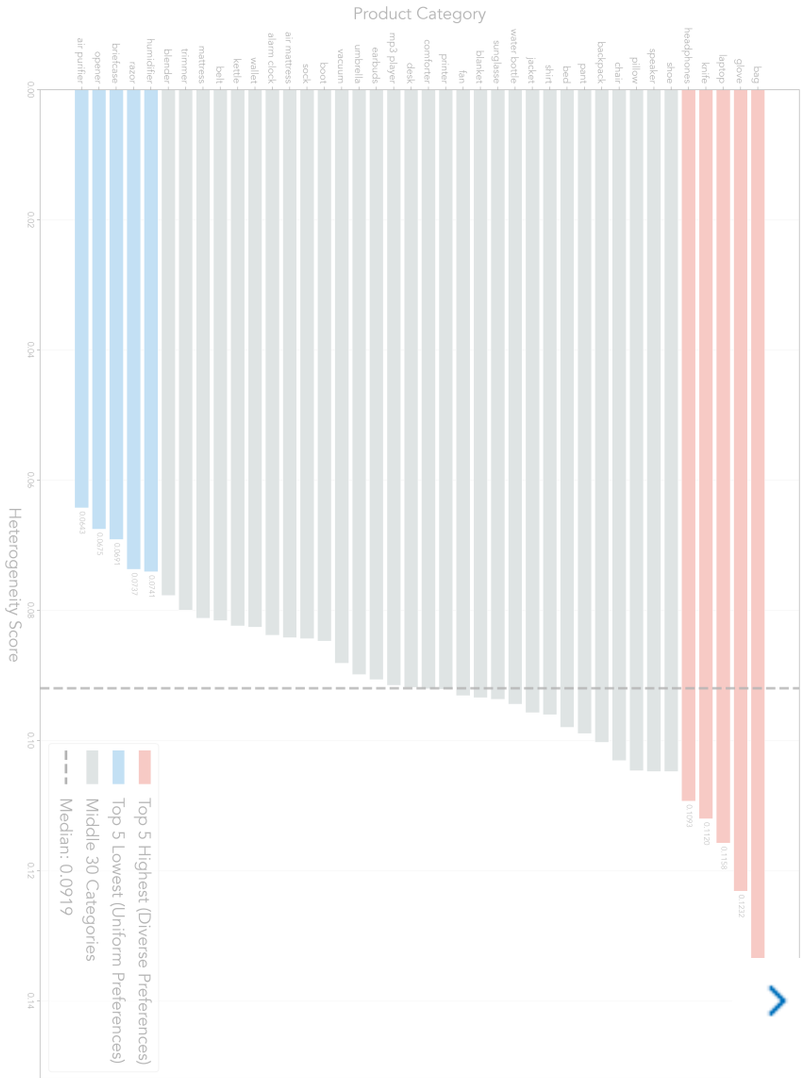
Third-party seller services

Amazon's marketplace is powered by third-party sellers, or individuals and businesses who list and sell their products on the platform. In fact, over 60% of all physical units sold on Amazon now come from these sellers.

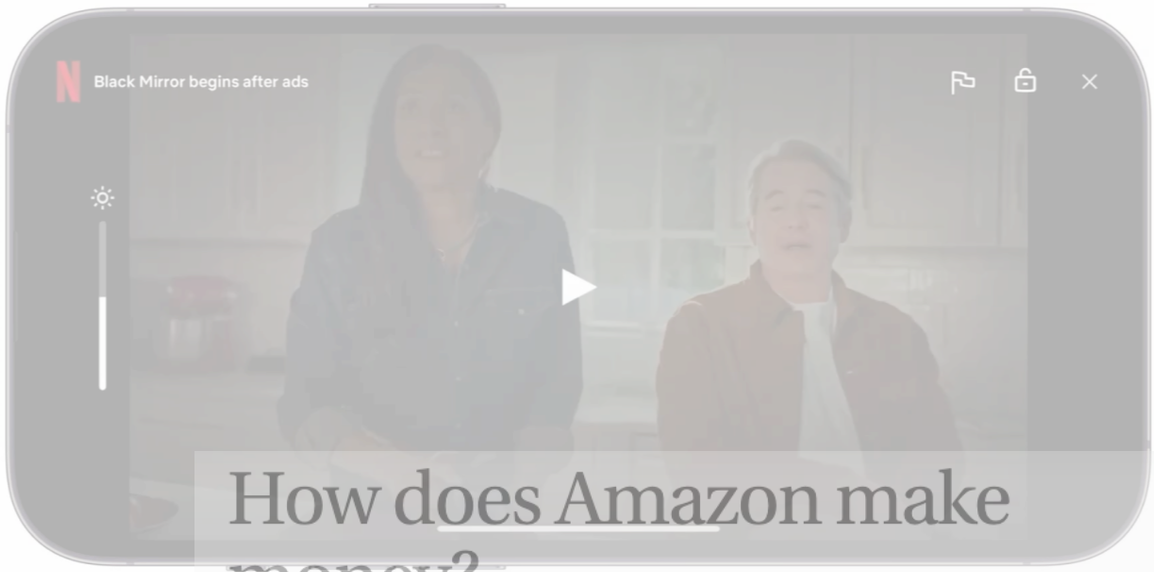
Amazon earns revenue from these sales by collecting referral fees (a.k.a. commission) on each sale. These fees vary by category and can range from 6% to 45%. Many sellers also use FBA, which means they pay Amazon to store, pack, and ship their orders. Add in sponsored product ads, seller subscriptions, and other upsells, and you've got a highly profitable, multi-layered revenue engine.

This is the core of Amazon's marketplace model: create the infrastructure and visibility sellers need, then monetize the tools that help them succeed.

What forces drive differing strategies?



Cognitive cost of elicitation



How does Amazon make money?

MEDIA

Netflix's advertising strategy shift is starting to pay off

Amazon has built a remarkably effective revenue model from which platform founders can learn. The company earns money at multiple points in its ecosystem across retail, infrastructure, digital media, and services. Let's examine each major revenue stream.

PUBLISHED WED, JAN 21 2026·1:39 PM EST | UPDATED WED, JAN 21 2026·4:43 PM EST

Lillian Rizzo

SHARE

Third-party seller services

> J Exp Psychol Gen. 2010 Nov;139(4):665-82. doi: 10.1037/a0020198.

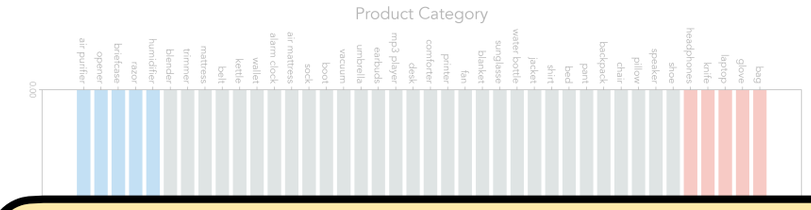
Decision making and the avoidance of cognitive demand

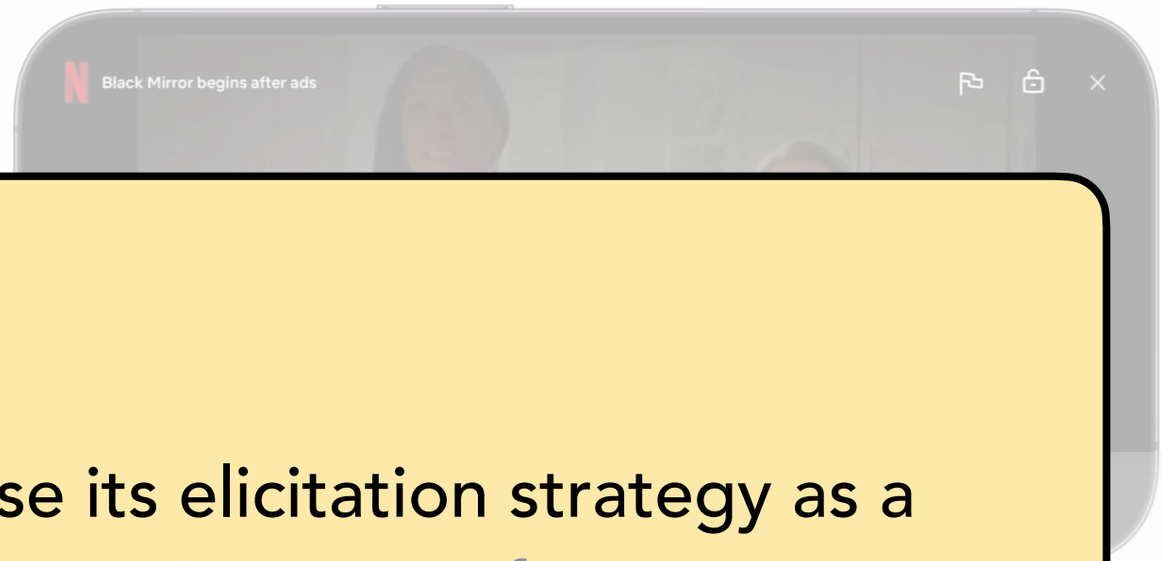
Wouter Kool [1], Joseph T McGuire, Zev B Rosen, Matthew M Botvinick

Affiliations + expand

PMID: 20853993 PMCID: PMC2970648 DOI: 10.1037/a0020198

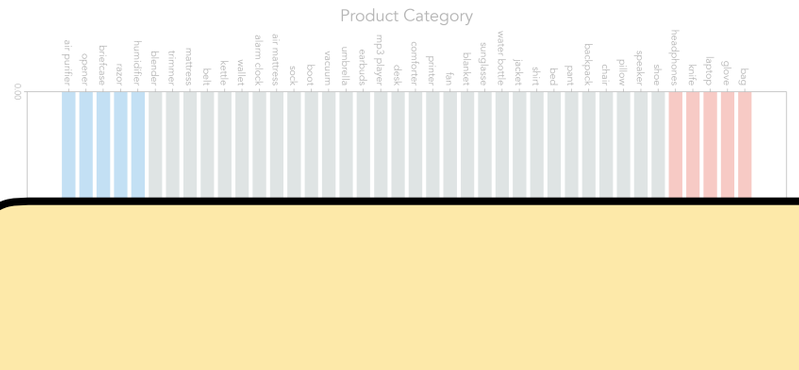
Research Question

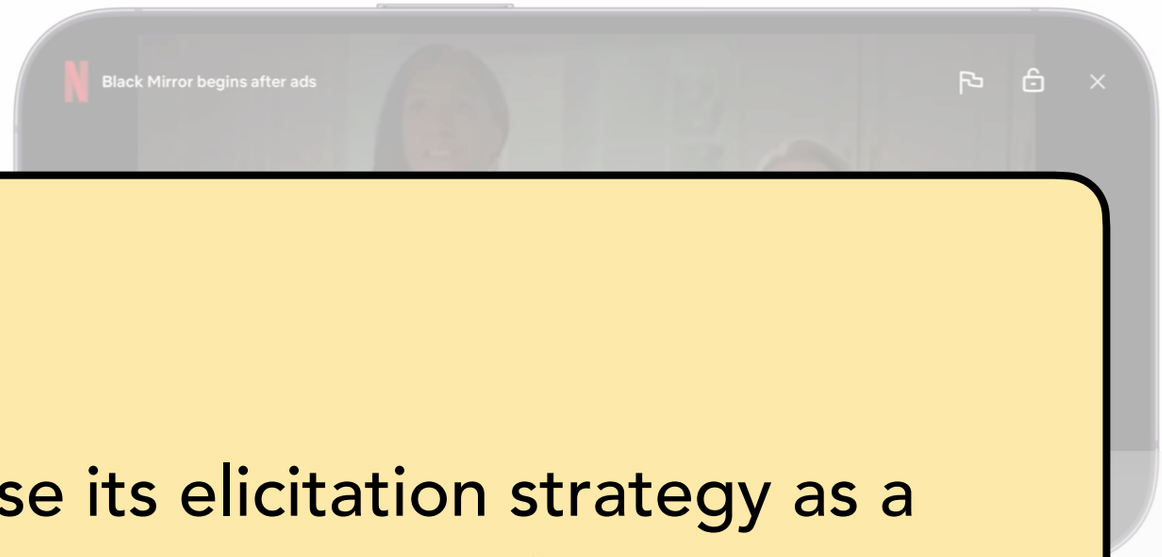




How should a platform choose its elicitation strategy as a function of market characteristics—preference heterogeneity, match value, and abandonment due to costly communication—and how do optimal strategies change with differing platform objectives?

Research Question





How should a platform choose its elicitation strategy as a function of **market characteristics**—preference heterogeneity, match value, and abandonment due to costly communication—and how do optimal strategies change with **differing platform objectives**?

In A Nutshell

- A parsimonious model of conversational elicitation
- LLM-based testbed for the value of elicitation



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
- LLM-based testbed for the value of elicitation



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
- LLM-based testbed for the value of elicitation



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
    - Conversion → quick recos, minimal probing (elicitation rarely pays)
- LLM-based testbed for the value of elicitation



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
    - Conversion → quick recos, minimal probing (elicitation rarely pays)
    - Commission → deeper elicitation is optimal can hurt consumer welfare
- LLM-based testbed for the value of elicitation



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
    - Conversion → quick recos, minimal probing (elicitation rarely pays)
    - Commission → deeper elicitation is optimal can hurt consumer welfare
- LLM-based testbed for the value of elicitation
  - Real long-form queries → personas that reveal preferences turn by turn



In A Nutshell

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
    - Conversion → quick recos, minimal probing (elicitation rarely pays)
    - Commission → deeper elicitation is optimal can hurt consumer welfare
- LLM-based testbed for the value of elicitation
  - Real long-form queries → personas that reveal preferences turn by turn
  - Measures incremental conversion and revenue as information accrues

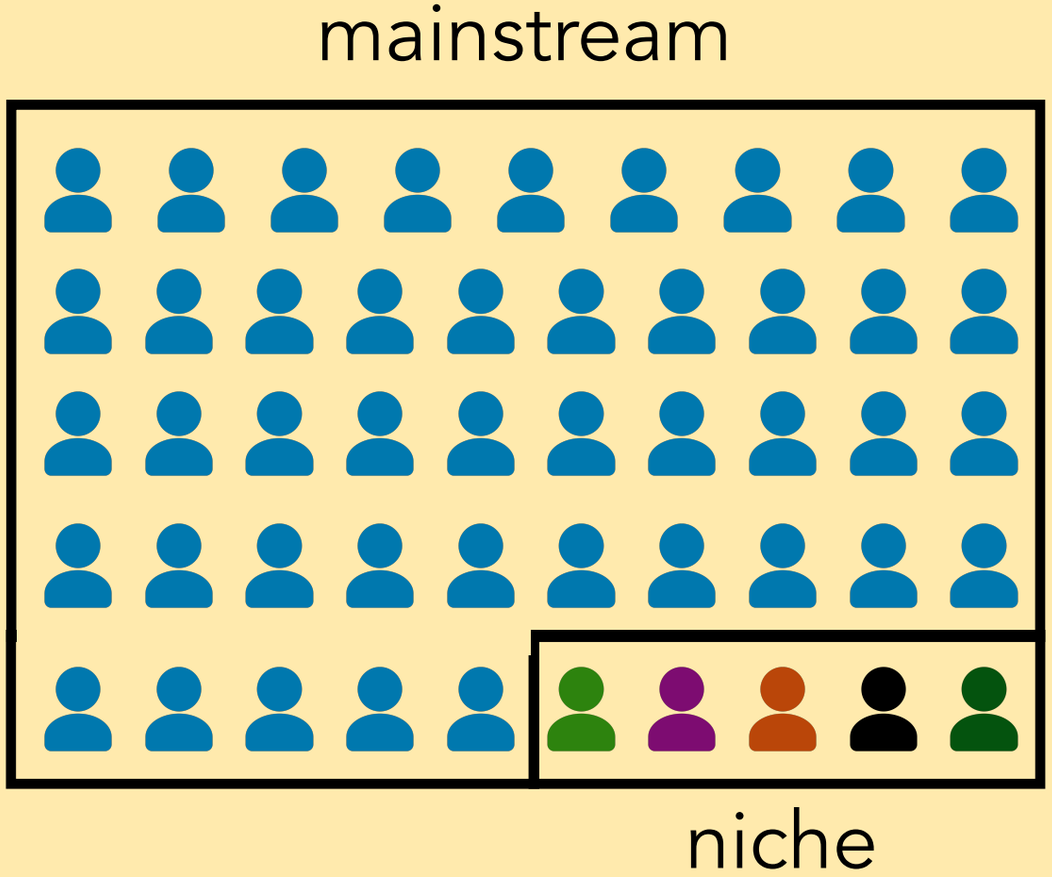


(Parsimonious) Model of conversational elicitation

Model

User Types

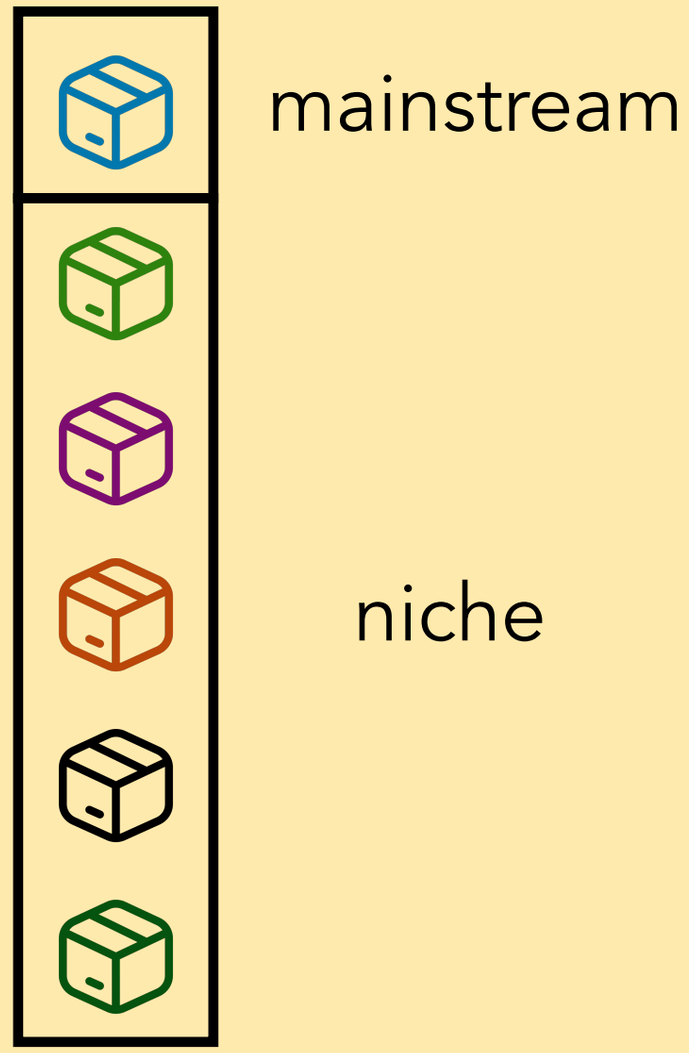
- There are $n + 1$ types of users
  - 1 mainstream ($\theta = \emptyset$)
  - n niches ($\theta = j$), $j = 1,2,\ldots,n$
- α : fraction of niche users



Model

Product Types

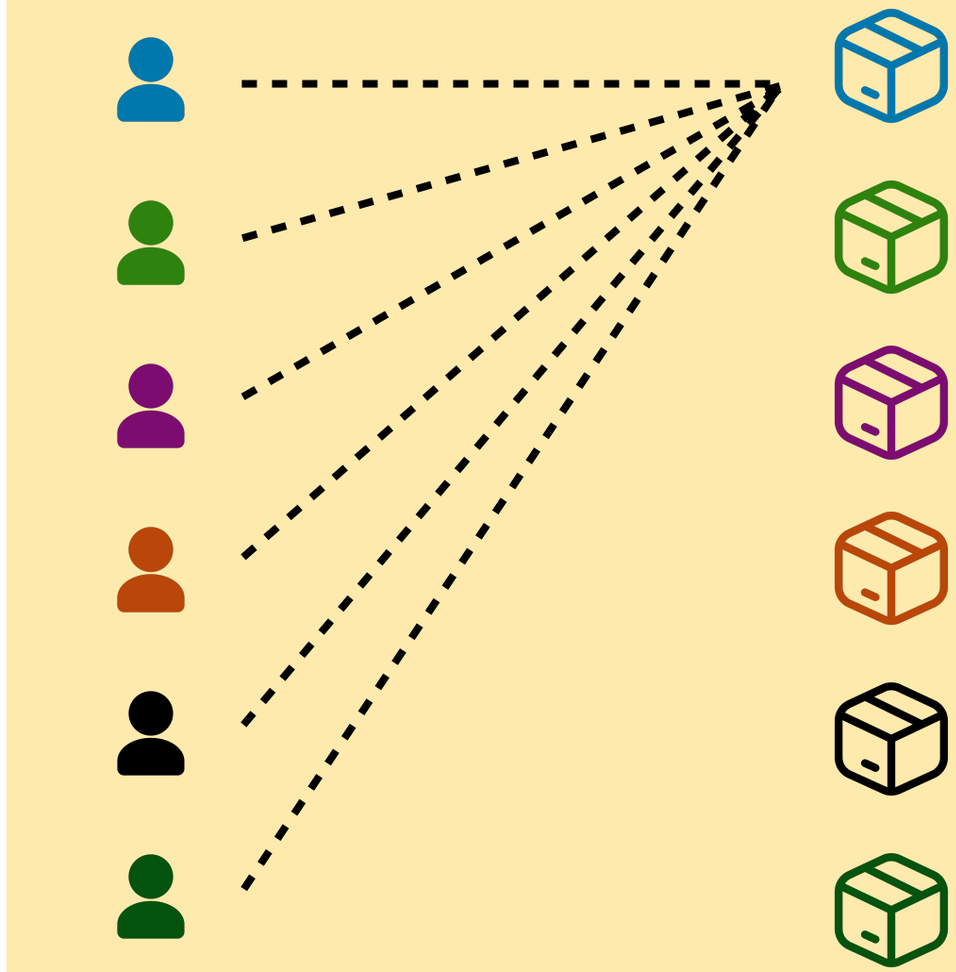
- There are $n + 1$ types of product
  - 1 mainstream (P)
  - n niche (j), $j = 1,2,\ldots,n$



Model

Match Value

$$u(\theta, P) = v_P > 0, \quad \forall \theta \in \Theta$$



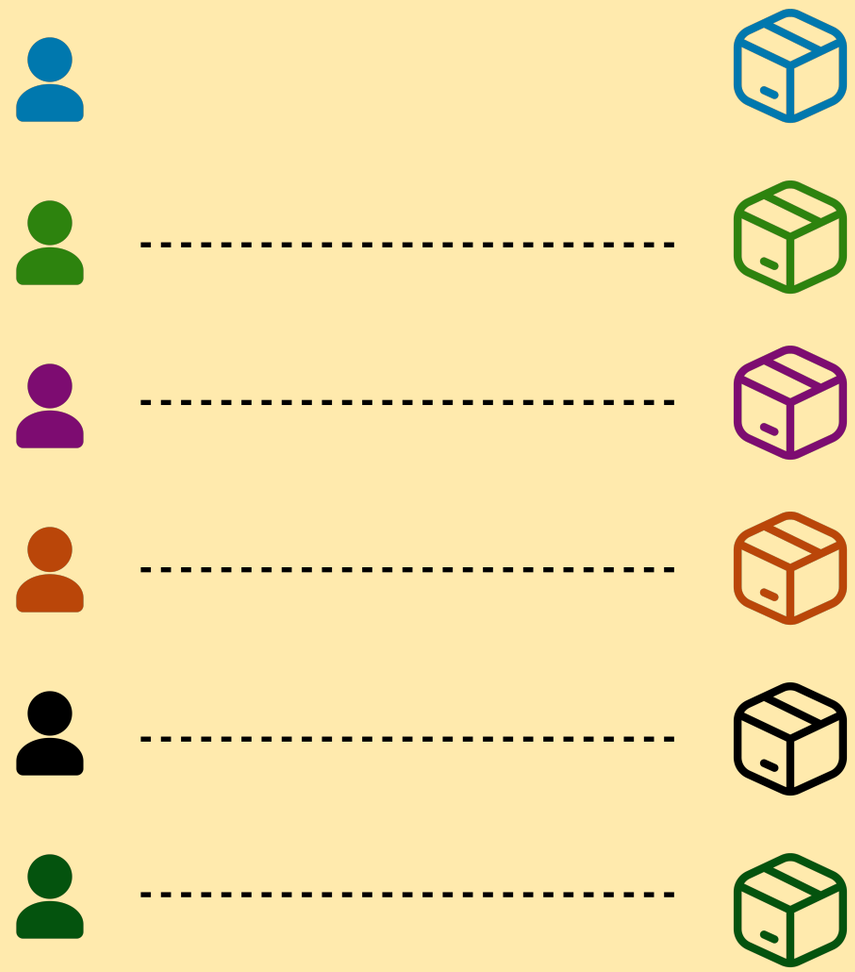
Model

Match Value

$$u(\theta, P) = v_P > 0, \quad \forall \theta \in \Theta$$

$$u(j, j) = v_N \alpha^{-\beta}, \quad \forall j \in [n]$$

$v_N > 0,\ \beta \geq 0$



Model

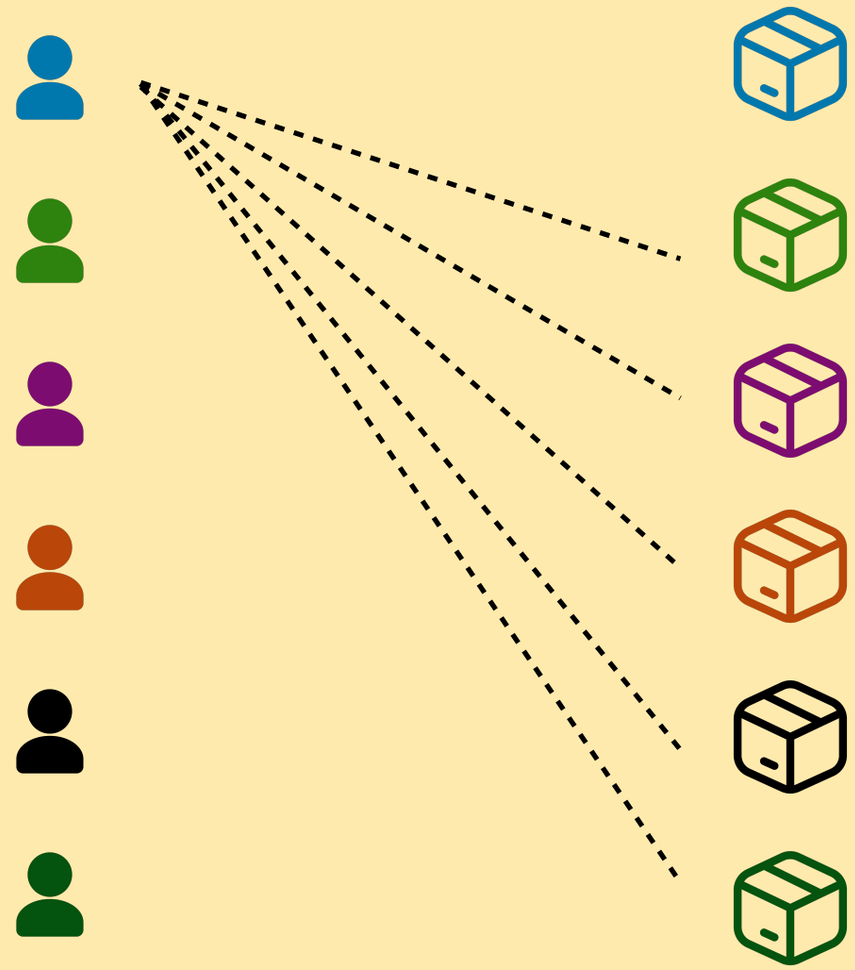
Match Value

$$u(\theta, P) = v_P > 0, \quad \forall \theta \in \Theta$$

$$u(j, j) = v_N \alpha^{-\beta}, \quad \forall j \in [n]$$

$$v_N > 0,\ \beta \geq 0$$

$$u(\theta, j) = -1\{\theta \neq j\}, \quad \forall \theta \in \Theta, j \in [n]$$



Model

User Θ

Rec. System

Model

User Θ

Rec. System

Decision:
elicitation intensity f, rec. policy π

(for commission, sellers price based on f, π)

Model

User Θ

Rec. System

Decision:
elicitation intensity f, rec. policy π

(for commission, sellers price based on f, π)

Noisy signal $S \mid \Theta$ with "accuracy" f,
but incurs cost/abandons.

Model

User Θ

Rec. System

Decision:
elicitation intensity f, rec. policy π

(for commission, sellers price based on f, π)

Noisy signal $S \mid \Theta$ with "accuracy" f,
but incurs cost/abandons.

Recommends $\pi(s)$

Model

User Θ

Rec. System

Decision:
elicitation intensity f, rec. policy π

(for commission, sellers price based on f, π)

Noisy signal $S \mid \Theta$ with "accuracy" f,
but incurs cost/abandons.

Recommends $\pi(s)$

Adopts $\pi(s)$ or outside option (logit choice model)

Model

User Θ

Rec. System

Decision:
elicitation intensity f, rec. policy π

(for commission, sellers price based on f, π)

Noisy signal $S \mid \Theta$ with "accuracy" f,
but incurs cost/abandons.

Recommends $\pi(s)$

Adopts $\pi(s)$ or outside option (logit choice model)

Observes conversion/commission outcome

Model

Platform Objectives

- Conversion
  - Platform earns some constant payoff whenever the user adopts platform recommendation
  - e.g. Netflix earns ad revenue as long as you watch ***something***

constant payoff

probability user adopts recommendation

$$\Gamma(f, \pi) = R \cdot s(f) \cdot \mathbb{E}[\sigma(u(\theta, \pi))]$$

probability user doesn't abandon

Model

Platform Objectives

- Commission
  - Platform earns a commission based on the price of the product that is chosen
  - e.g. Amazon earns commission on the price you pay for chosen item

demand for product j with price p_j

probability user doesn't abandon

$$D_j(f, \pi, p) = s(f) \cdot \mathbb{E}[1\{\pi = j\} \cdot \sigma(u(\theta, j) - p_j)]$$

probability user adopts recommendation

seller j chooses price p_j to max rev.

$$\Pi_j(p_j; f, \pi) = (1 - \tau) \cdot p_j \cdot D_j(f, \pi, p)$$

commission rate; equilibrium demand

$$\Pi^{\text{com}}(f, \pi) = \tau \cdot \sum_j p_j^\star D_j(f, \pi, p^\star)$$

equilibrium price

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Intuition

Purchase probability is quickly saturated, so the gains to elicitation are bounded

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Intuition

Purchase probability is quickly saturated, so the gains to elicitation are bounded

Implication for user welfare

User welfare optimization requires deep elicitation which implies misalignment between engagement & welfare optimization

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Intuition

Purchase probability is quickly saturated, so the gains to elicitation are bounded

Implication for user welfare

User welfare optimization requires deep elicitation which implies misalignment between engagement & welfare optimization

Commission

Theorem 2 (Informal)

For α sufficiently small, it is commission optimal to have $f^\star \approx 1$ **(deep elicitation)**

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Intuition

Purchase probability is quickly saturated, so the gains to elicitation are bounded

Implication for user welfare

User welfare optimization requires deep elicitation which implies misalignment between engagement & welfare optimization

Commission

Theorem 2 (Informal)

For α sufficiently small, it is commission optimal to have $f^\star \approx 1$ **(deep elicitation)**

Intuition

Conversation improves screening. More niche customers are correctly matched which implies higher equilibrium prices

Analytical Result: stark differences in how much we ask

our analytical results focus on the small niche regime, $\alpha \to 0$

Engagement

Theorem 1 (Informal)

For α sufficiently small, it is engagement optimal is recommend P **(minimal probing)**

Intuition

Purchase probability is quickly saturated, so the gains to elicitation are bounded

Implication for user welfare

User welfare optimization requires deep elicitation which implies misalignment between engagement & welfare optimization

Commission

Theorem 2 (Informal)

For α sufficiently small, it is commission optimal to have $f^\star \approx 1$ **(deep elicitation)**

Intuition

Conversation improves screening. More niche customers are correctly matched which implies higher equilibrium prices

Implication for user welfare

Induced user welfare effect can be positive or negative depending on parameters.
(Paradox of Personalization)

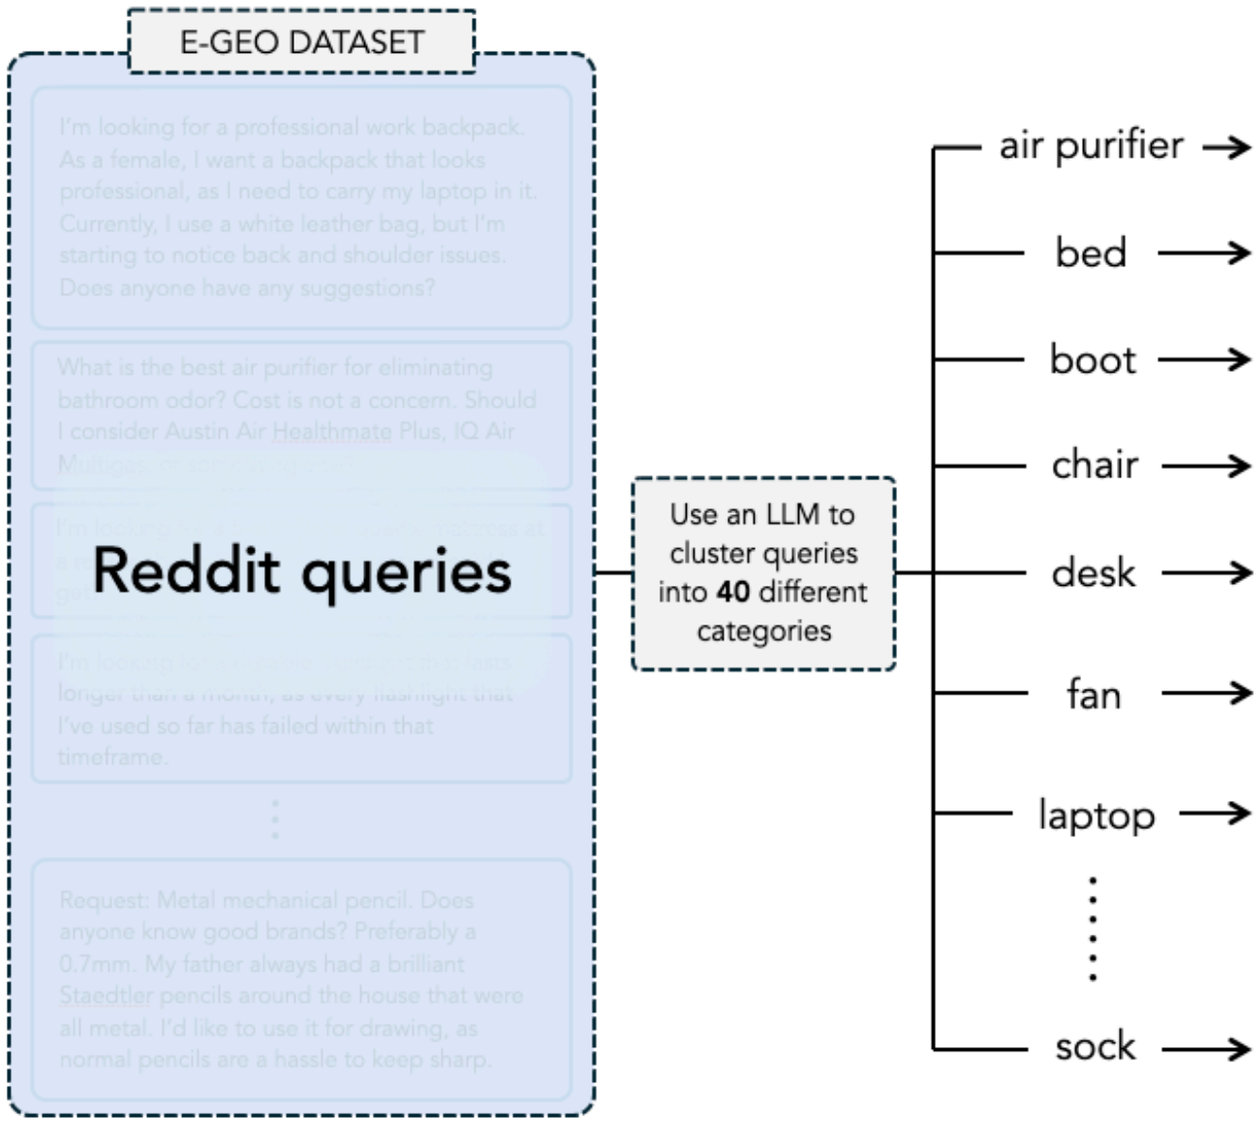
LLM-based testbed for the value of elicitation

Grounding LLMs to create structured data



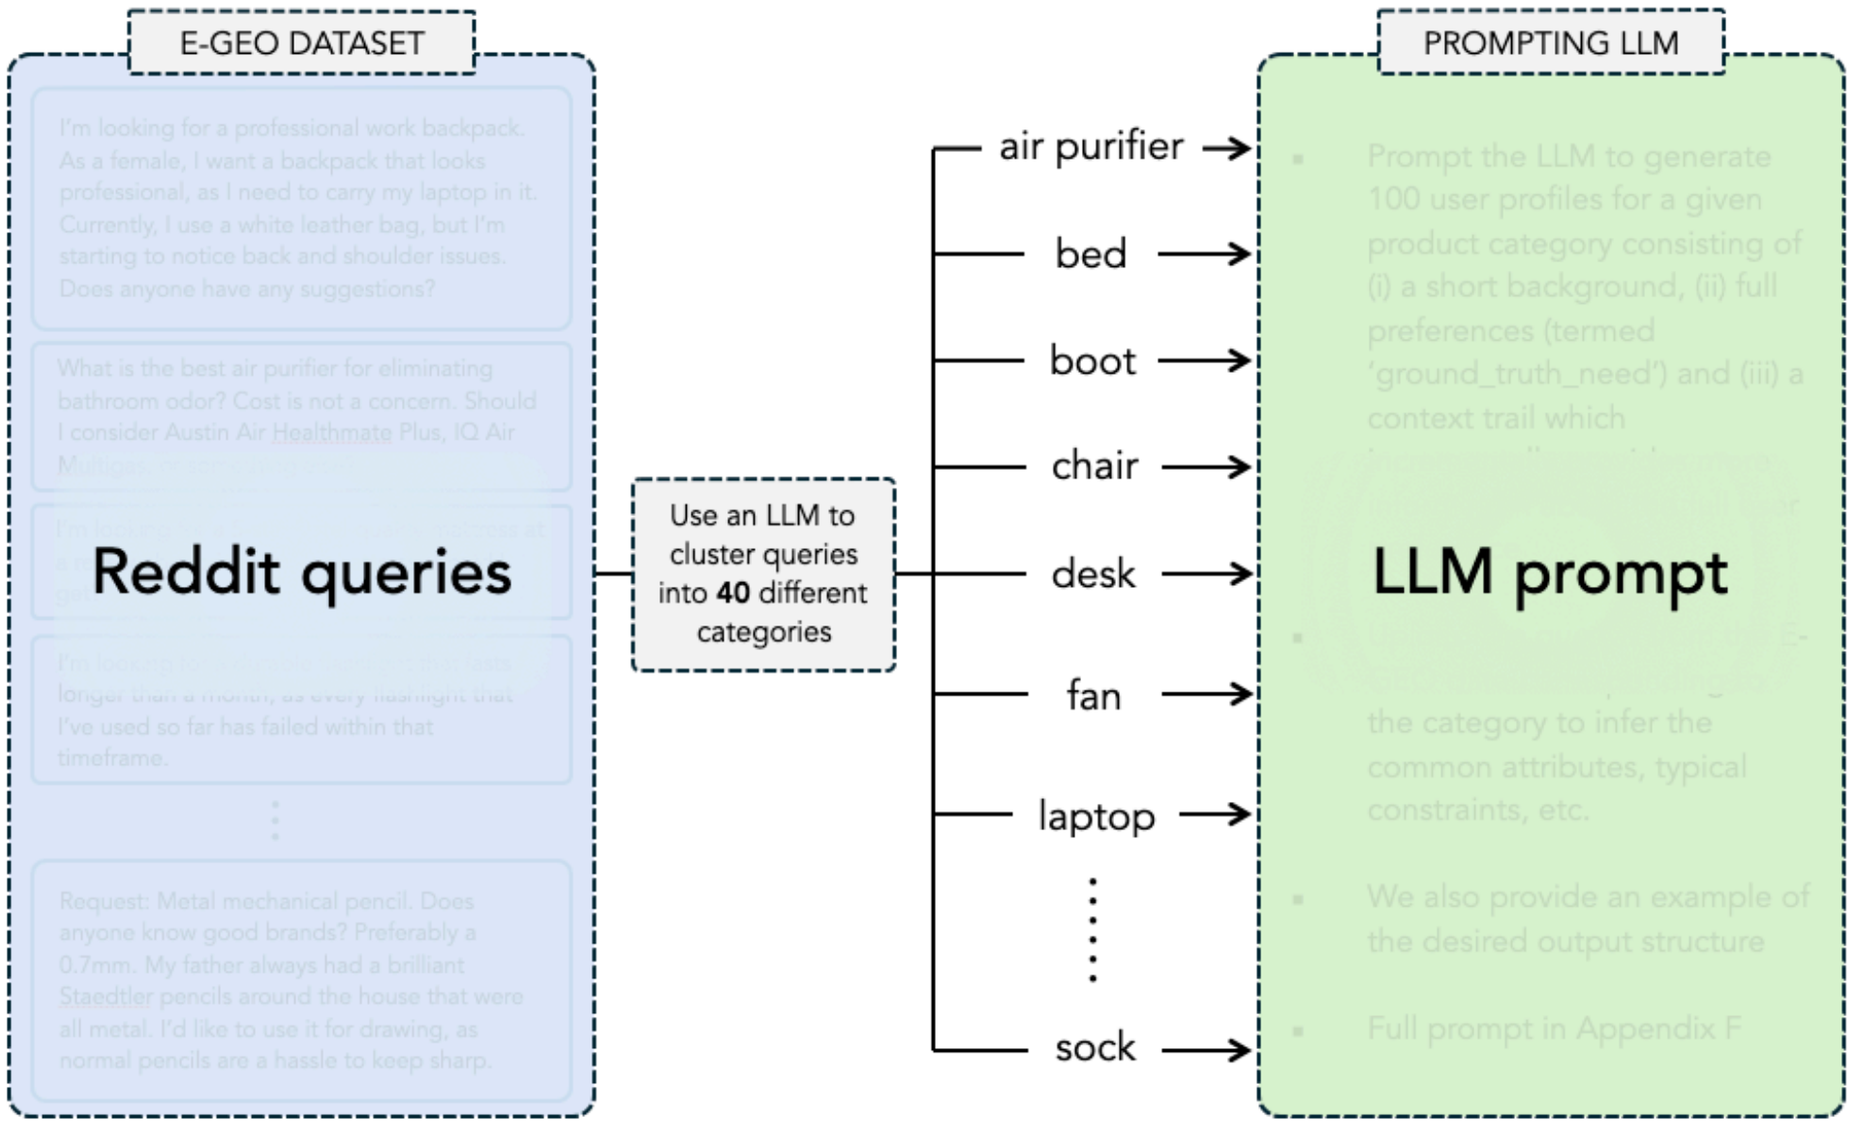
Bagga, Puneet S., Vivek F. Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. "E-GEO: A Testbed for Generative Engine Optimization in E-Commerce."

Grounding LLMs to create structured data



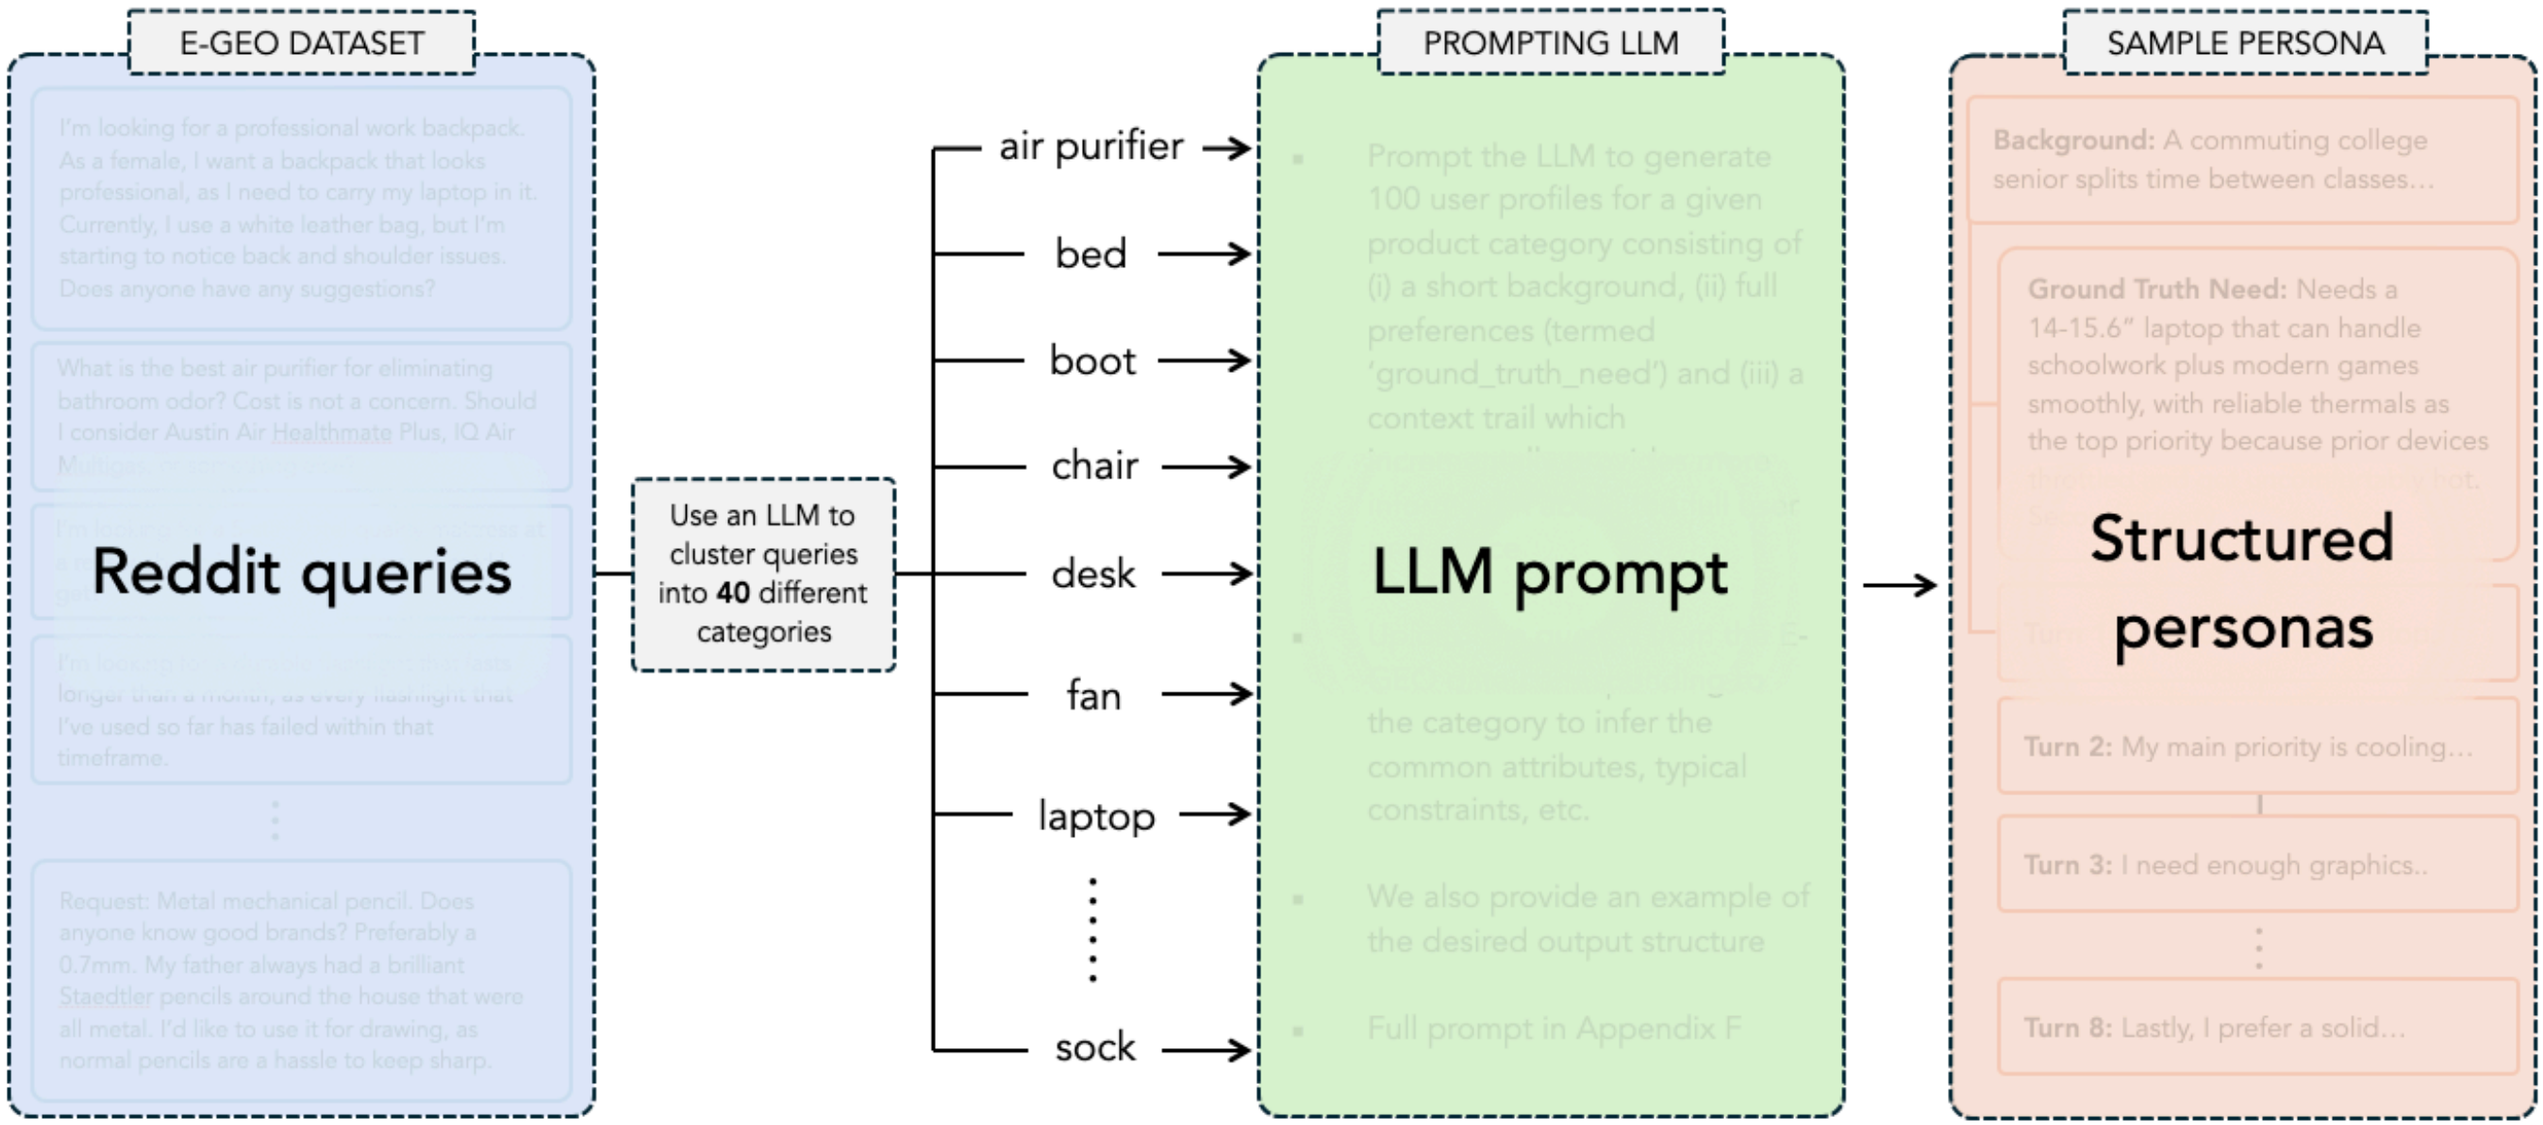
Bagga, Puneet S., Vivek F. Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. "E-GEO: A Testbed for Generative Engine Optimization in E-Commerce."

Grounding LLMs to create structured data



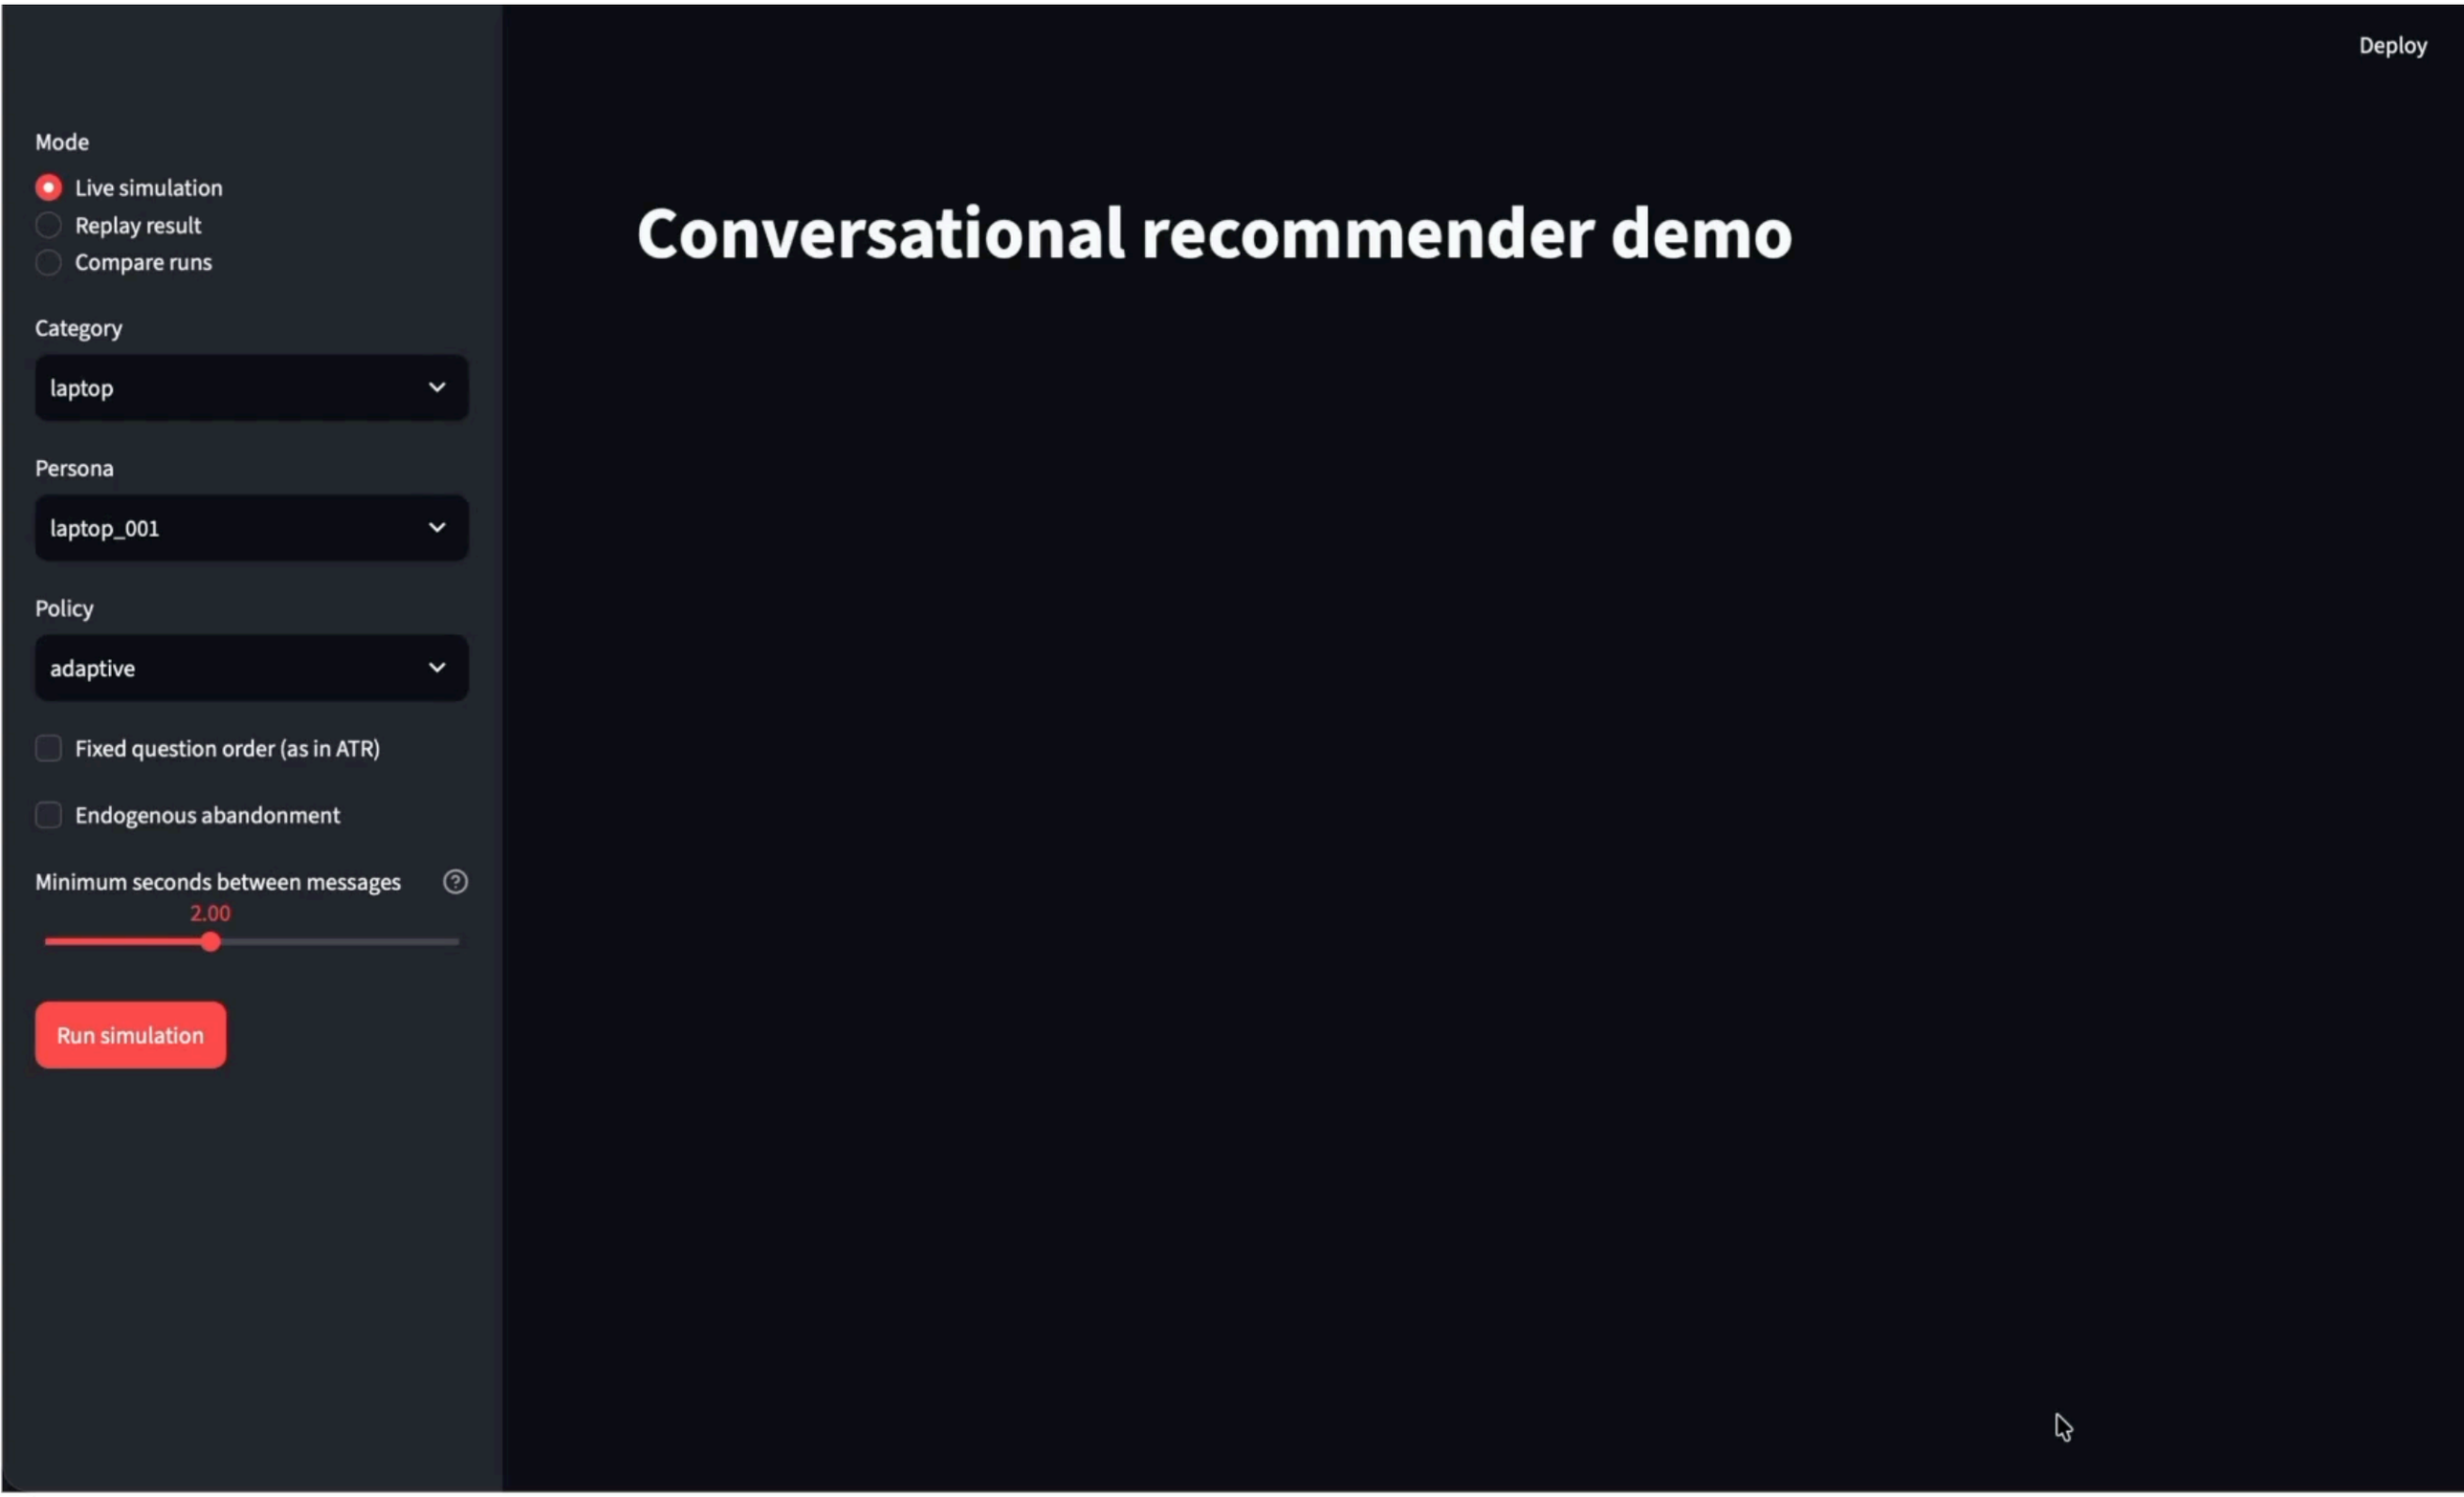
Bagga, Puneet S., Vivek F. Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. "E-GEO: A Testbed for Generative Engine Optimization in E-Commerce."

Grounding LLMs to create structured data

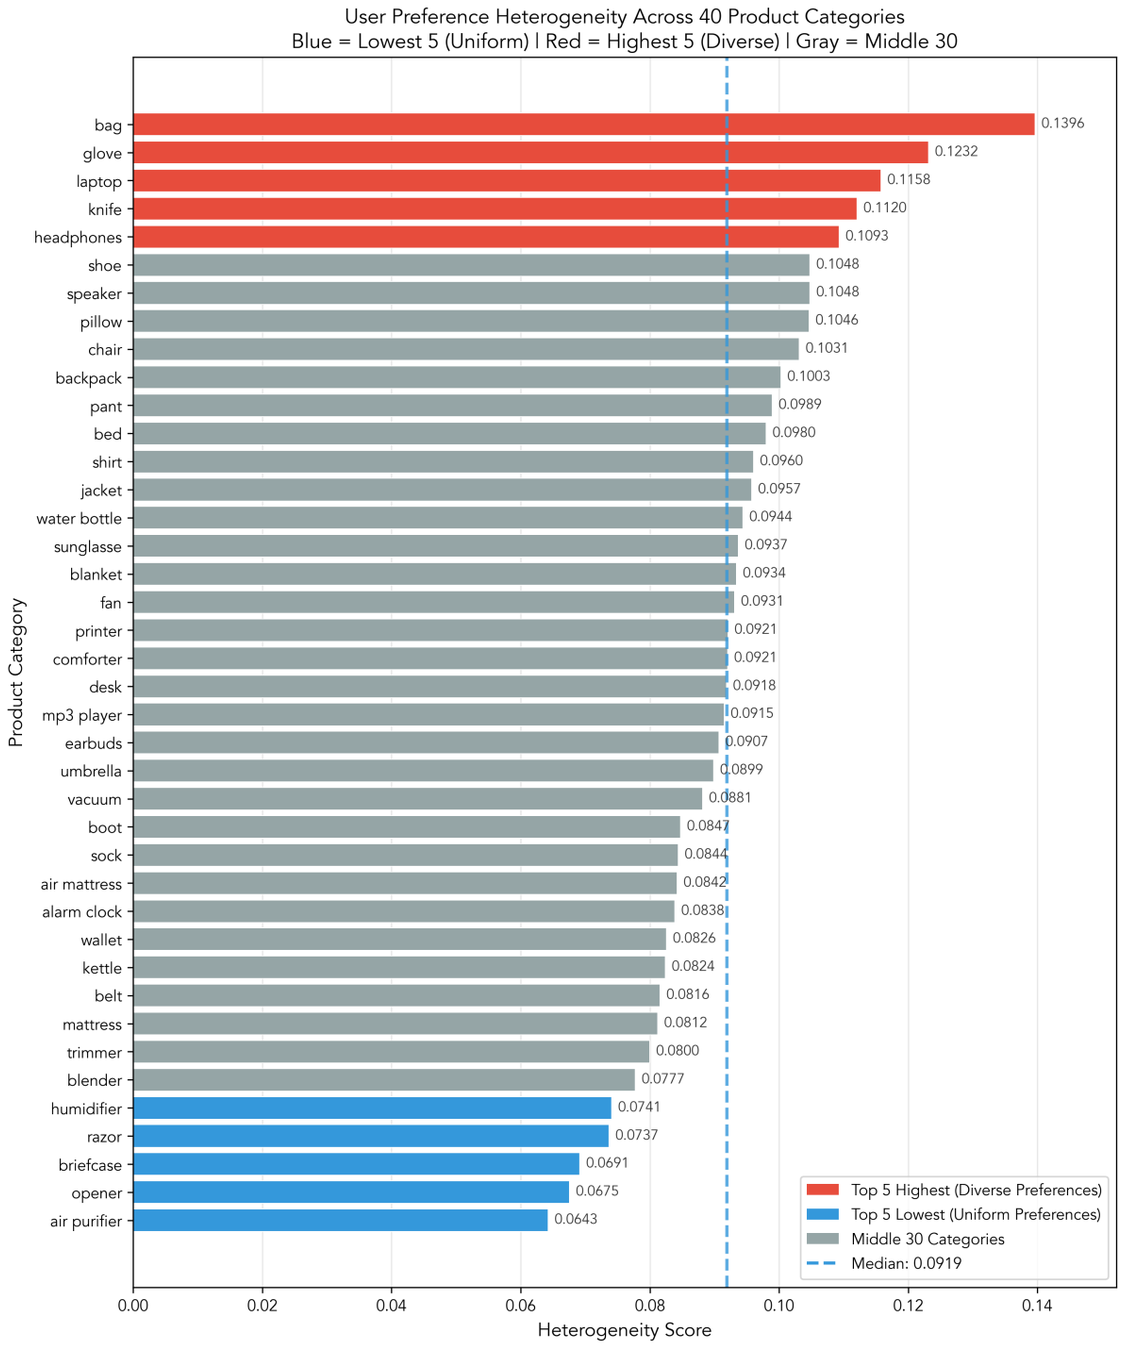


Bagga, Puneet S., Vivek F. Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu. "E-GEO: A Testbed for Generative Engine Optimization in E-Commerce."

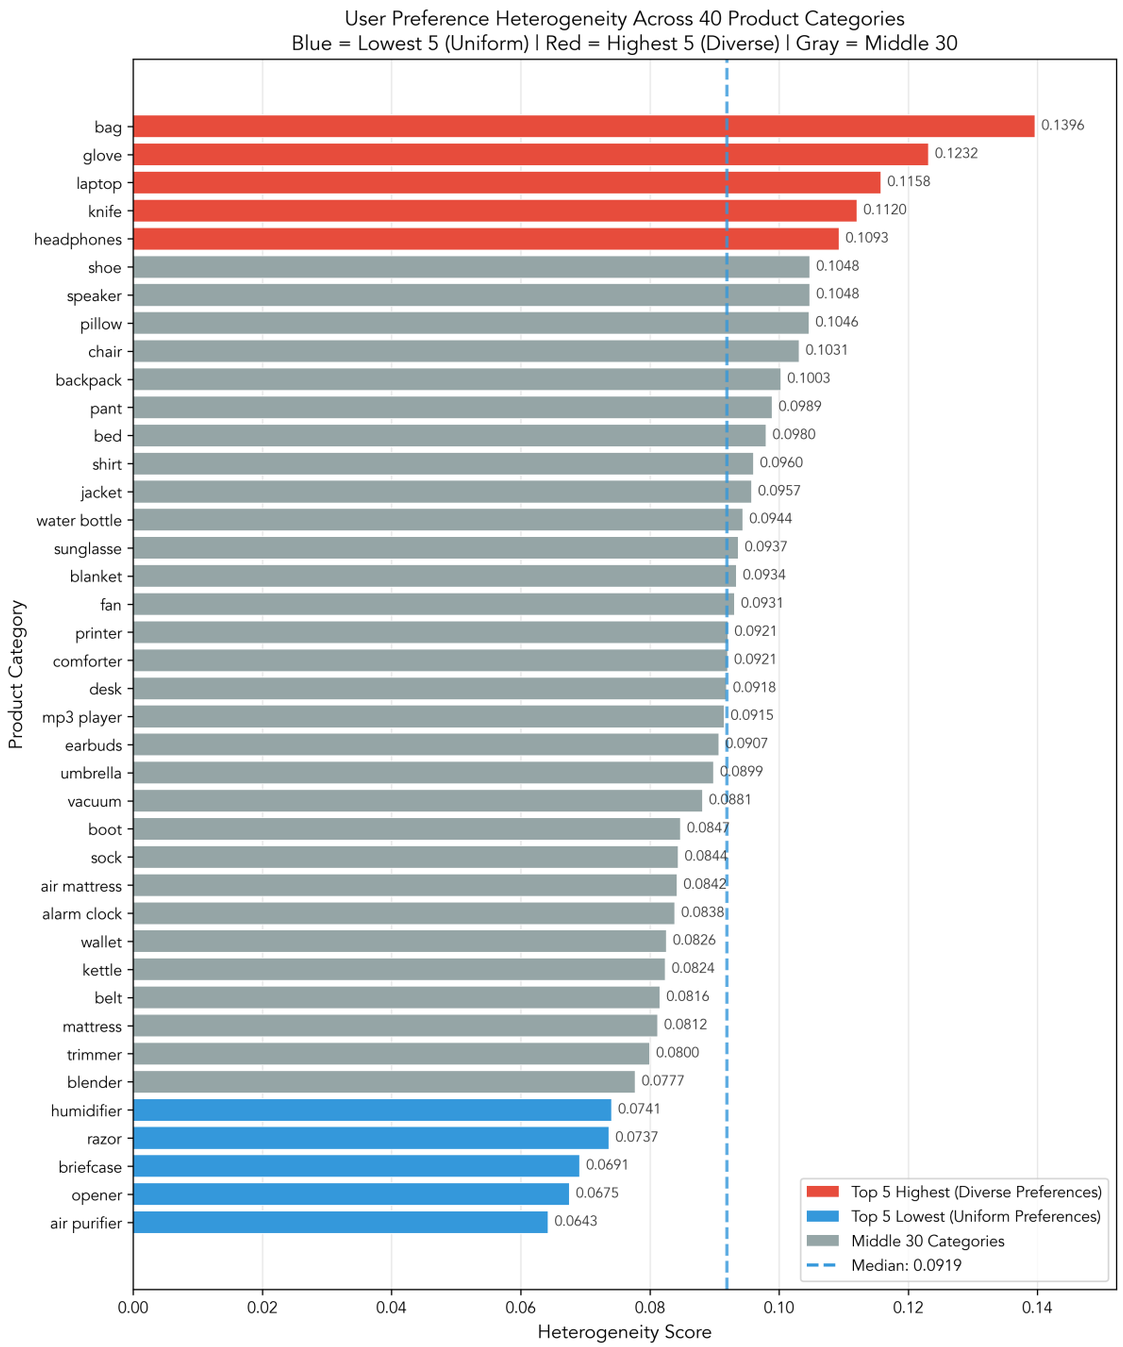
Simulating user decisions with LLMs



Preference Heterogeneity & its impact on WTP

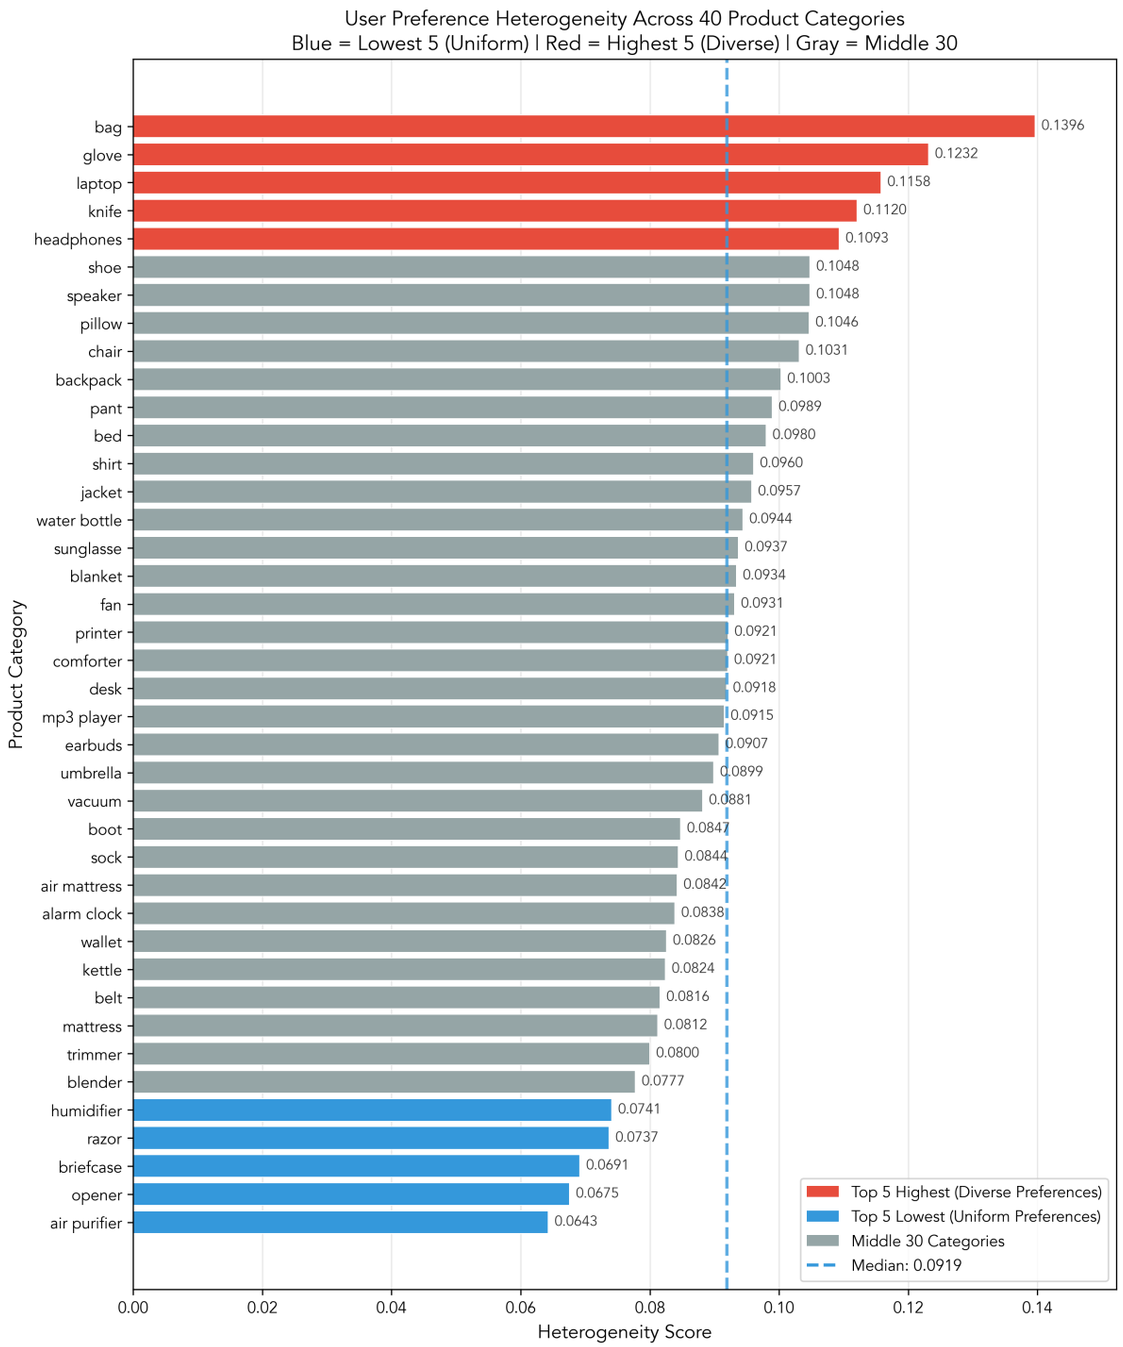


Preference Heterogeneity & its impact on WTP



We measure heterogeneity using semantic similarity in preferences

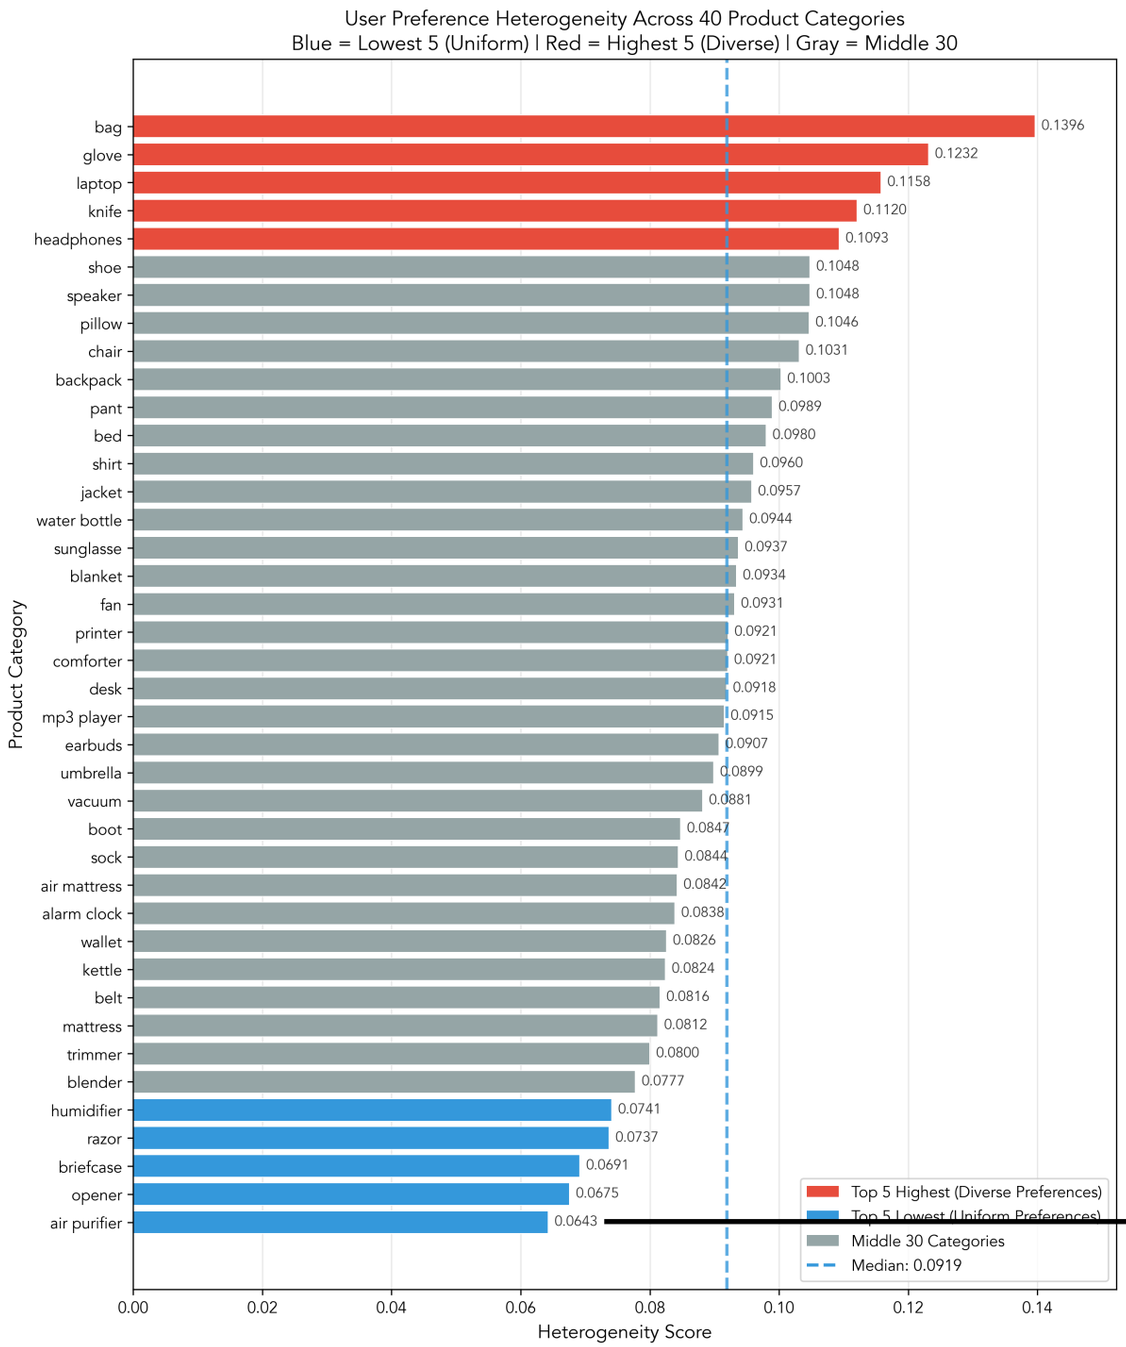
Preference Heterogeneity & its impact on WTP

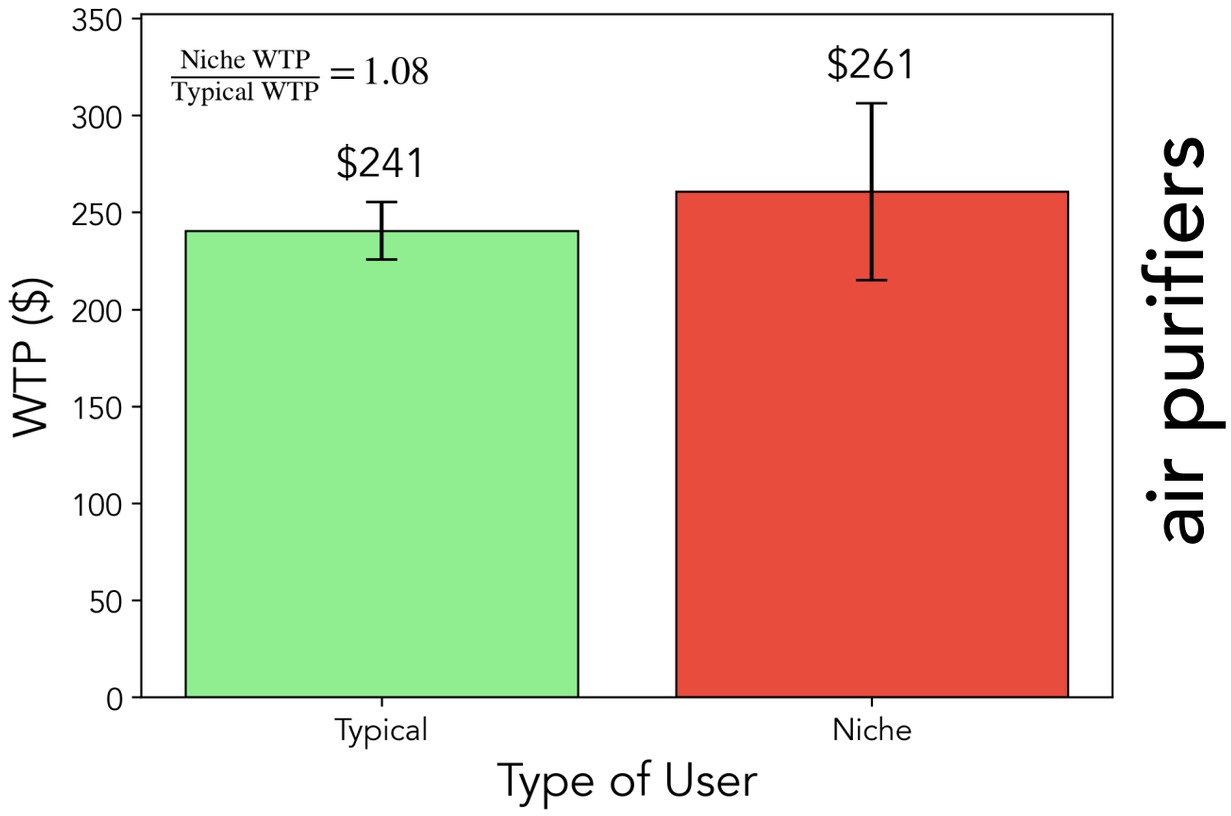


We measure heterogeneity using semantic similarity in preferences

Furthest 10% of users are niche

Preference Heterogeneity & its impact on WTP

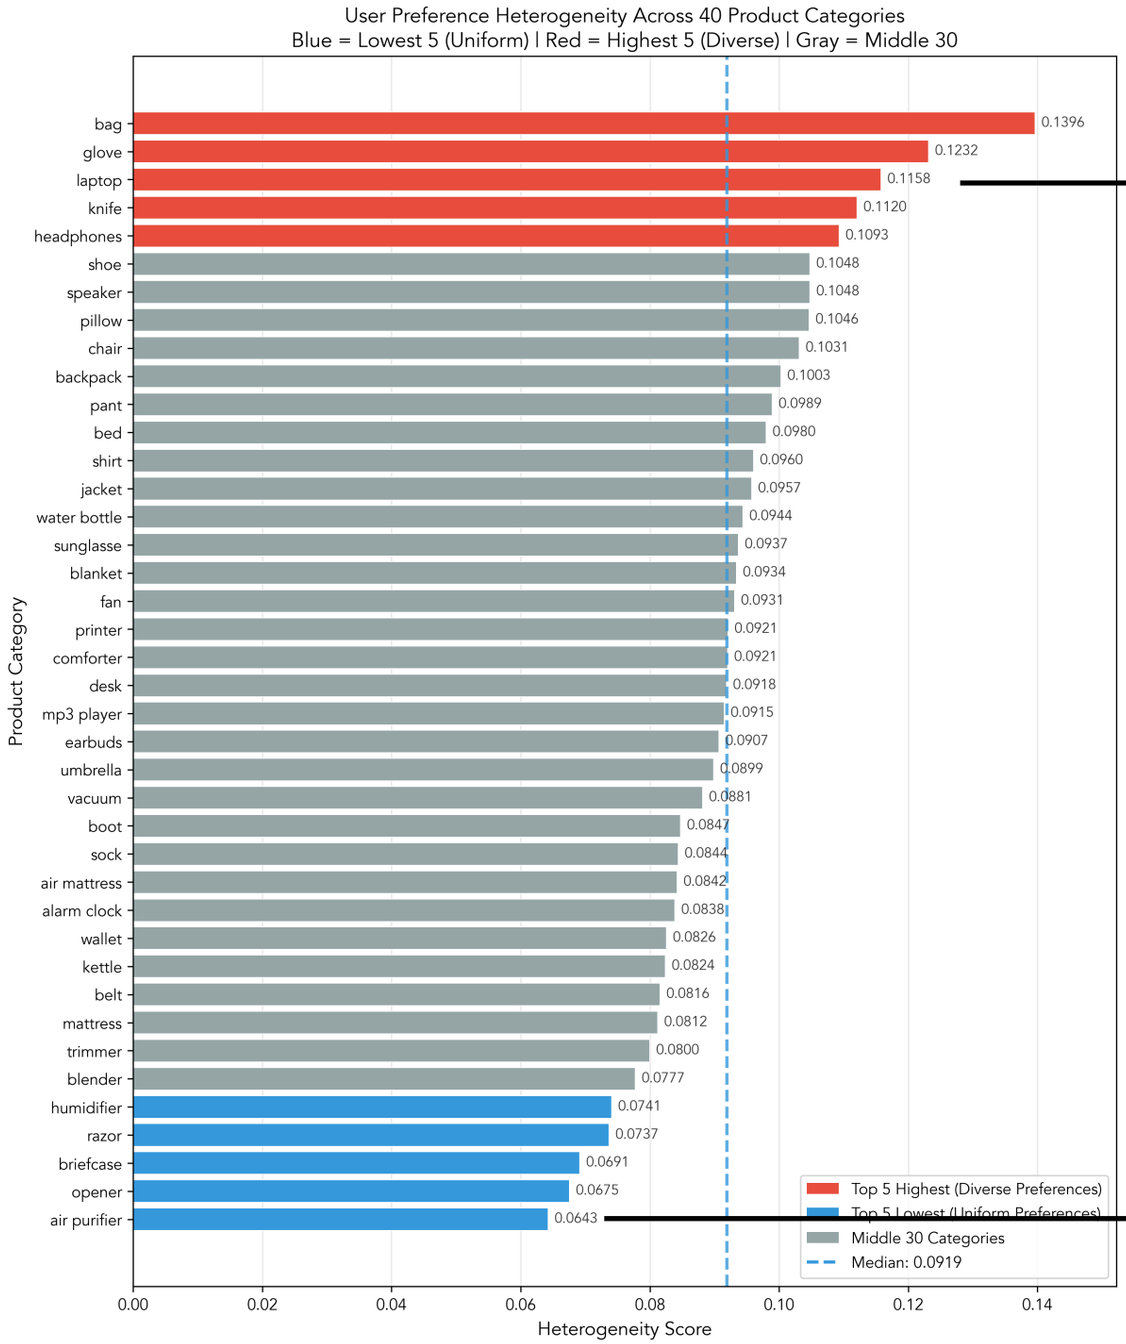


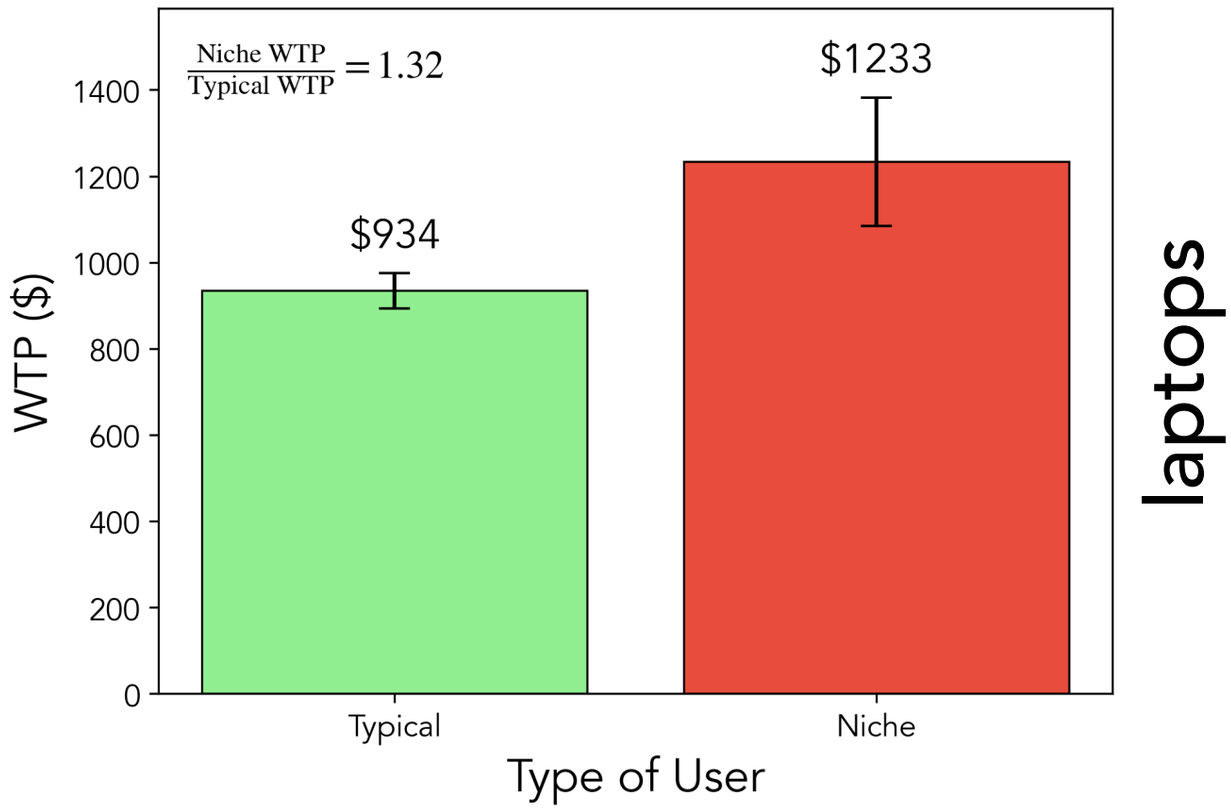


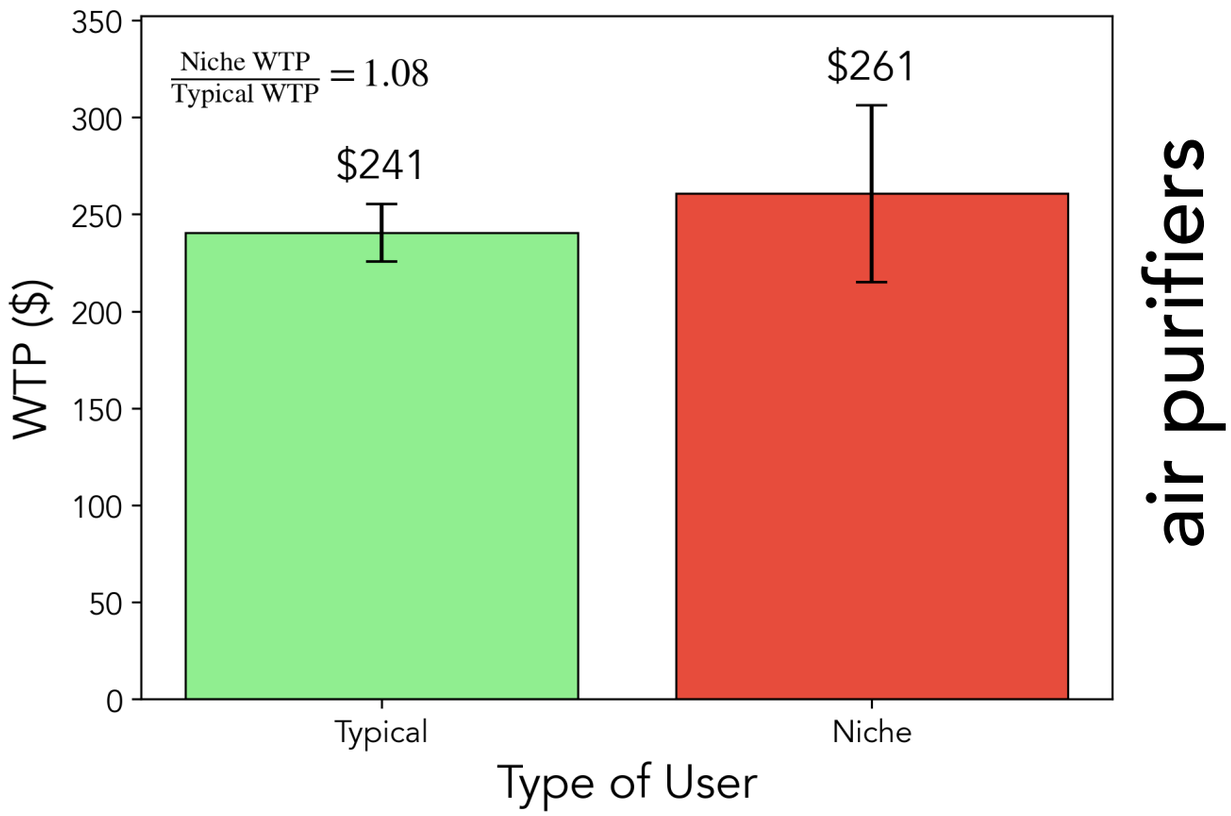
We measure heterogeneity using semantic similarity in preferences

Furthest 10% of users are niche

Preference Heterogeneity & its impact on WTP



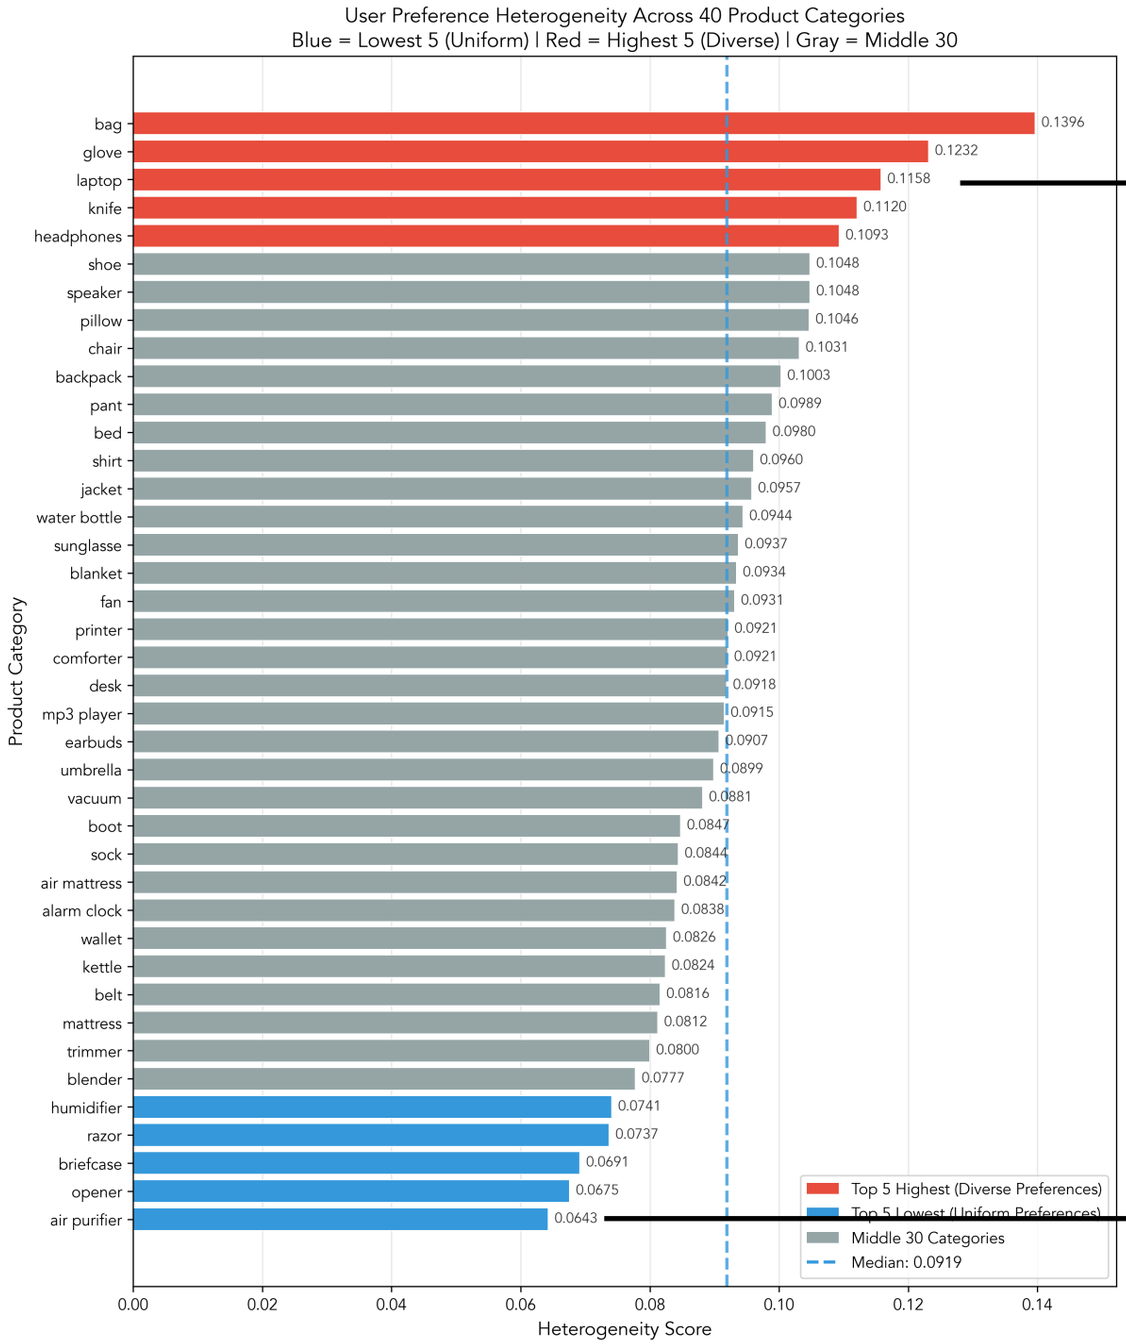


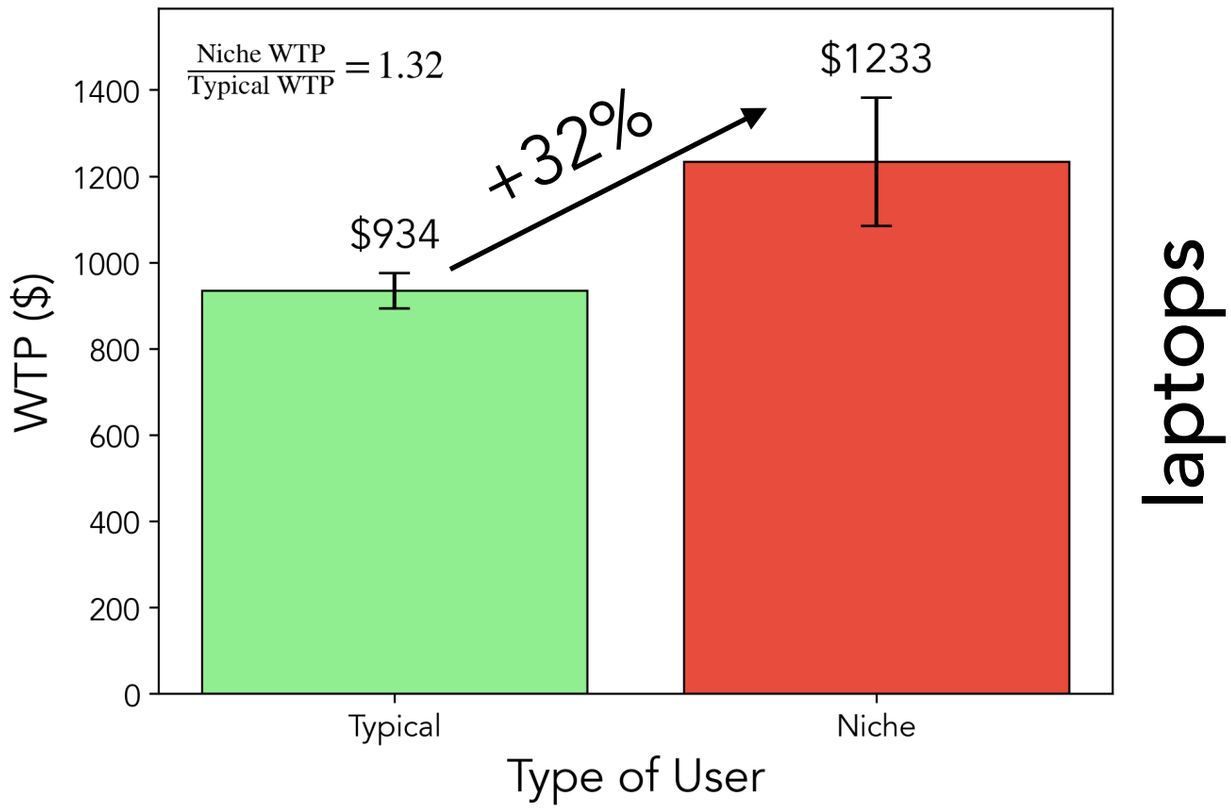


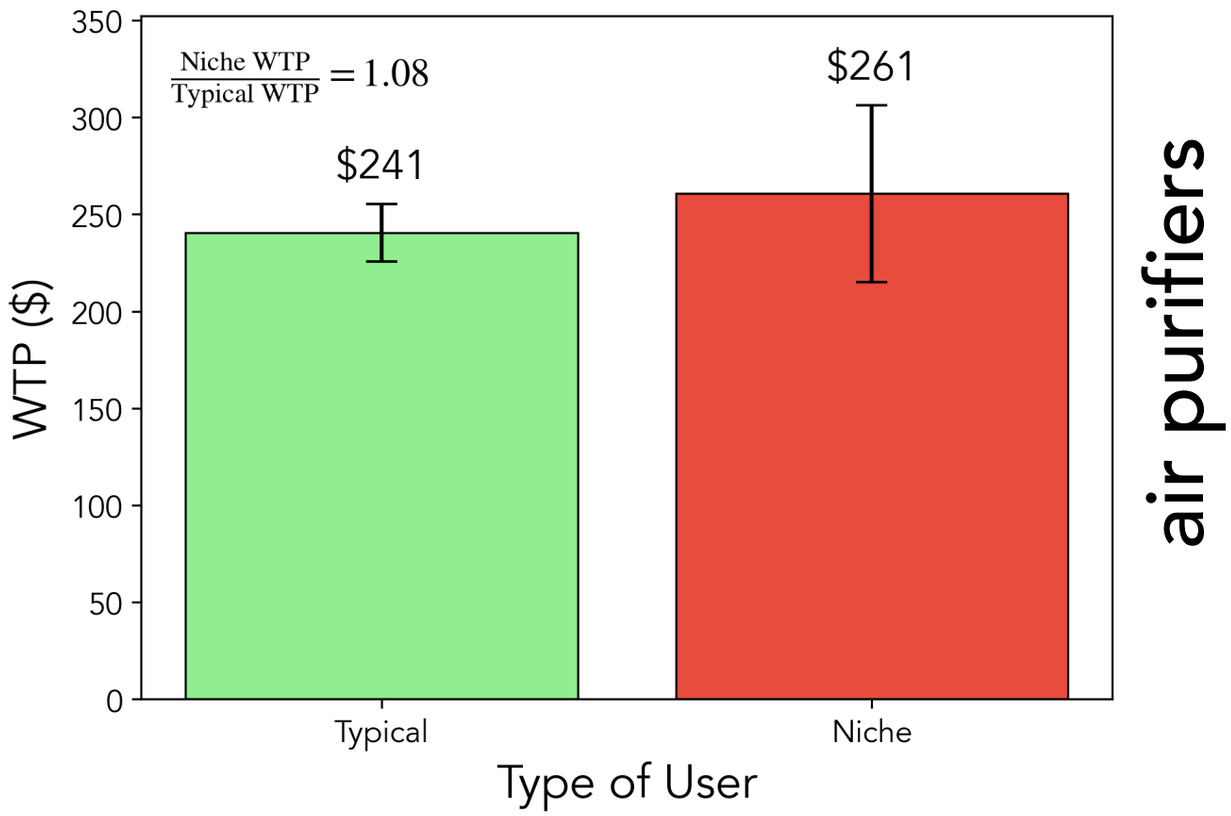
We measure heterogeneity using semantic similarity in preferences

Furthest 10% of users are niche

Preference Heterogeneity & its impact on WTP





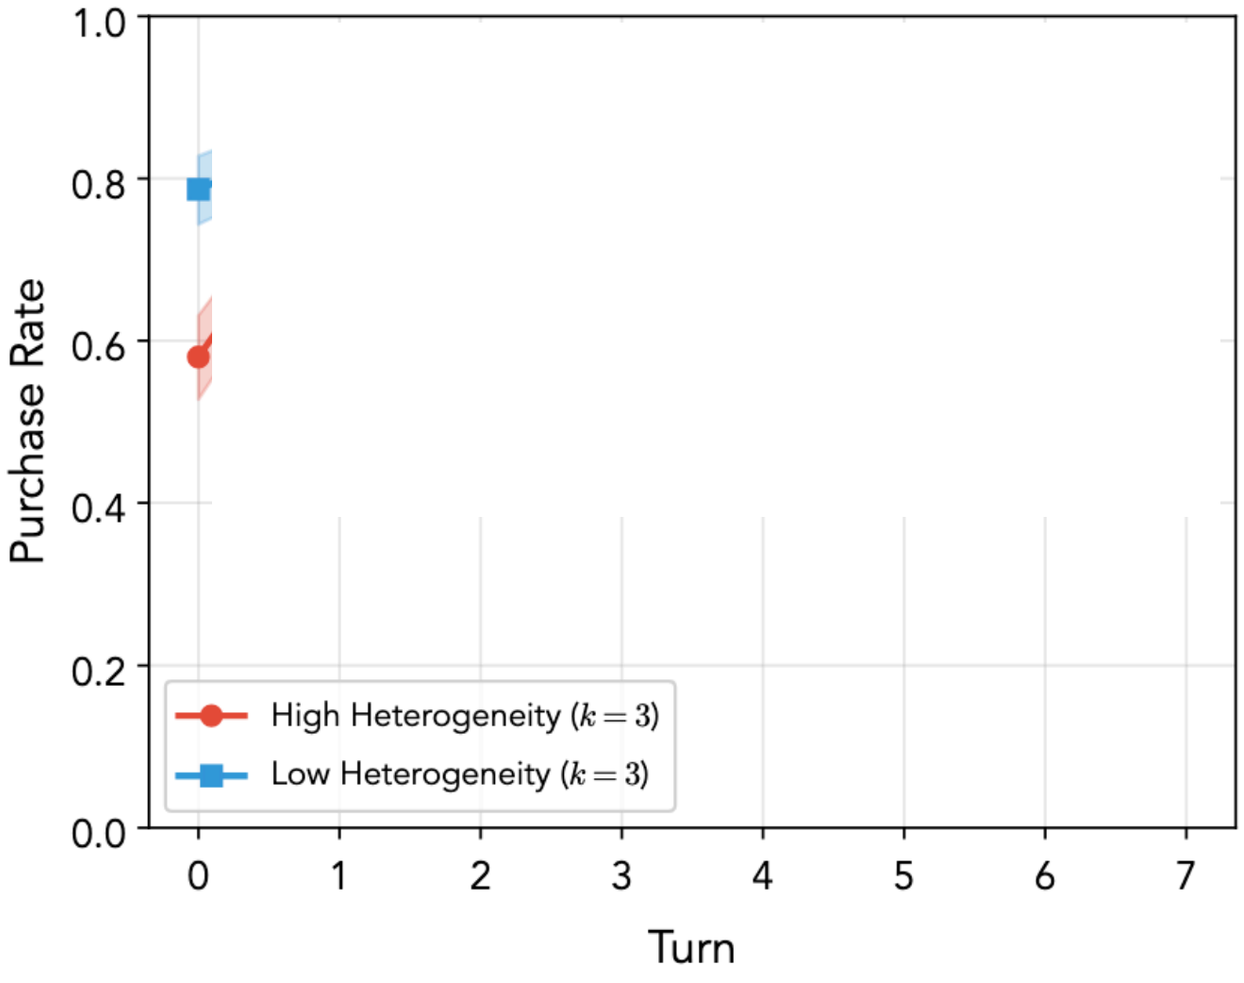


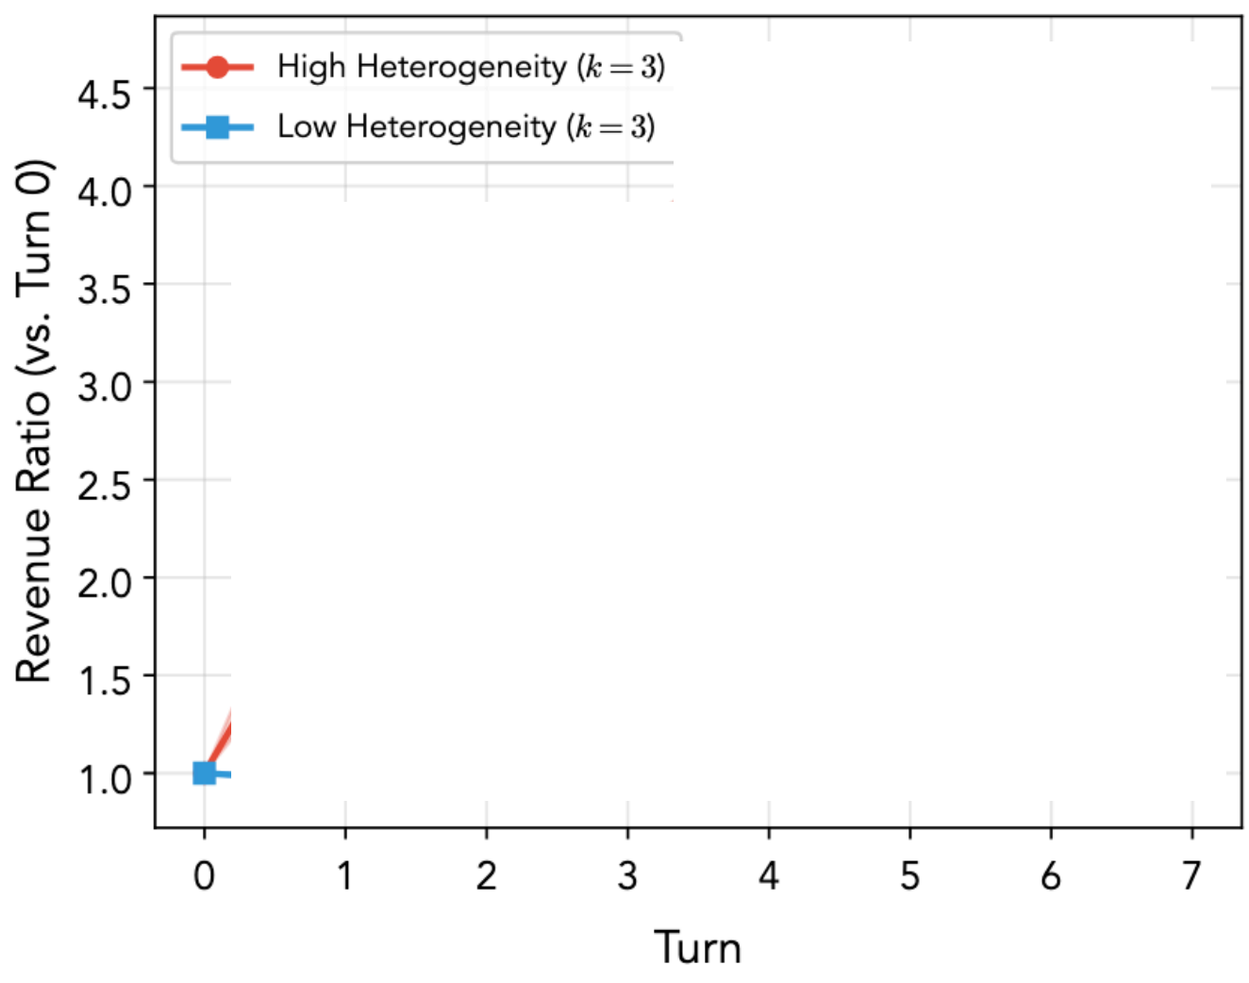
We measure heterogeneity using semantic similarity in preferences

Furthest 10% of users are niche

Conversation is essential for heterogeneous prefs.

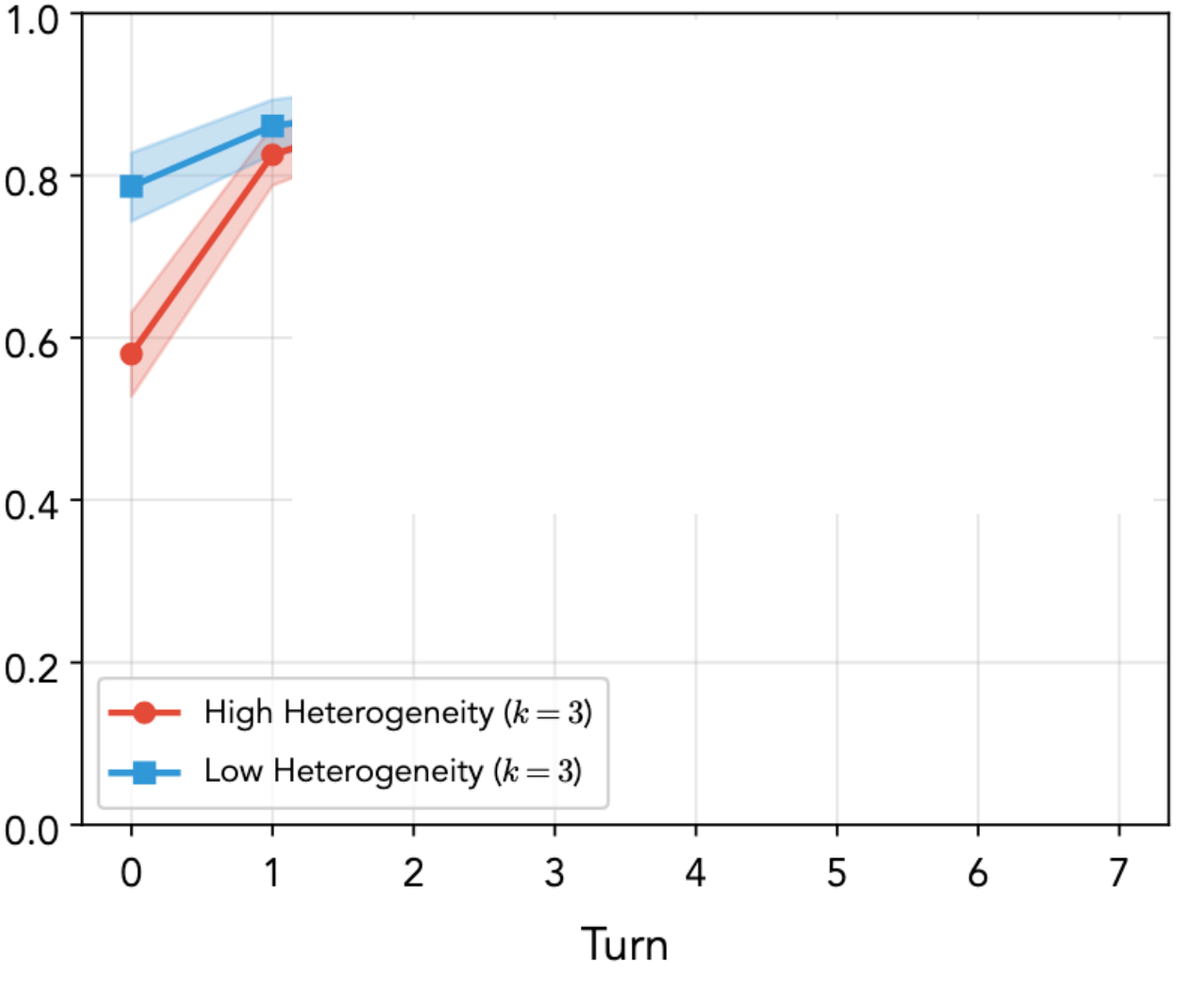
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

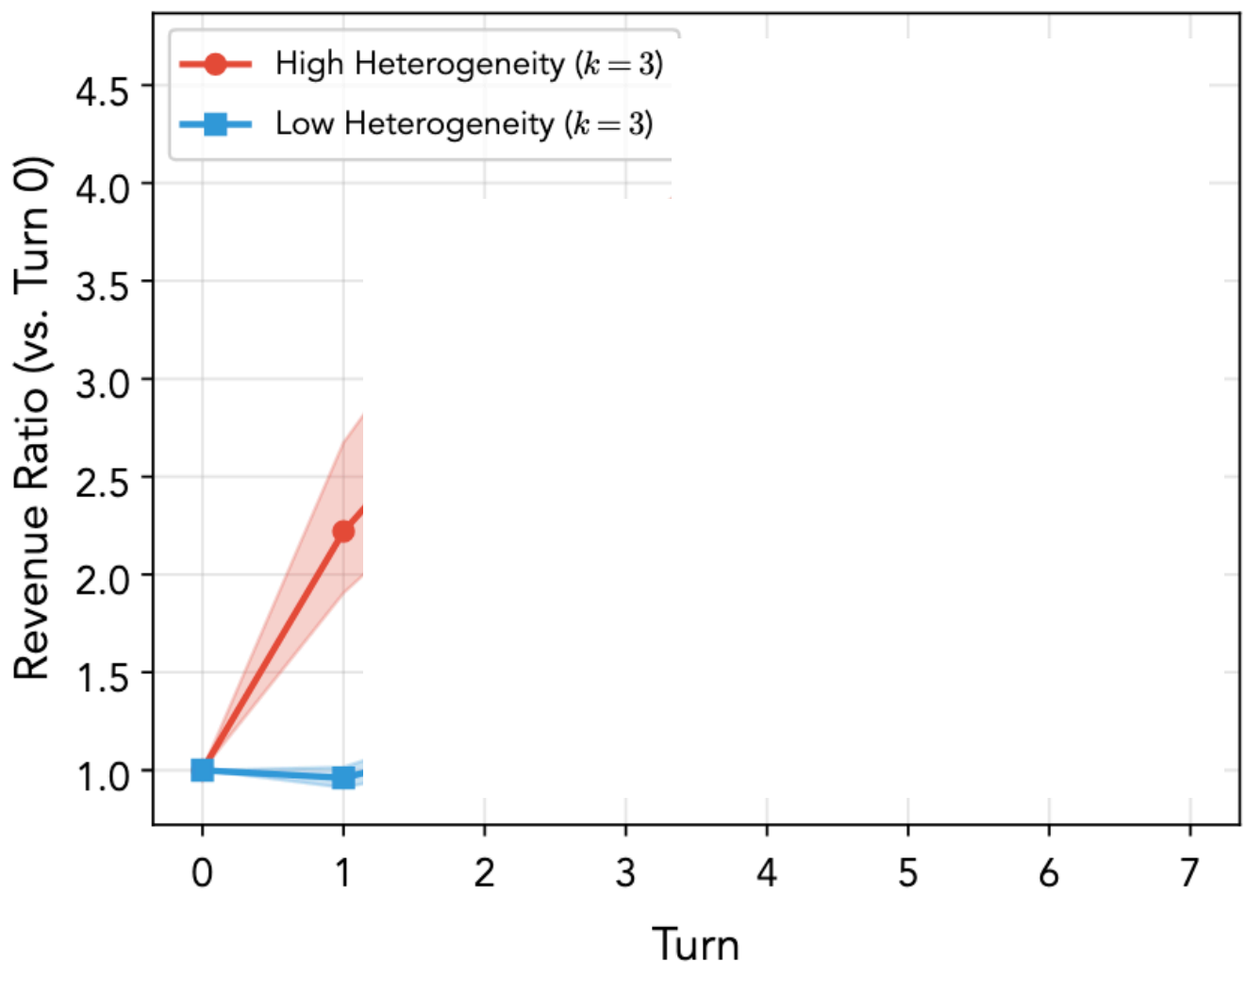




Conversation is essential for heterogeneous prefs.

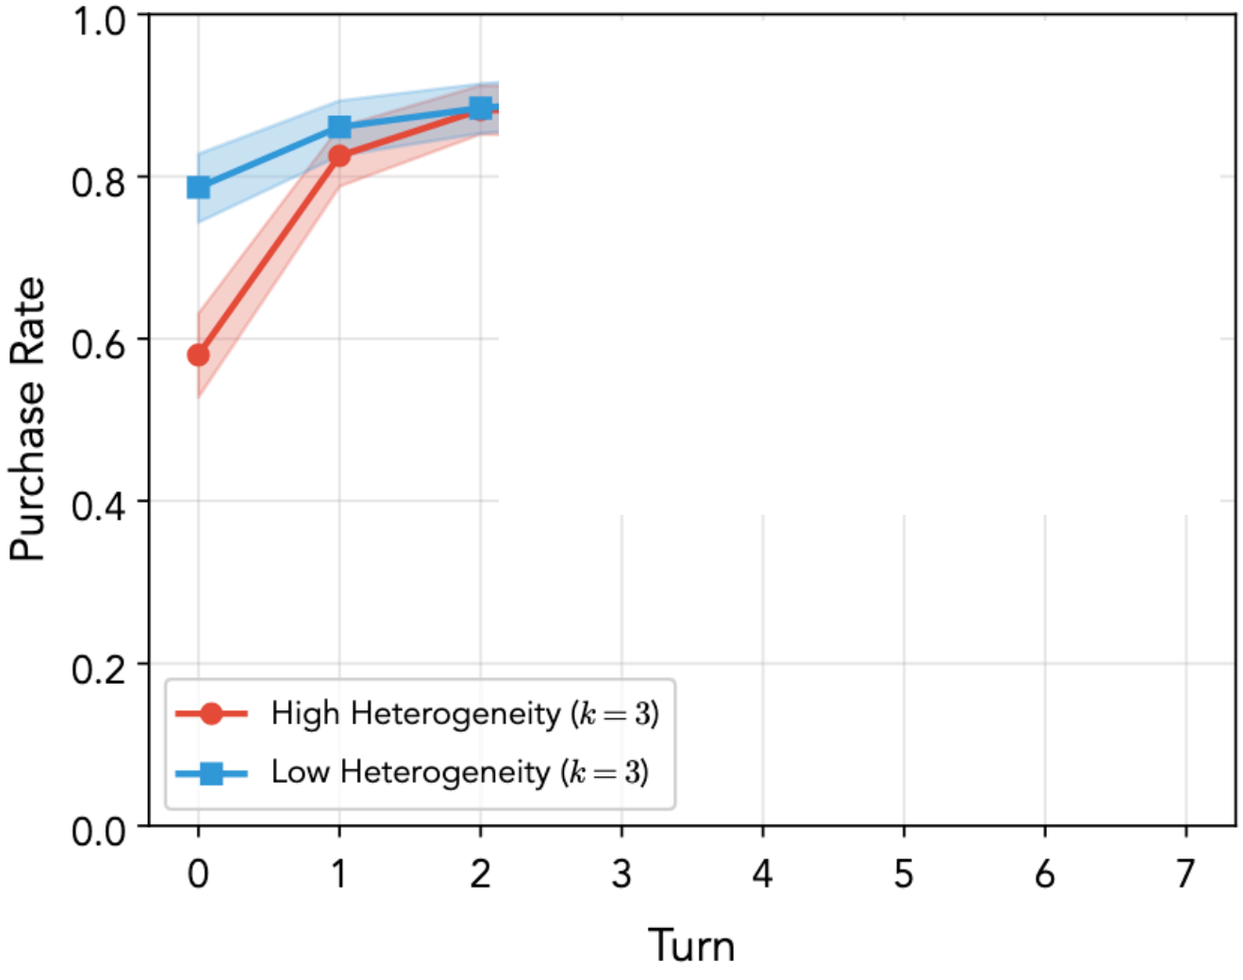
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

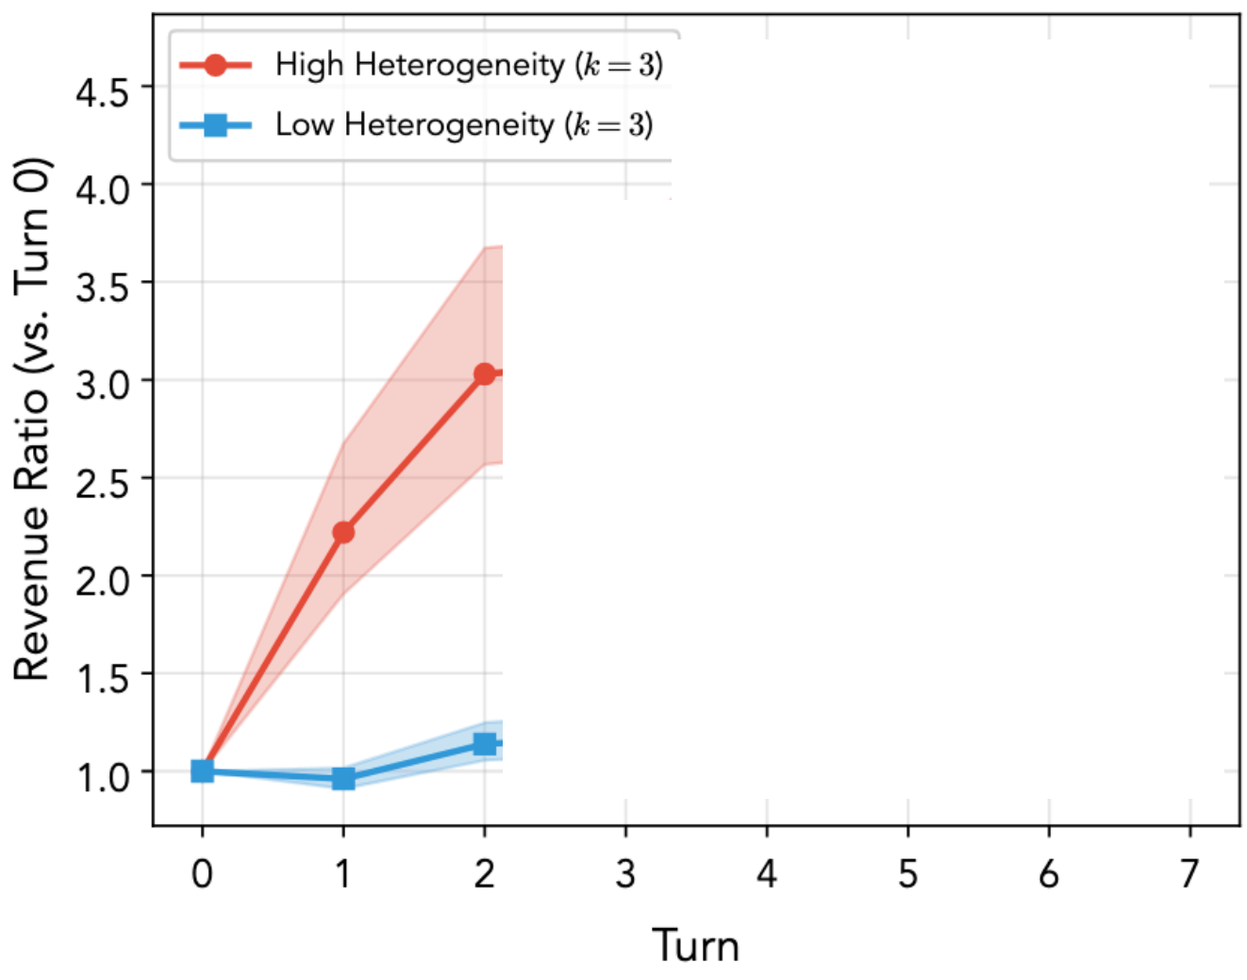




Conversation is essential for heterogeneous prefs.

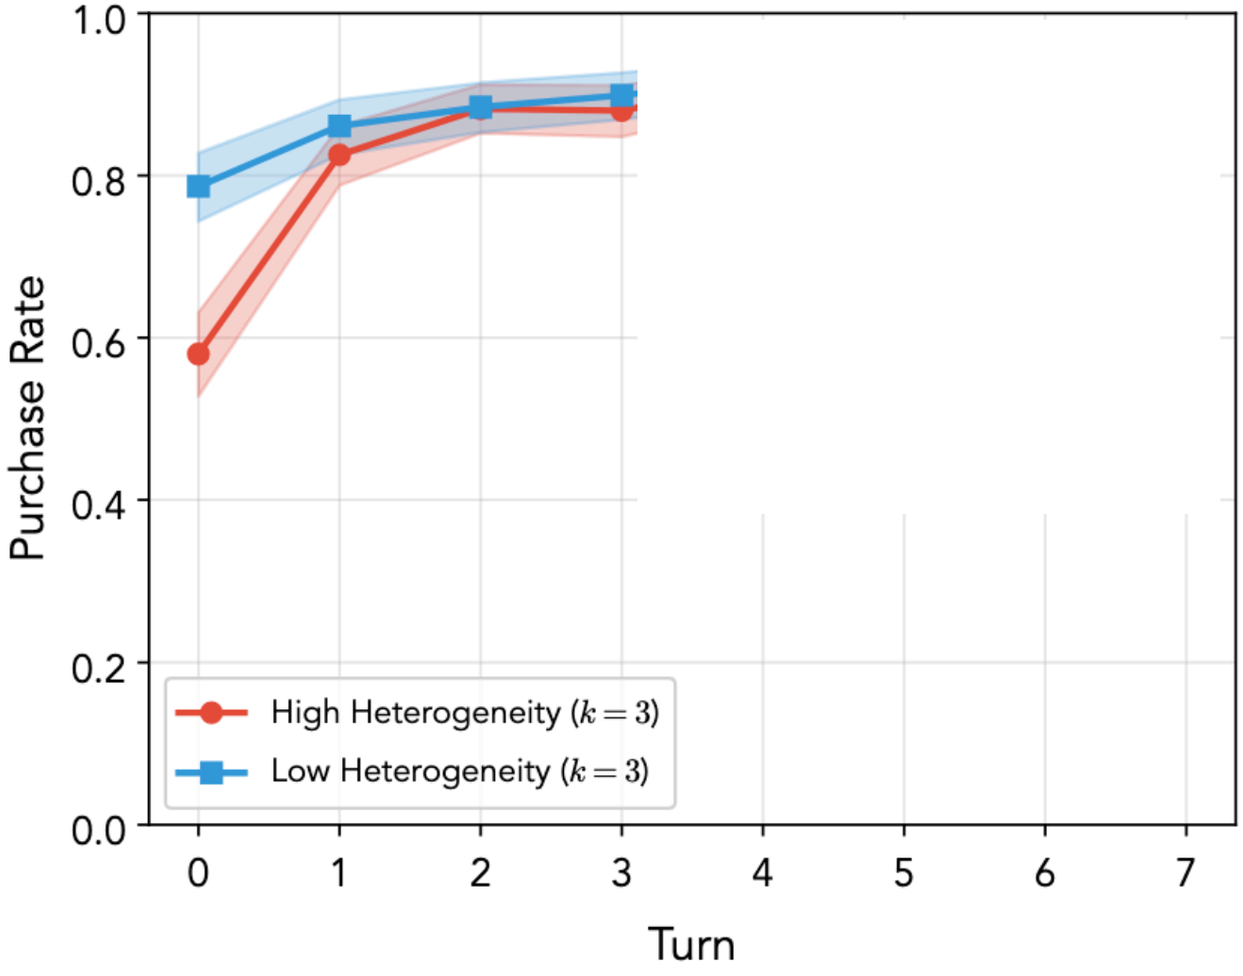
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

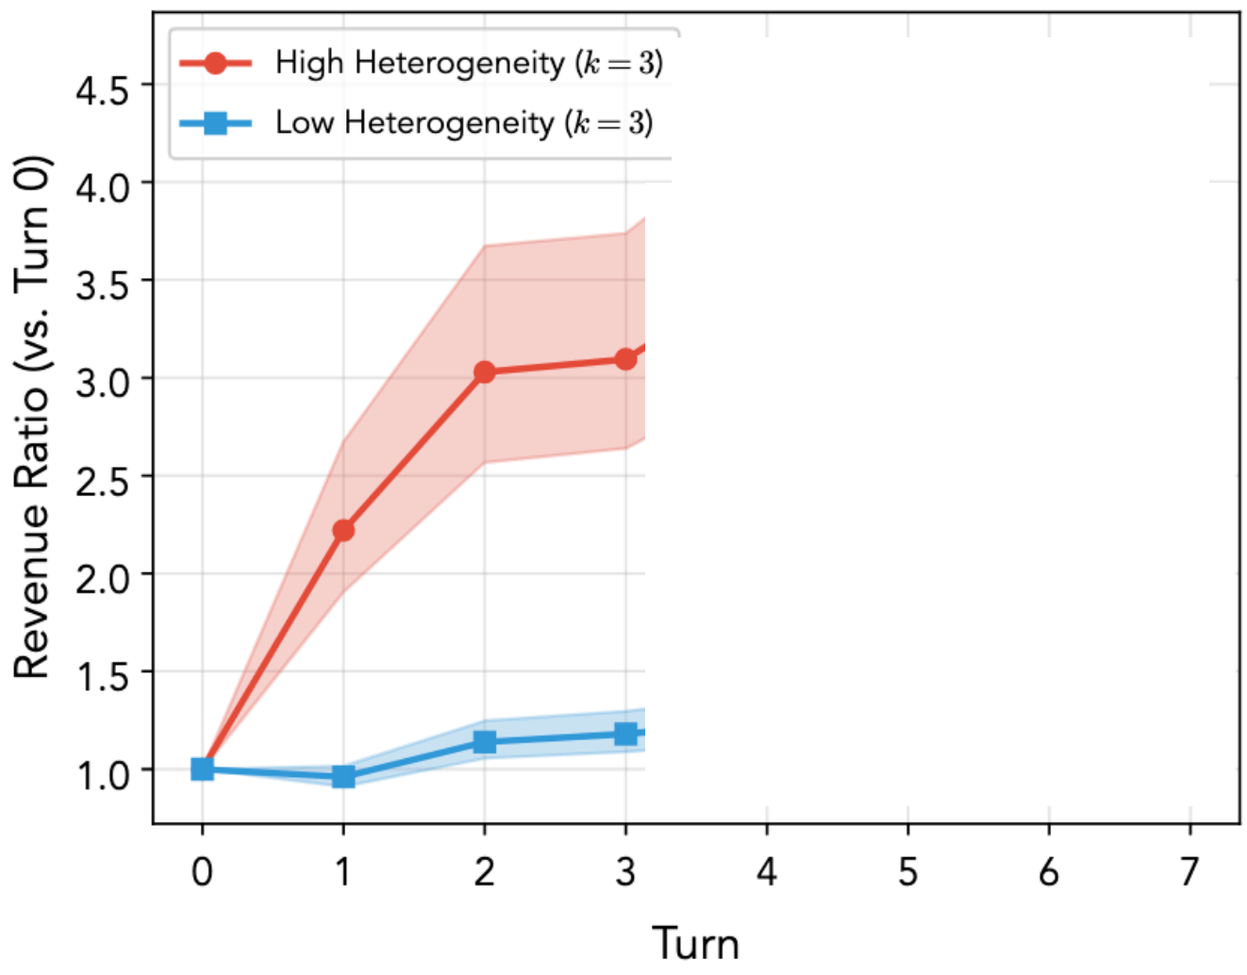




Conversation is essential for heterogeneous prefs.

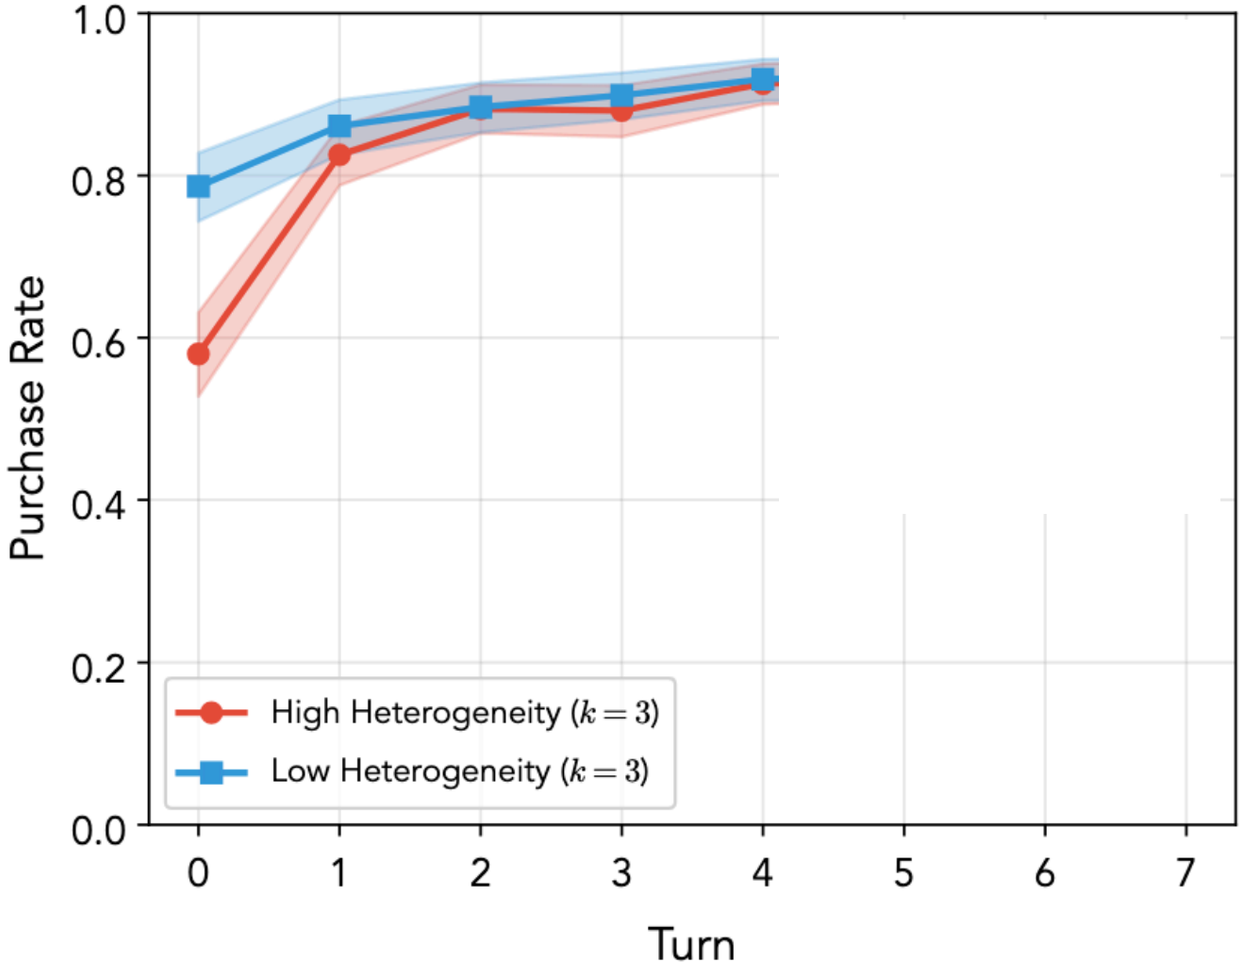
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

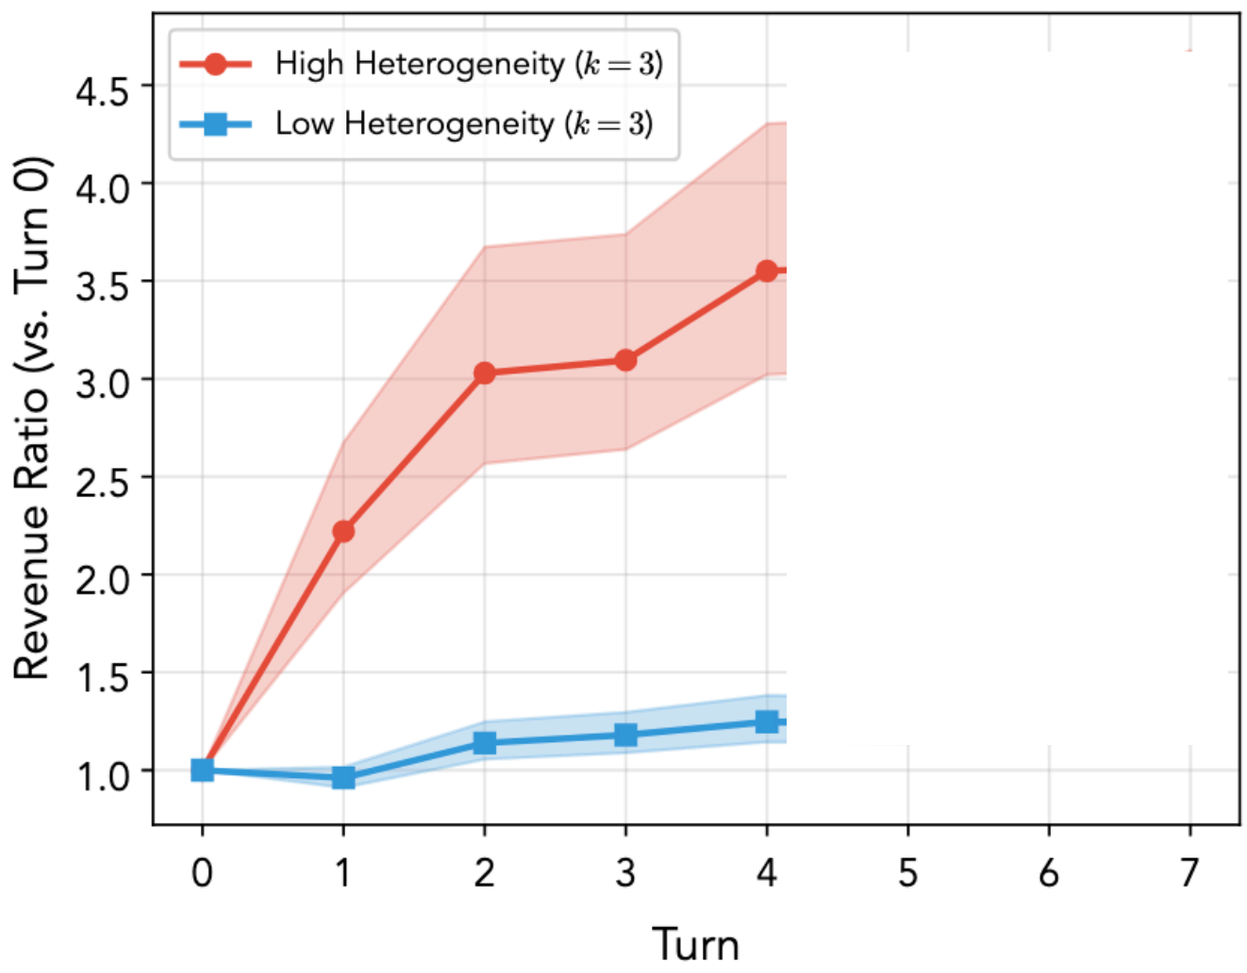




Conversation is essential for heterogeneous prefs.

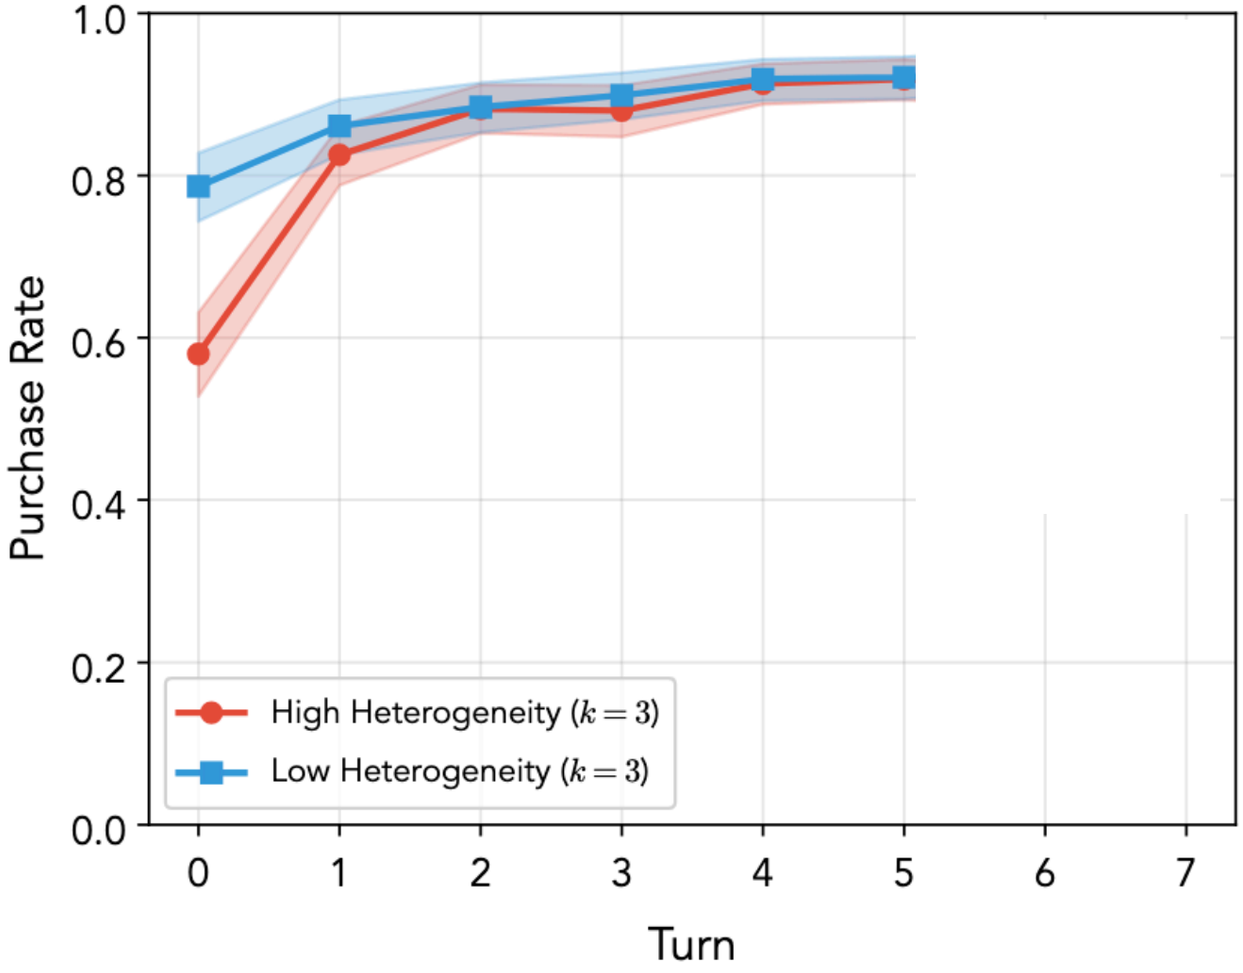
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

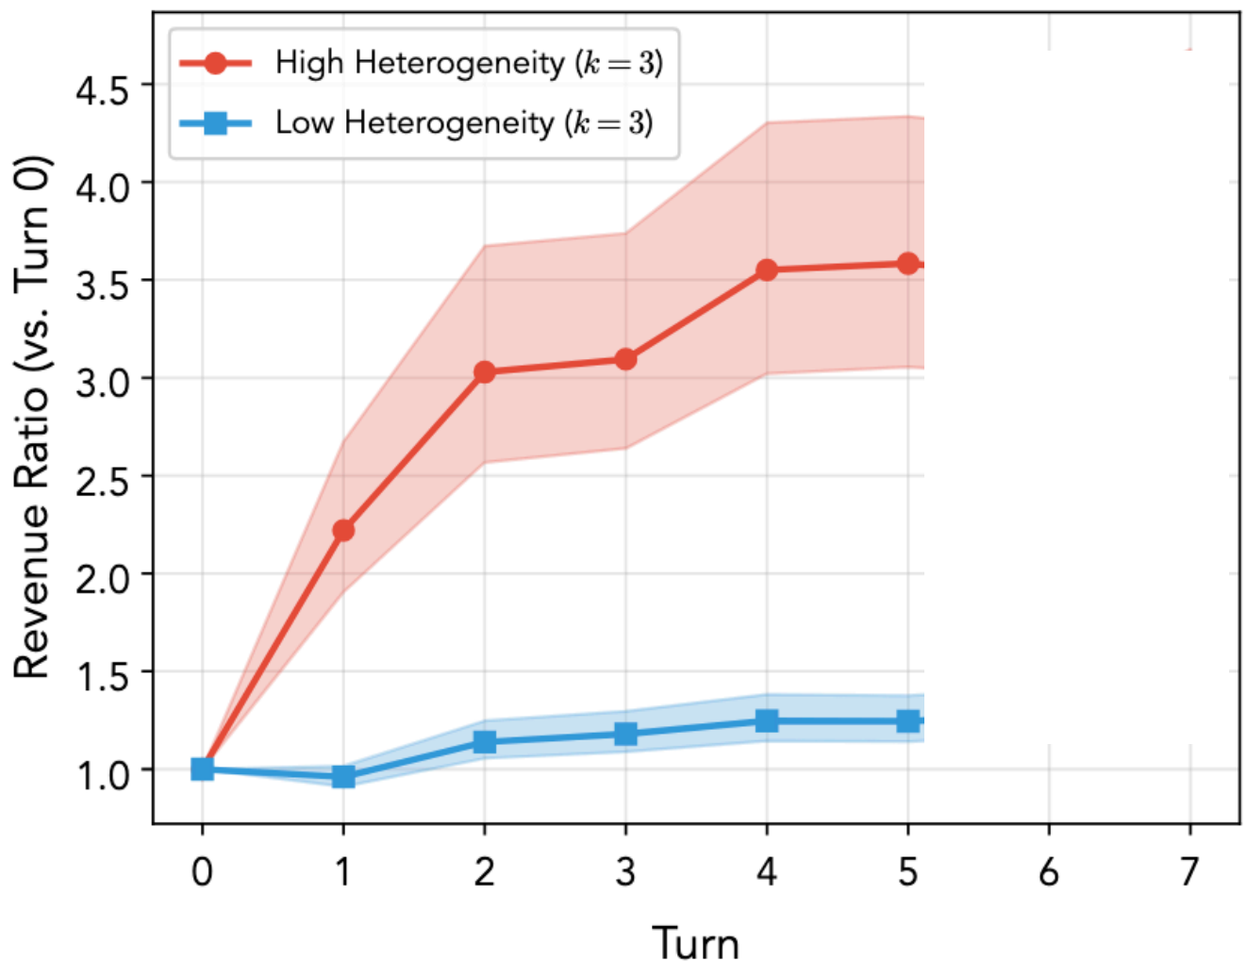




Conversation is essential for heterogeneous prefs.

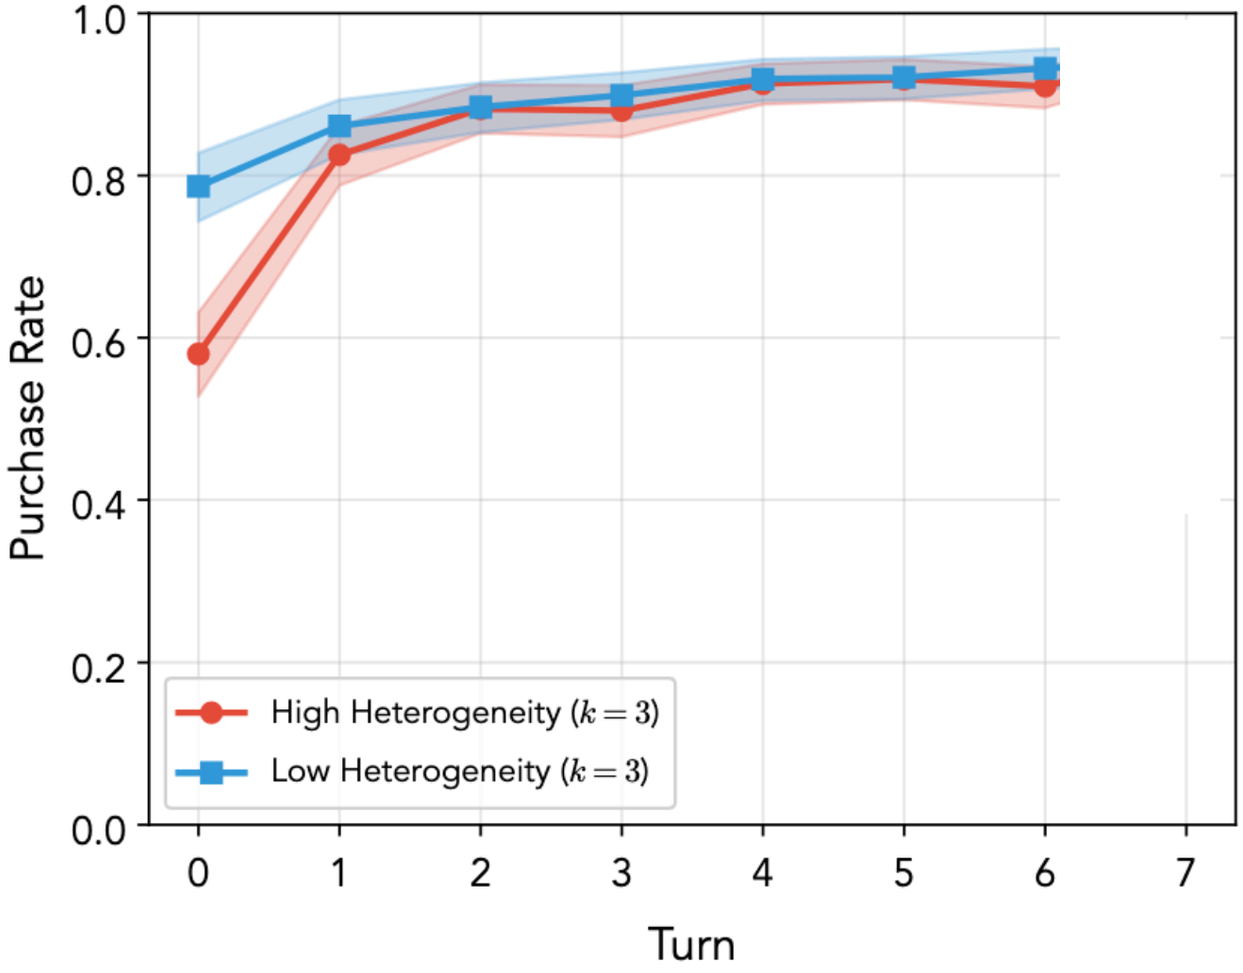
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

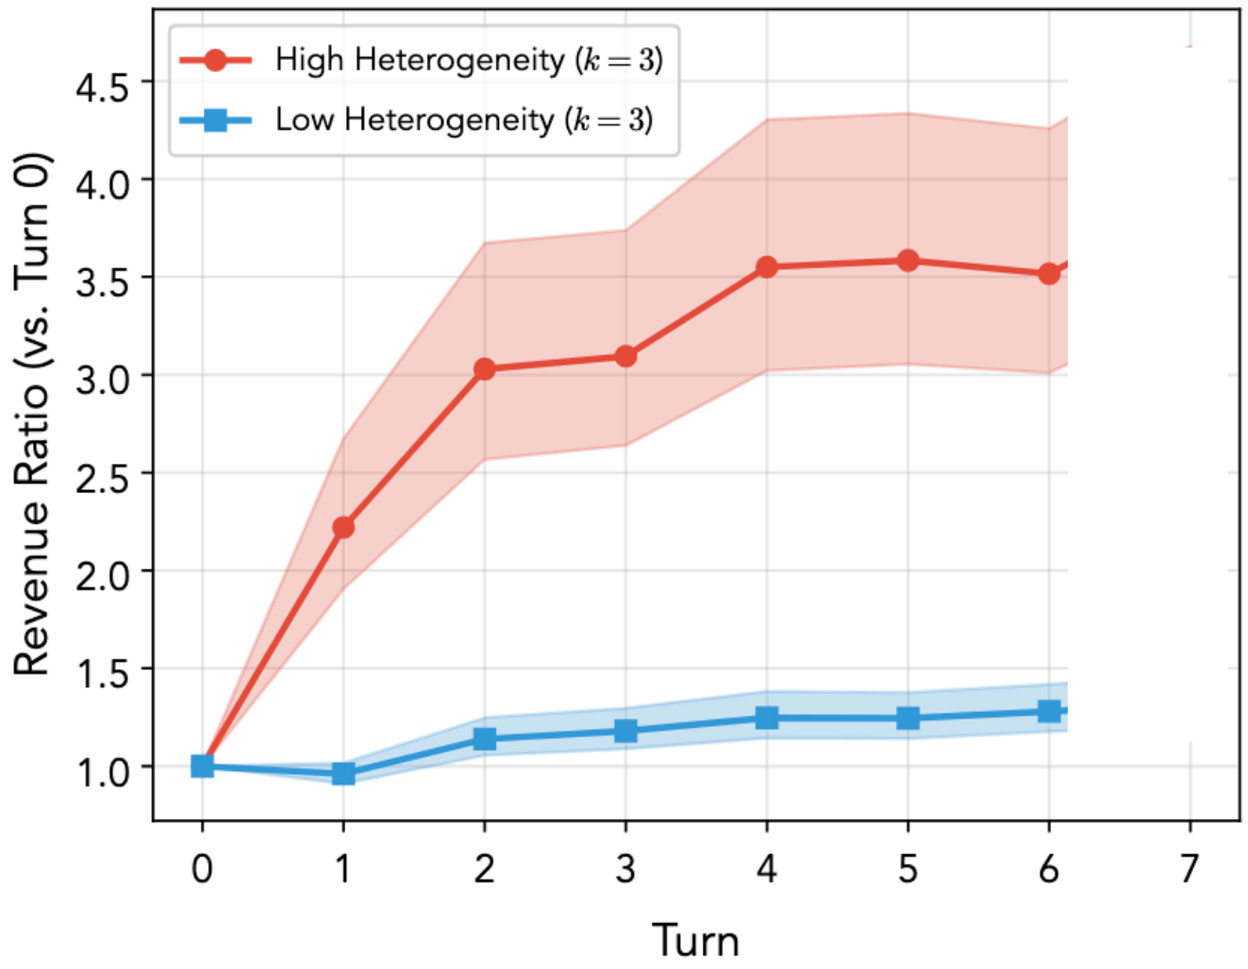




Conversation is essential for heterogeneous prefs.

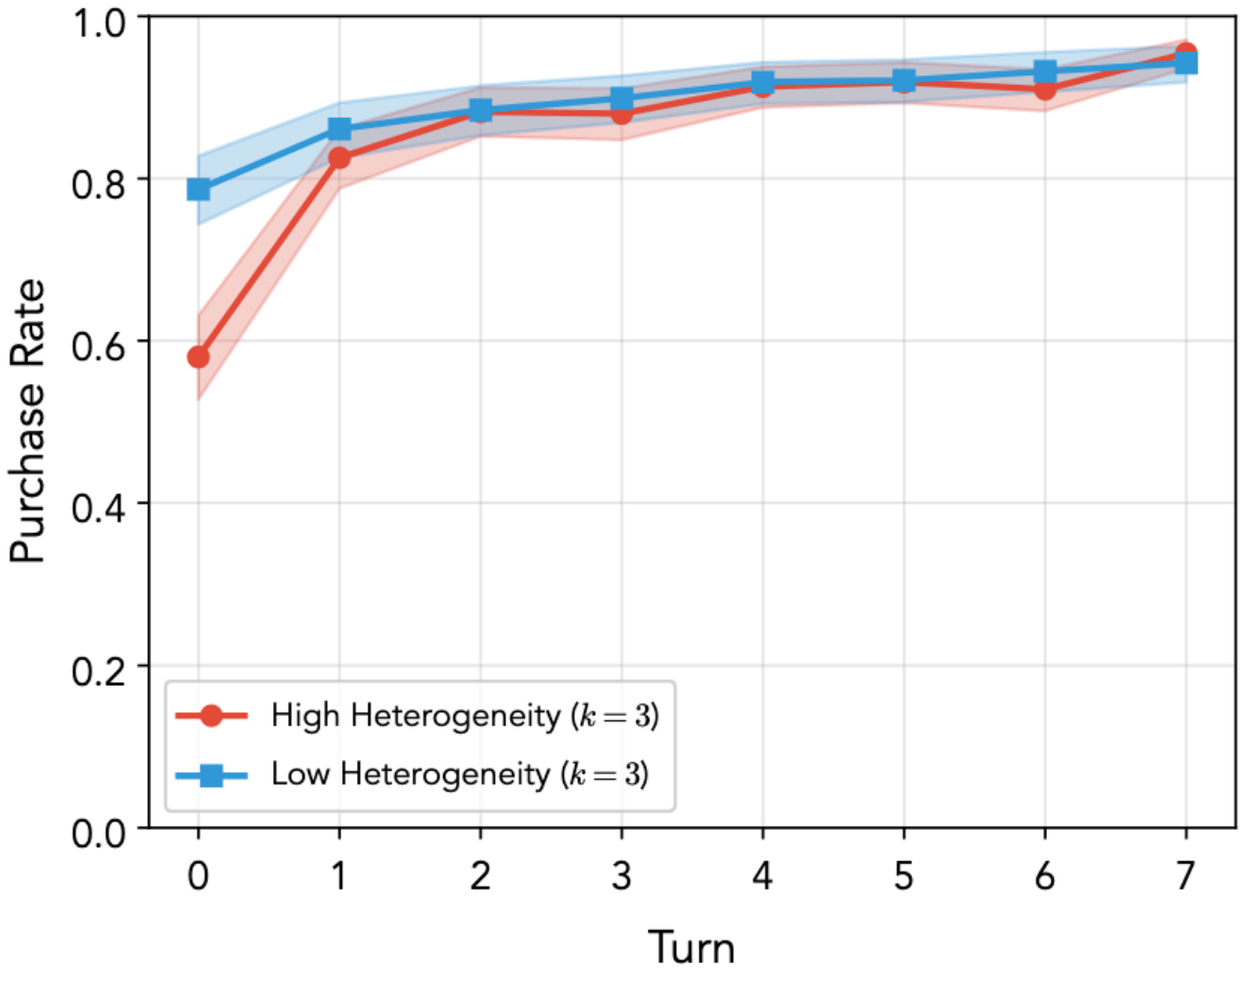
- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

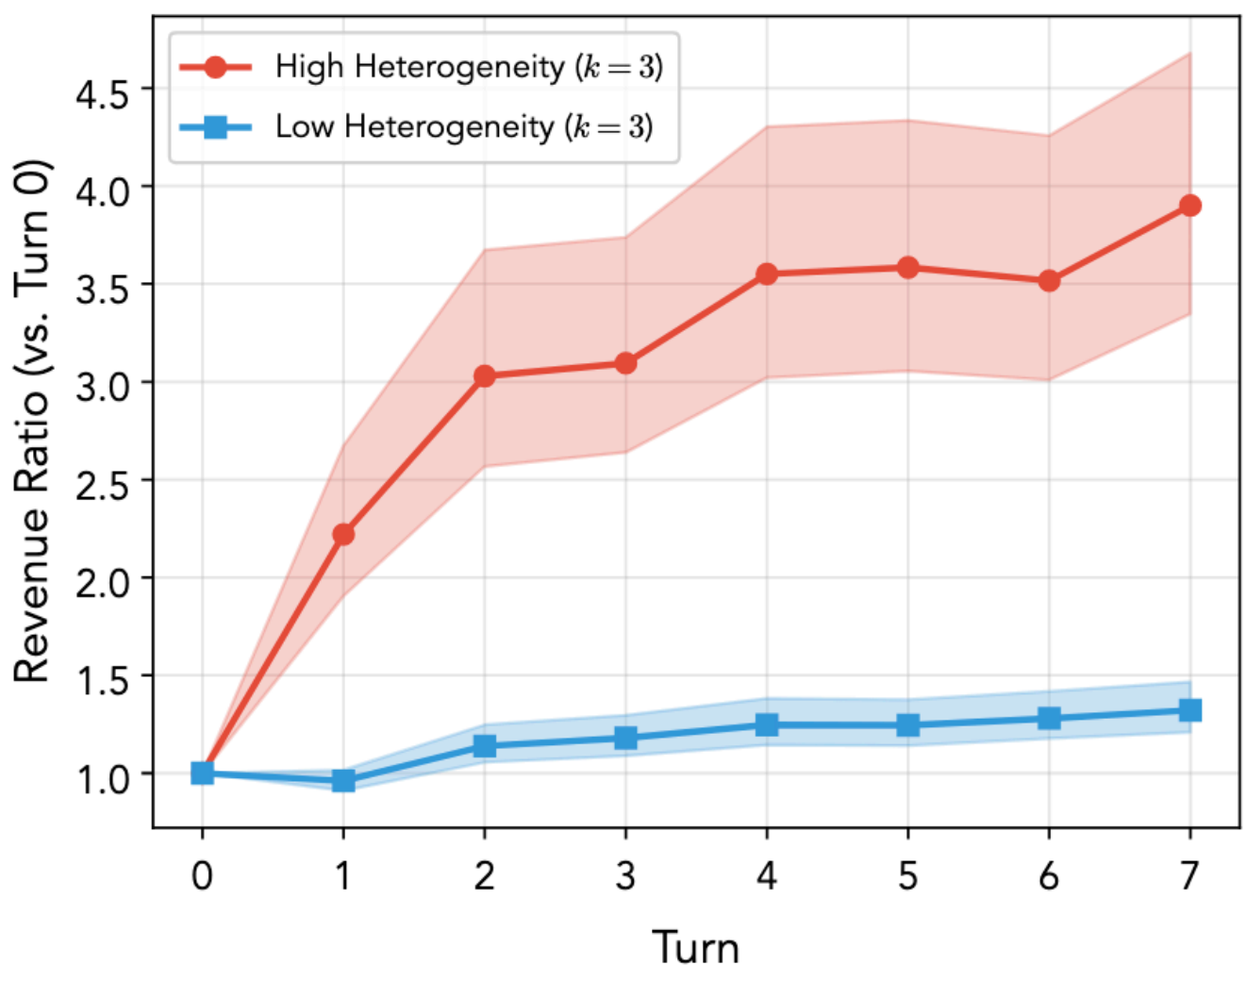




Conversation is essential for heterogeneous prefs.

- **More information** leads to **more purchases/higher revenue.**
- **Gains are substantially larger** in categories with **high heterogeneity**.

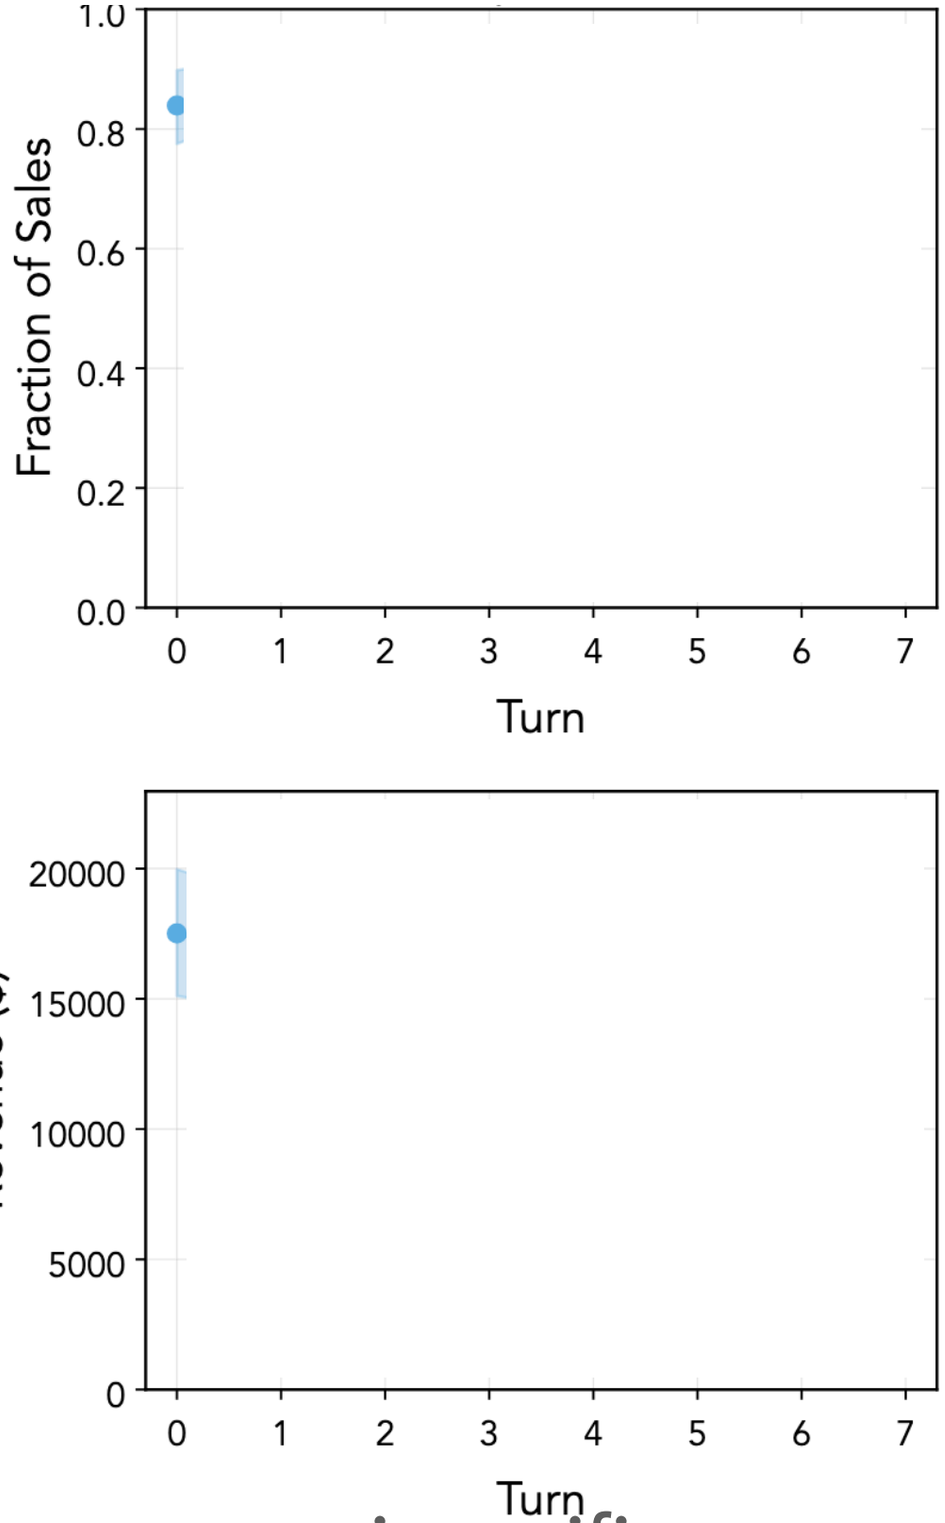


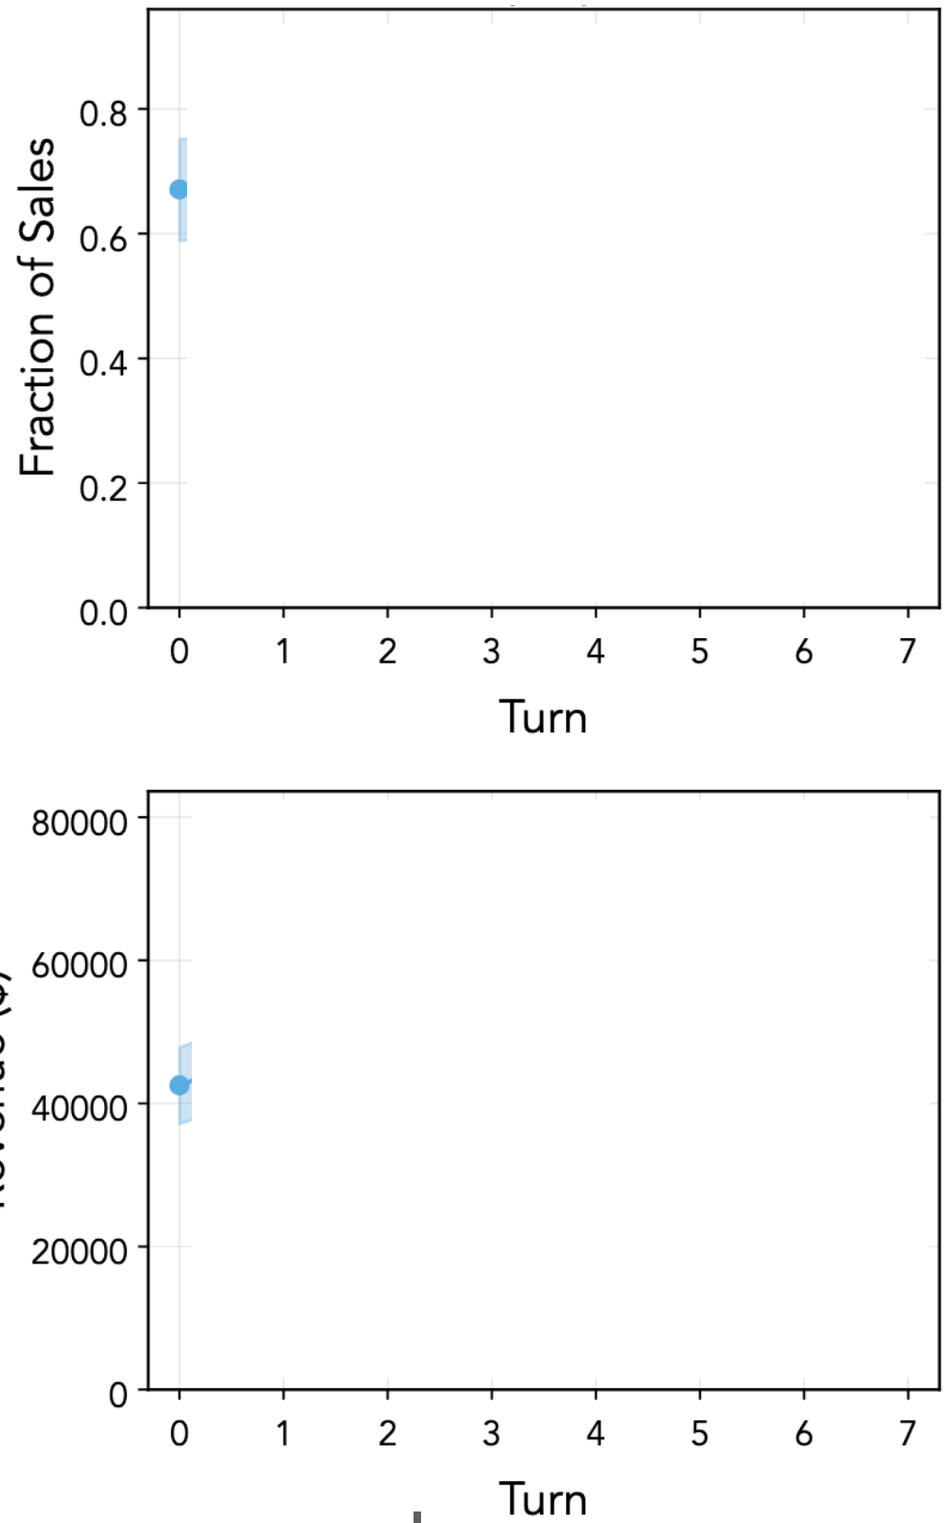


Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

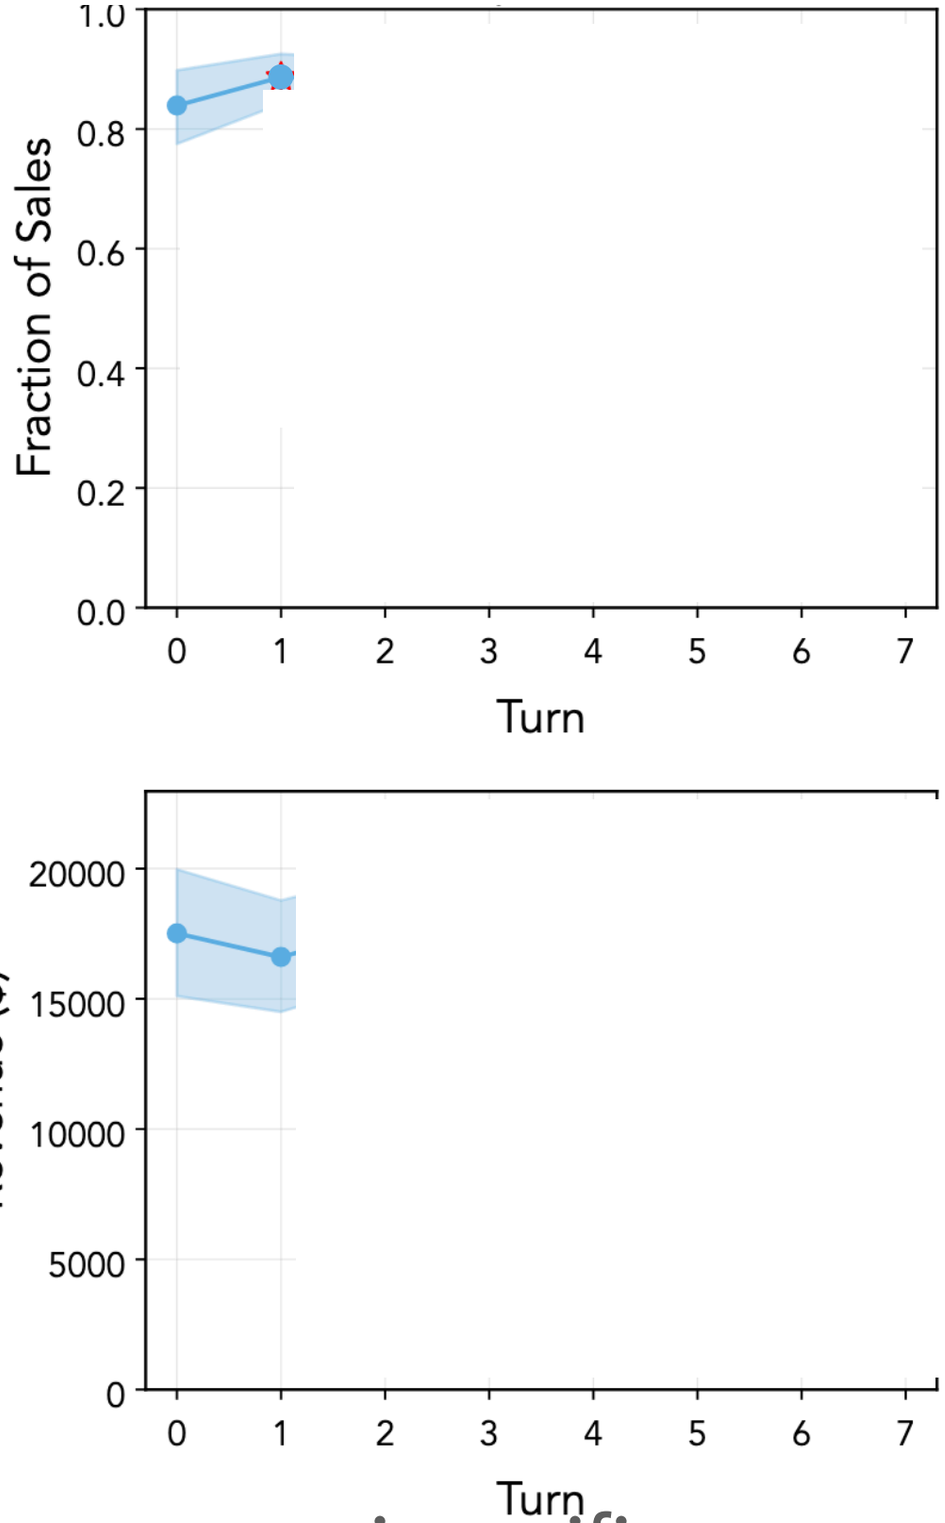
air purifiers

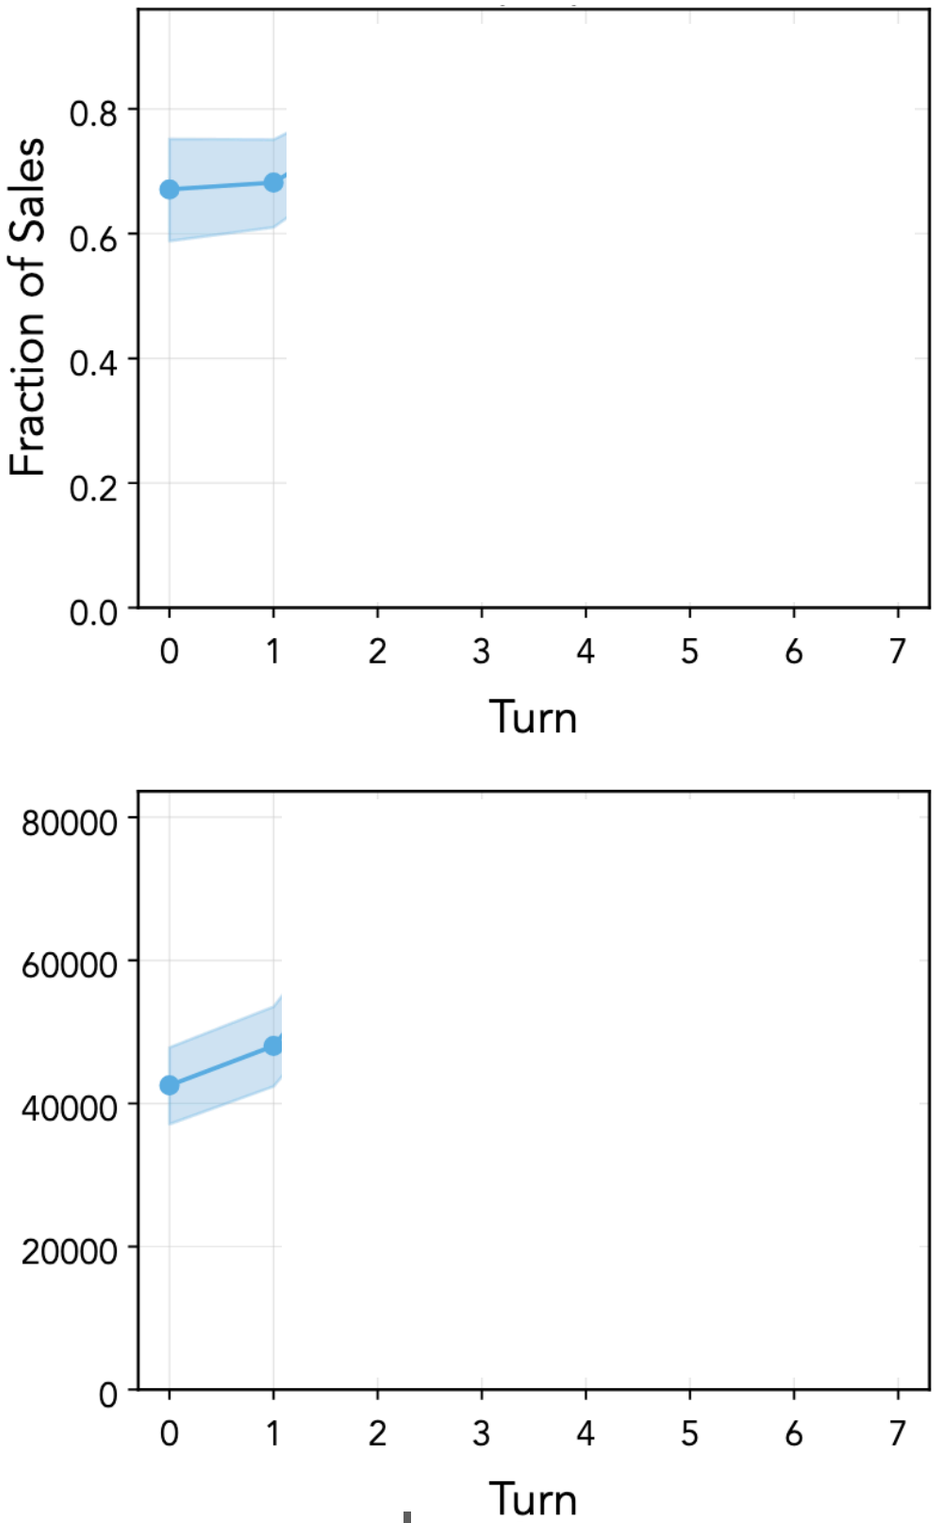
laptops

Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

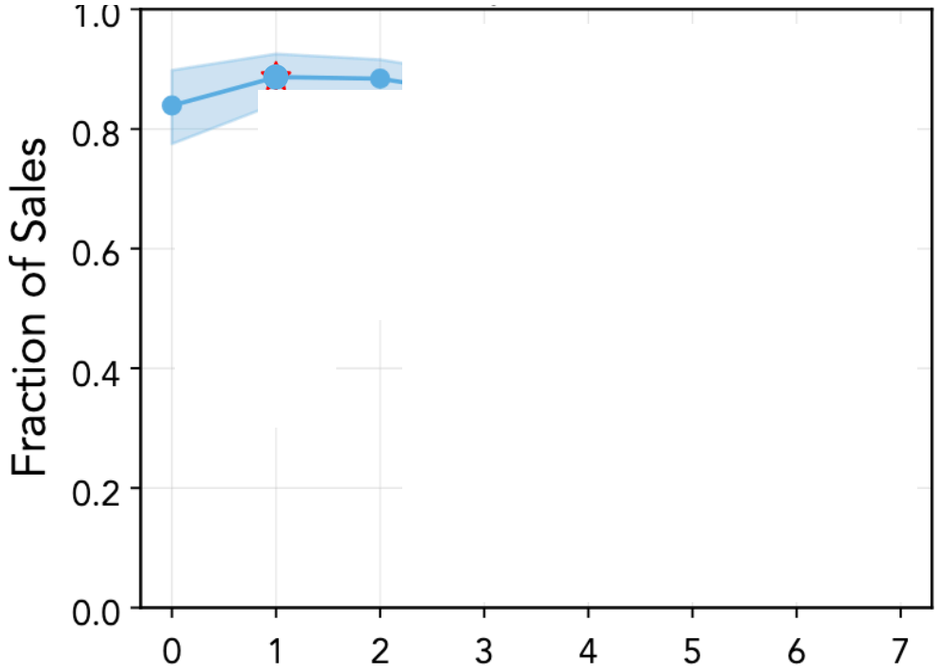
air purifiers

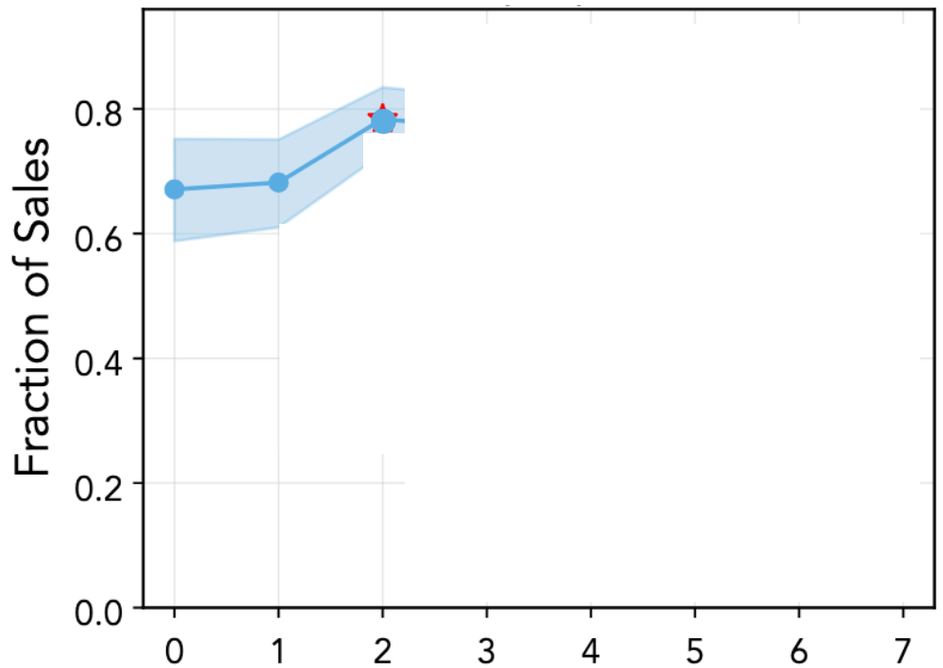
laptops

Testbed with simulated user abandonment

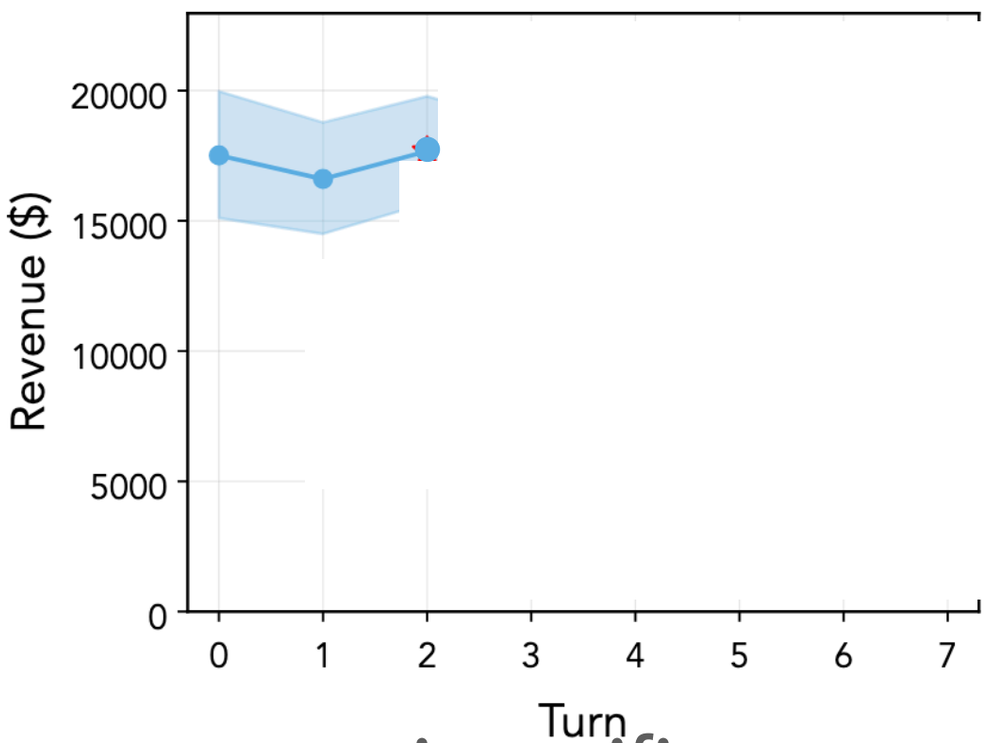
at each turn 3% of users who haven't purchased drop-out

Conversion

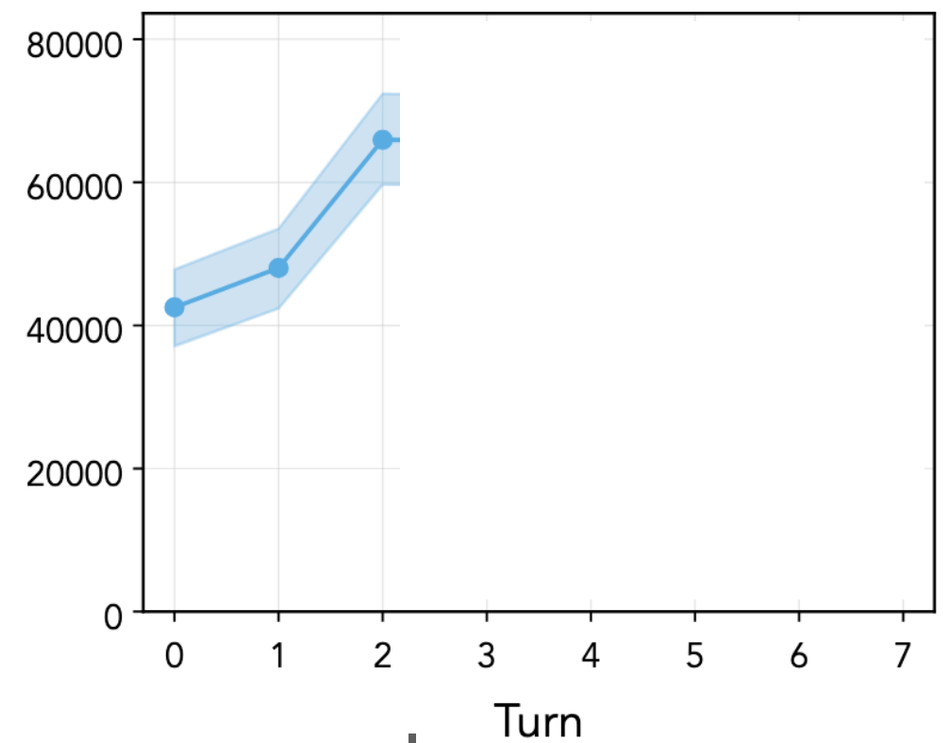




Commission



air purifiers

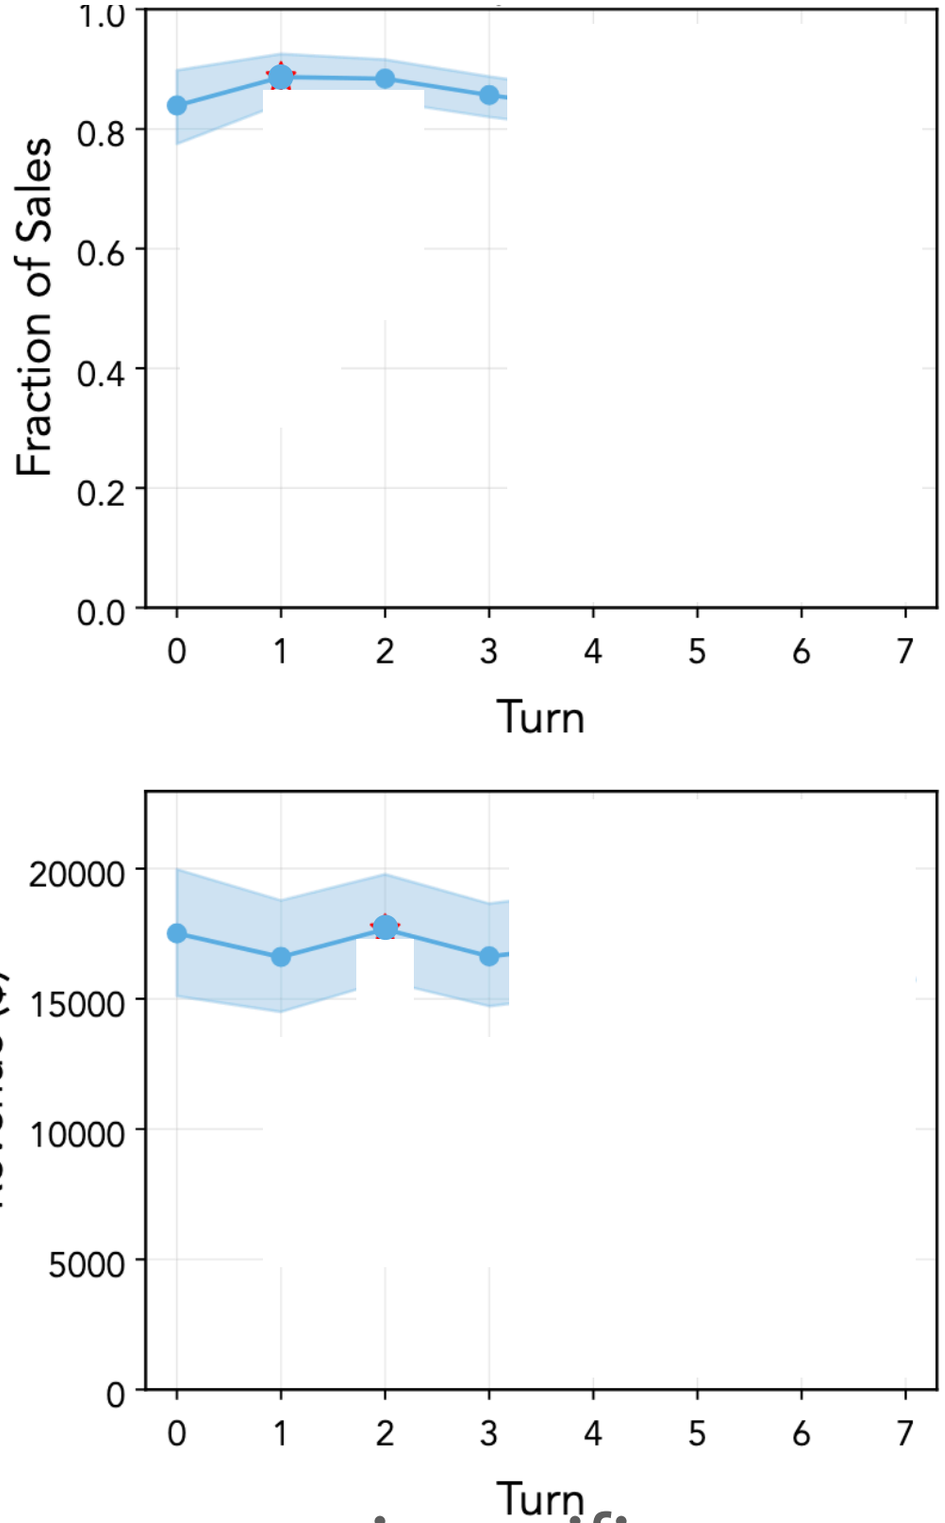


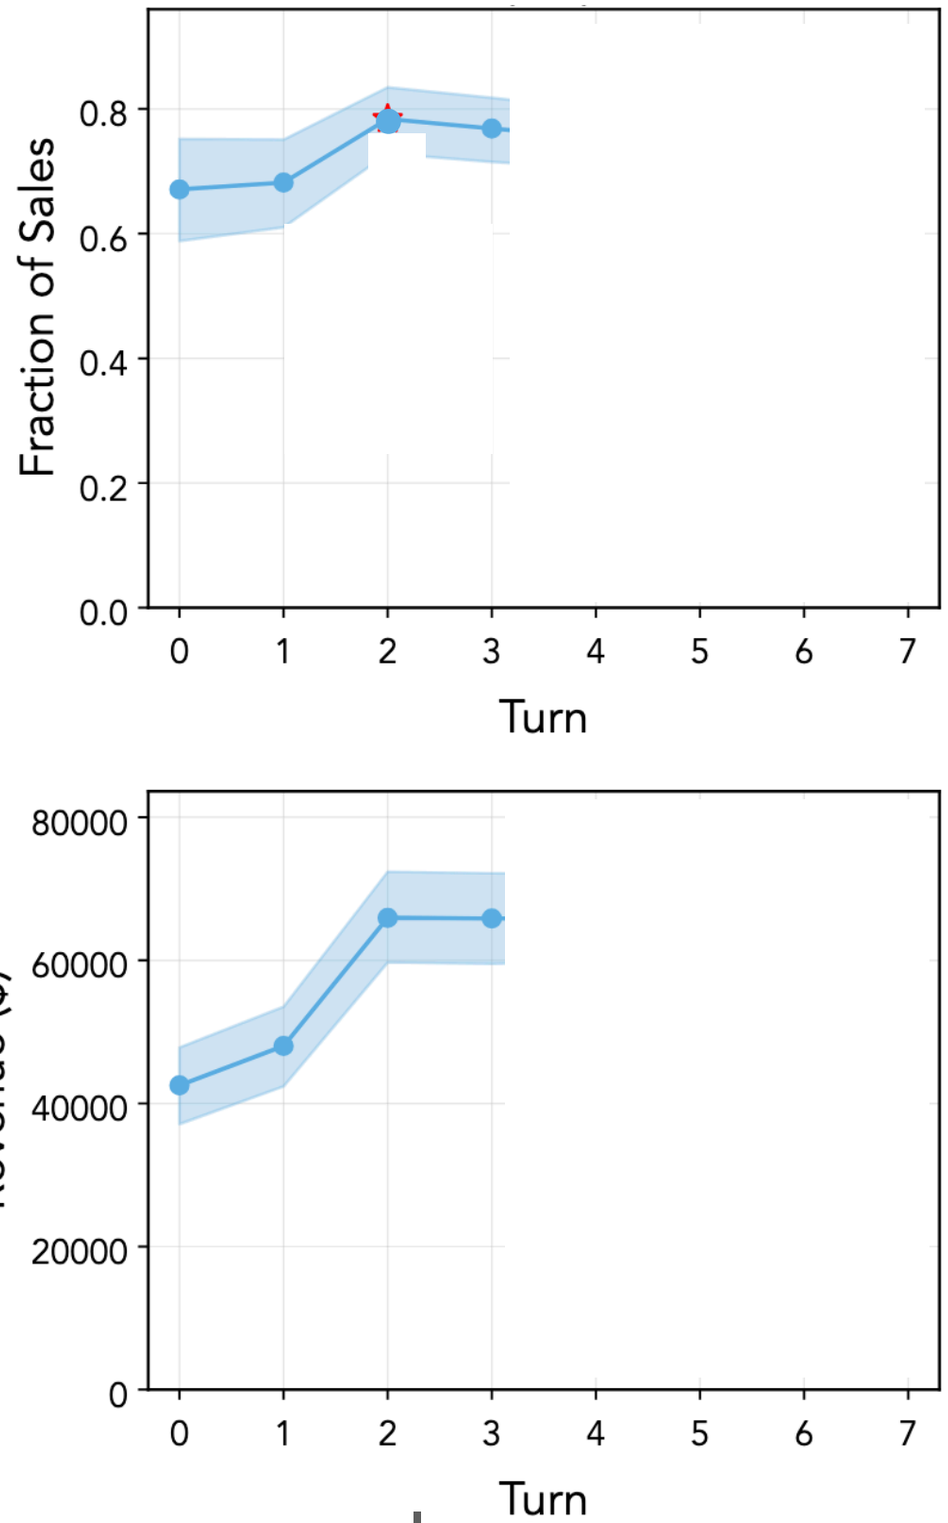
laptops

Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

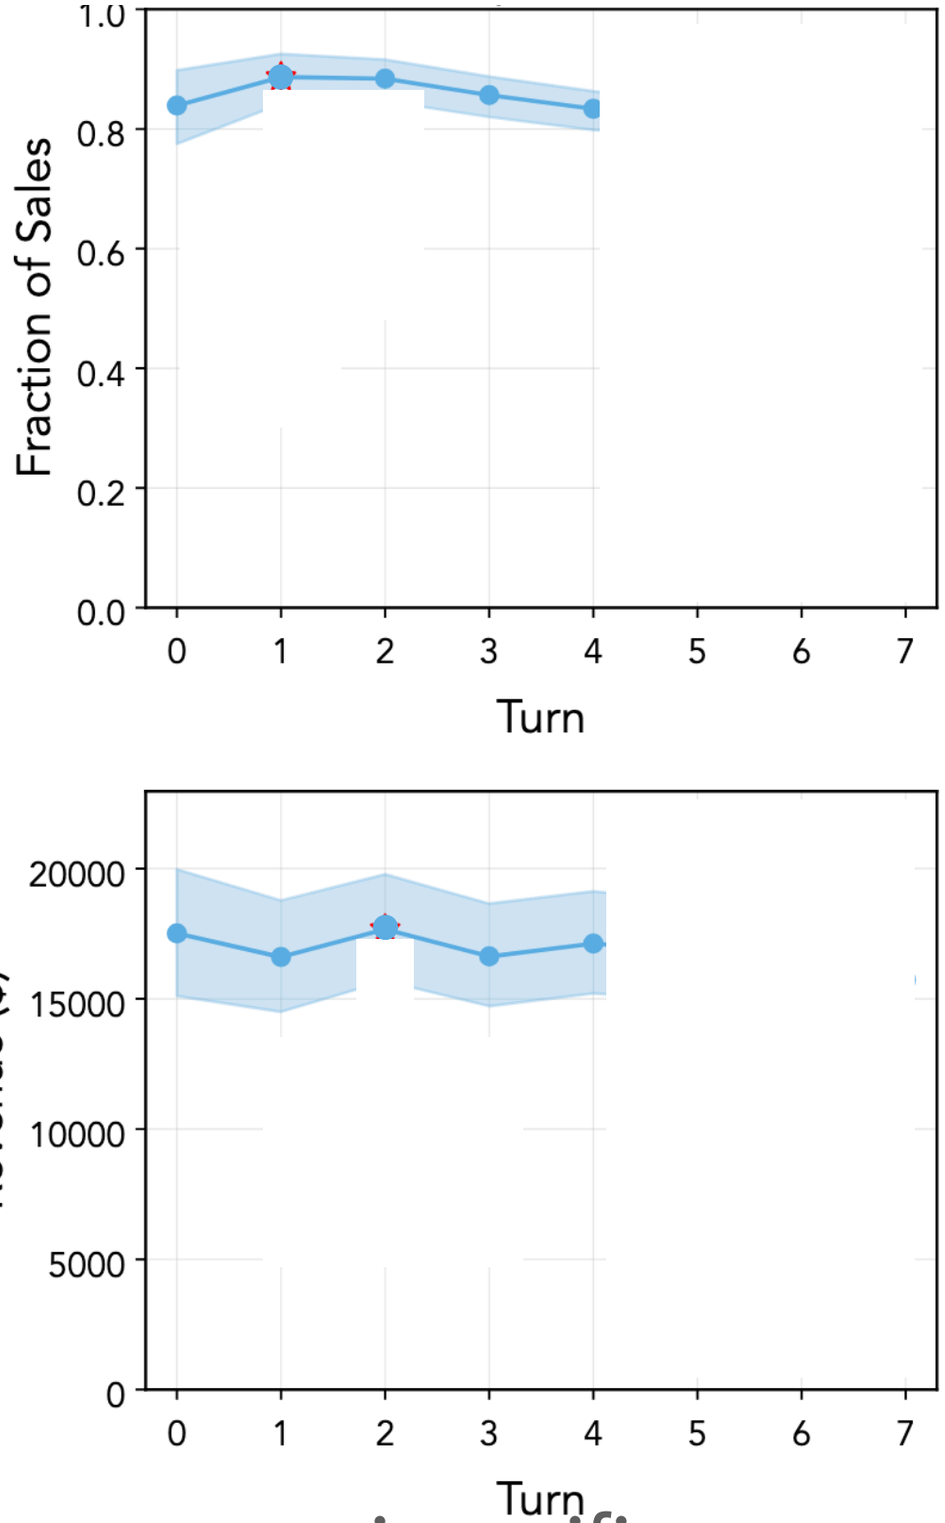
air purifiers

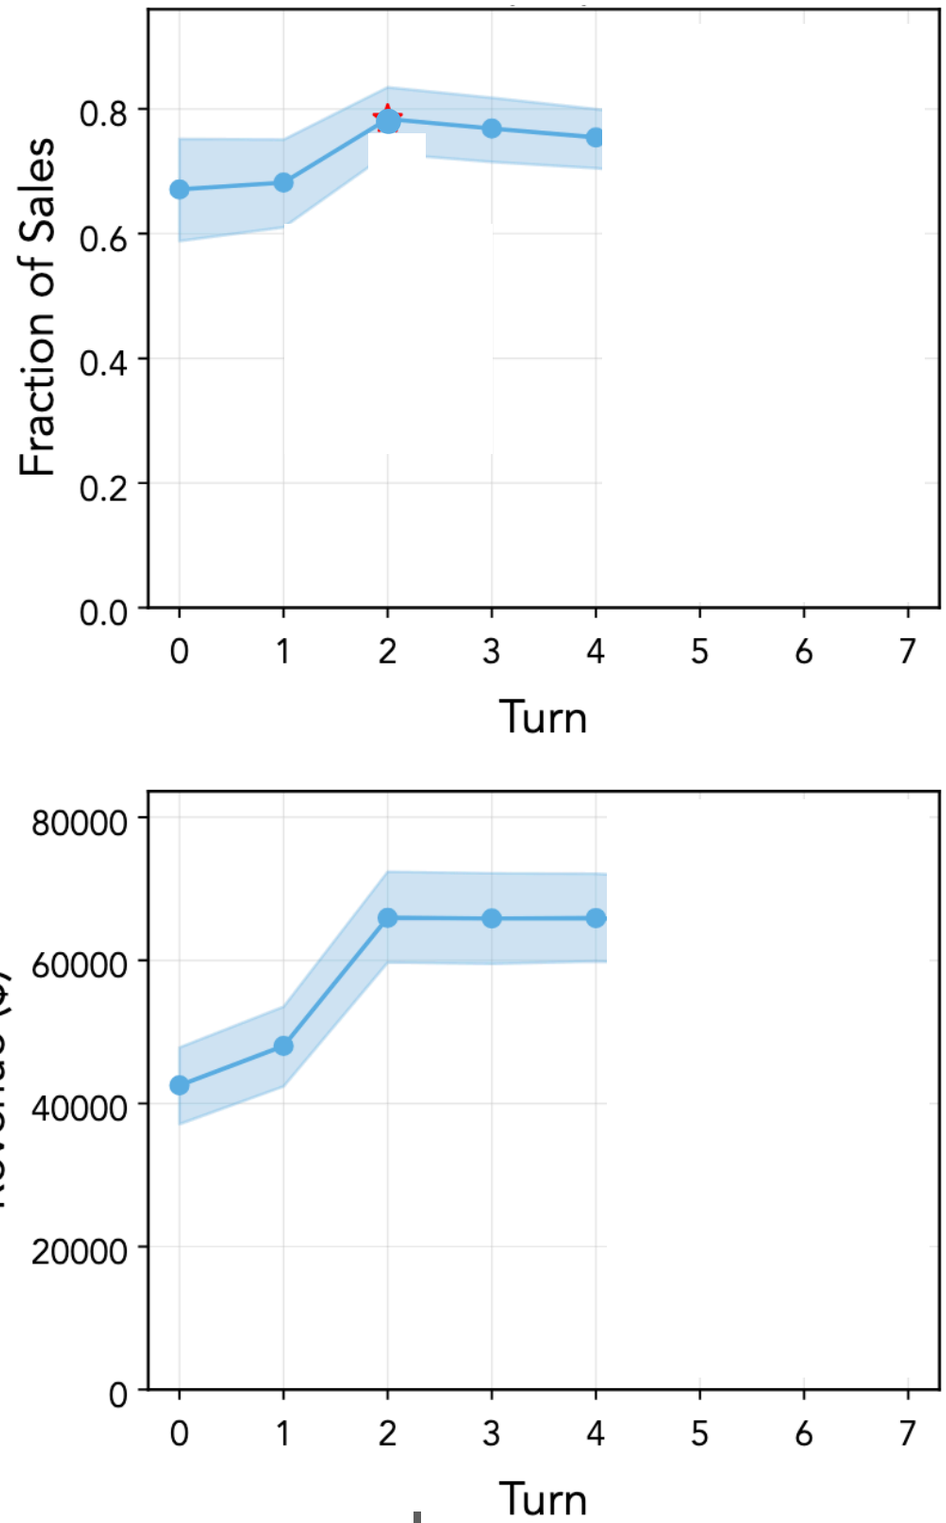
laptops

Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

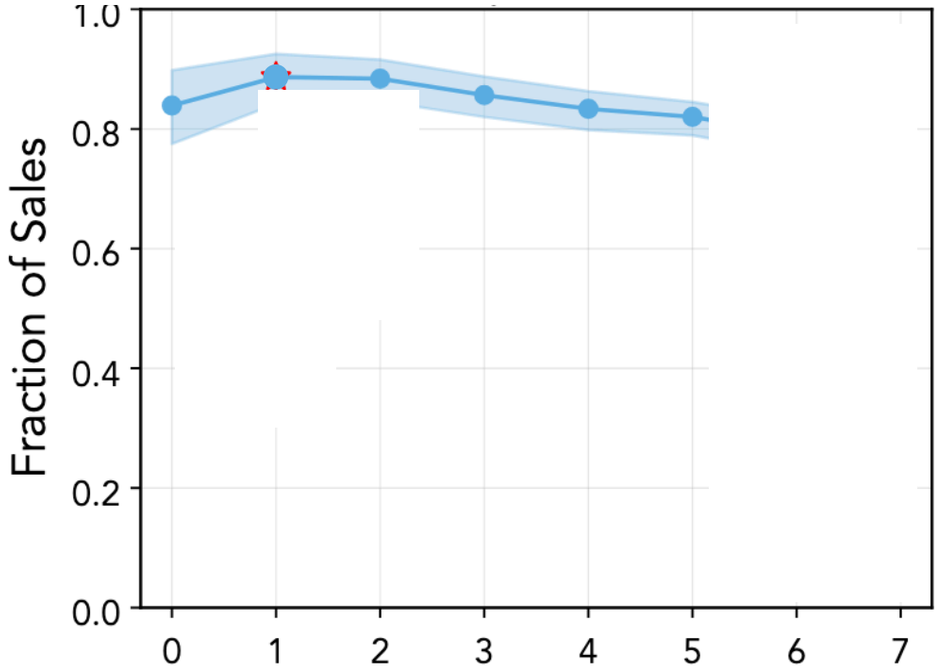
air purifiers

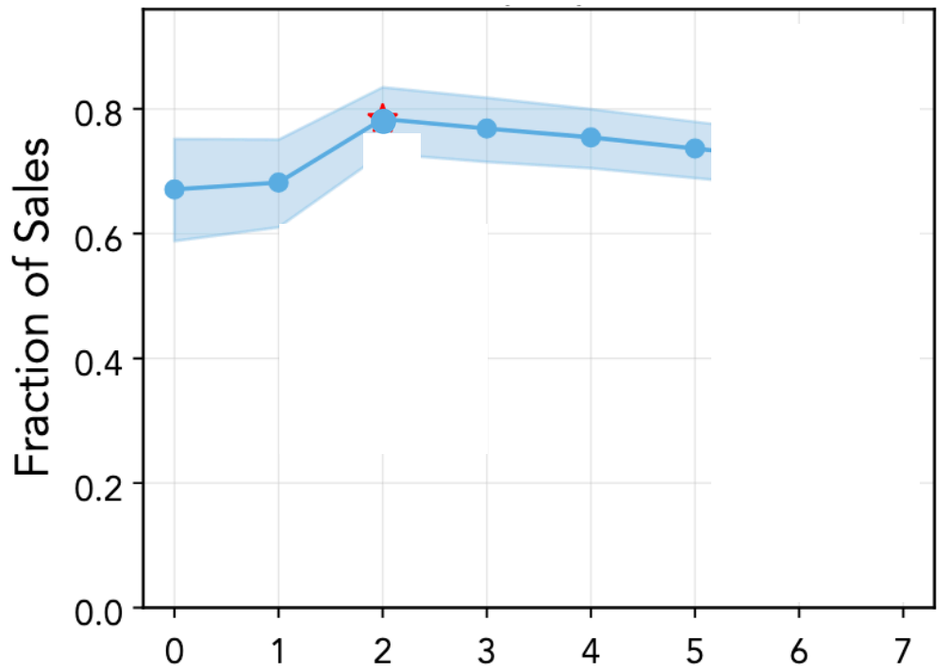
laptops

Testbed with simulated user abandonment

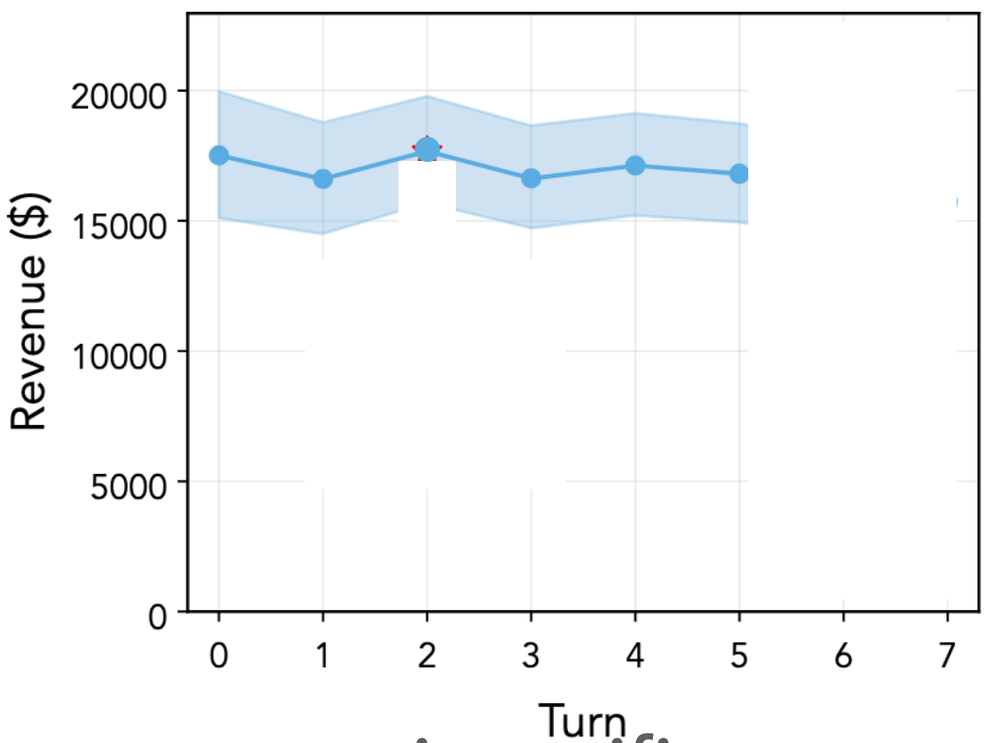
at each turn 3% of users who haven't purchased drop-out

Conversion

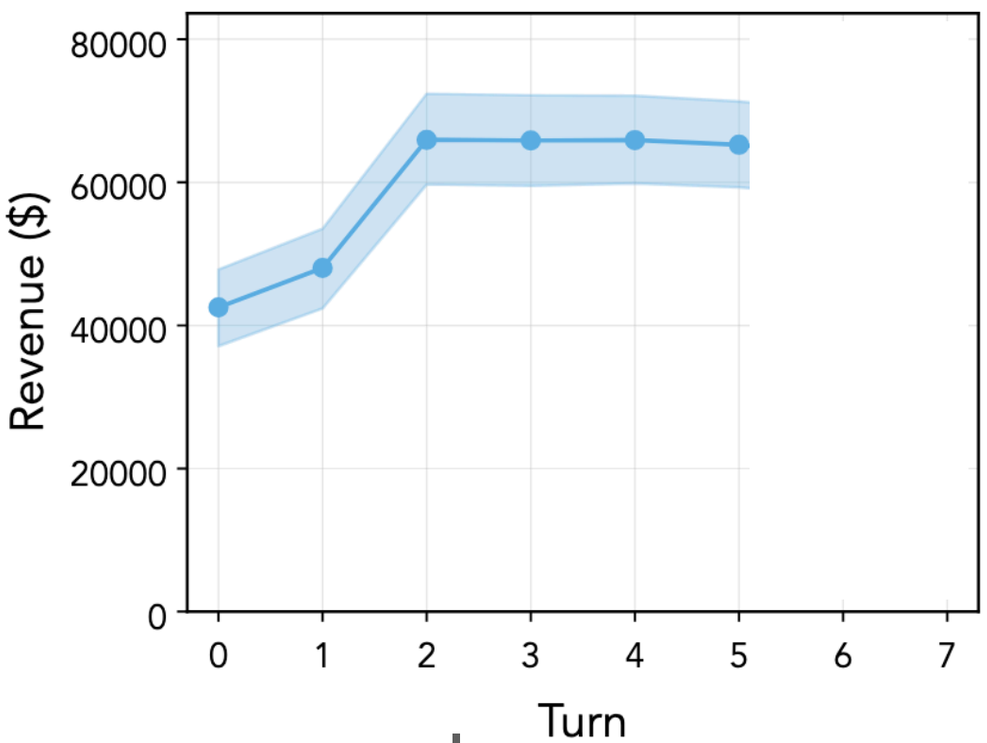




Commission



air purifiers

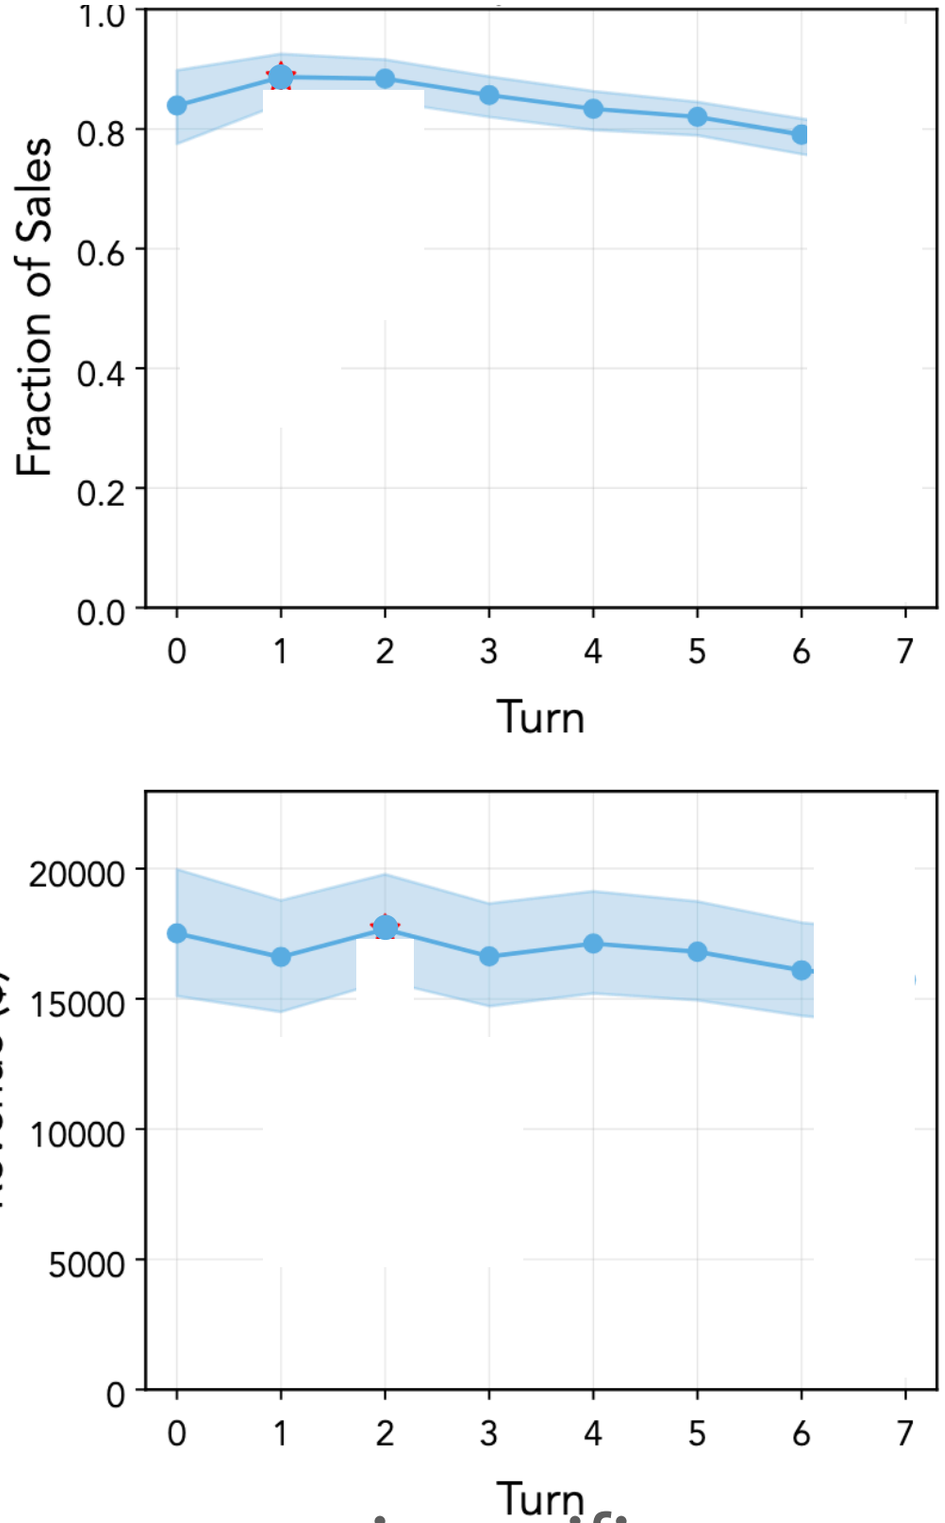


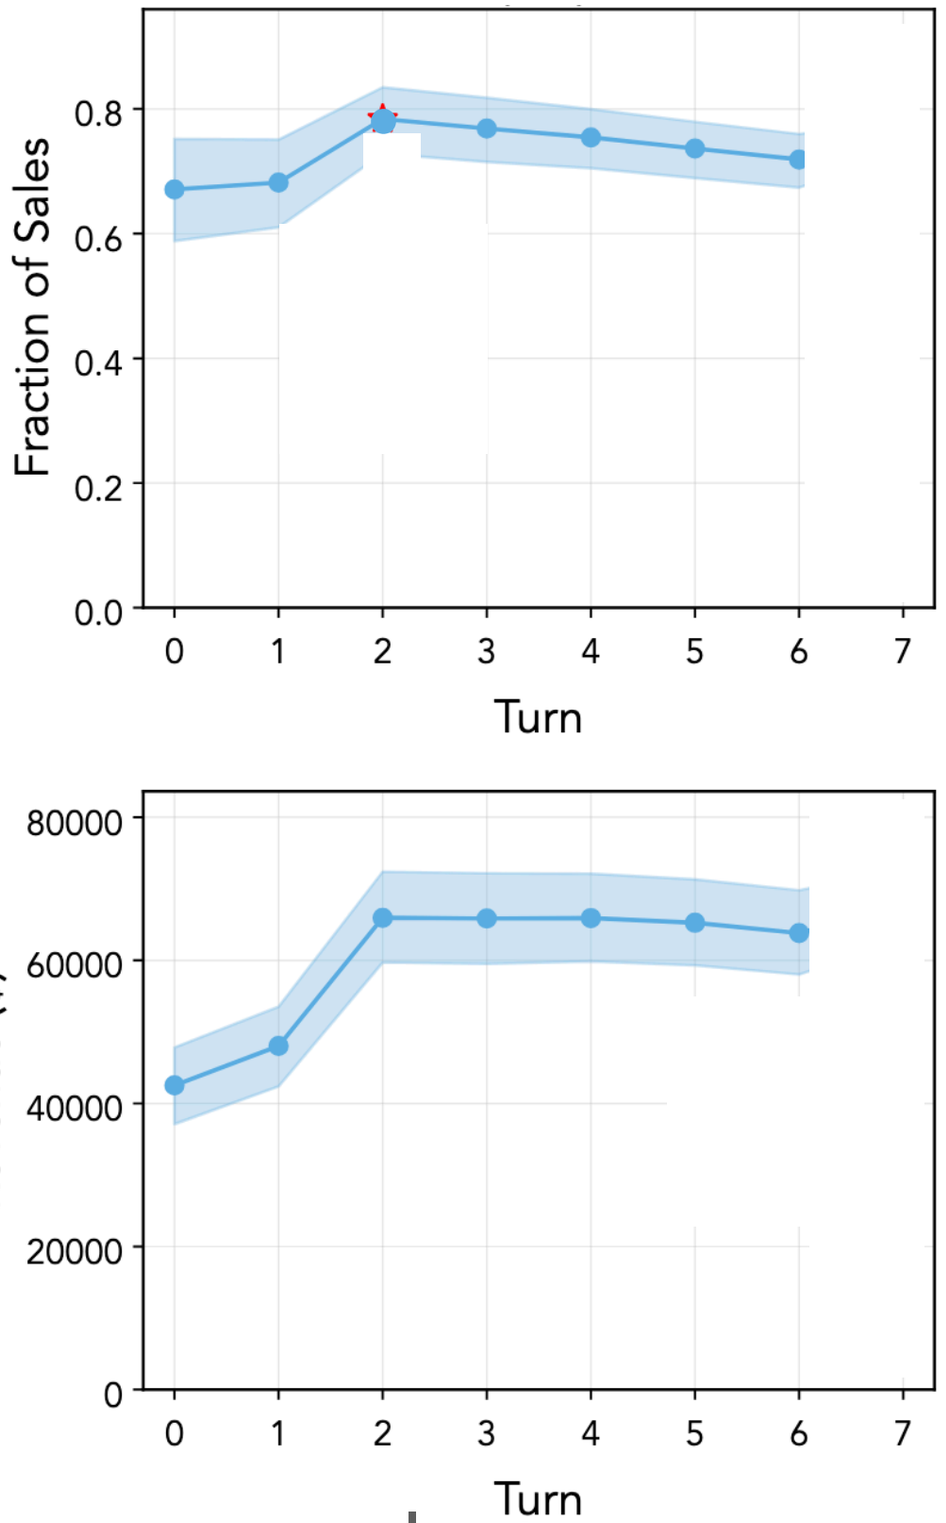
laptops

Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

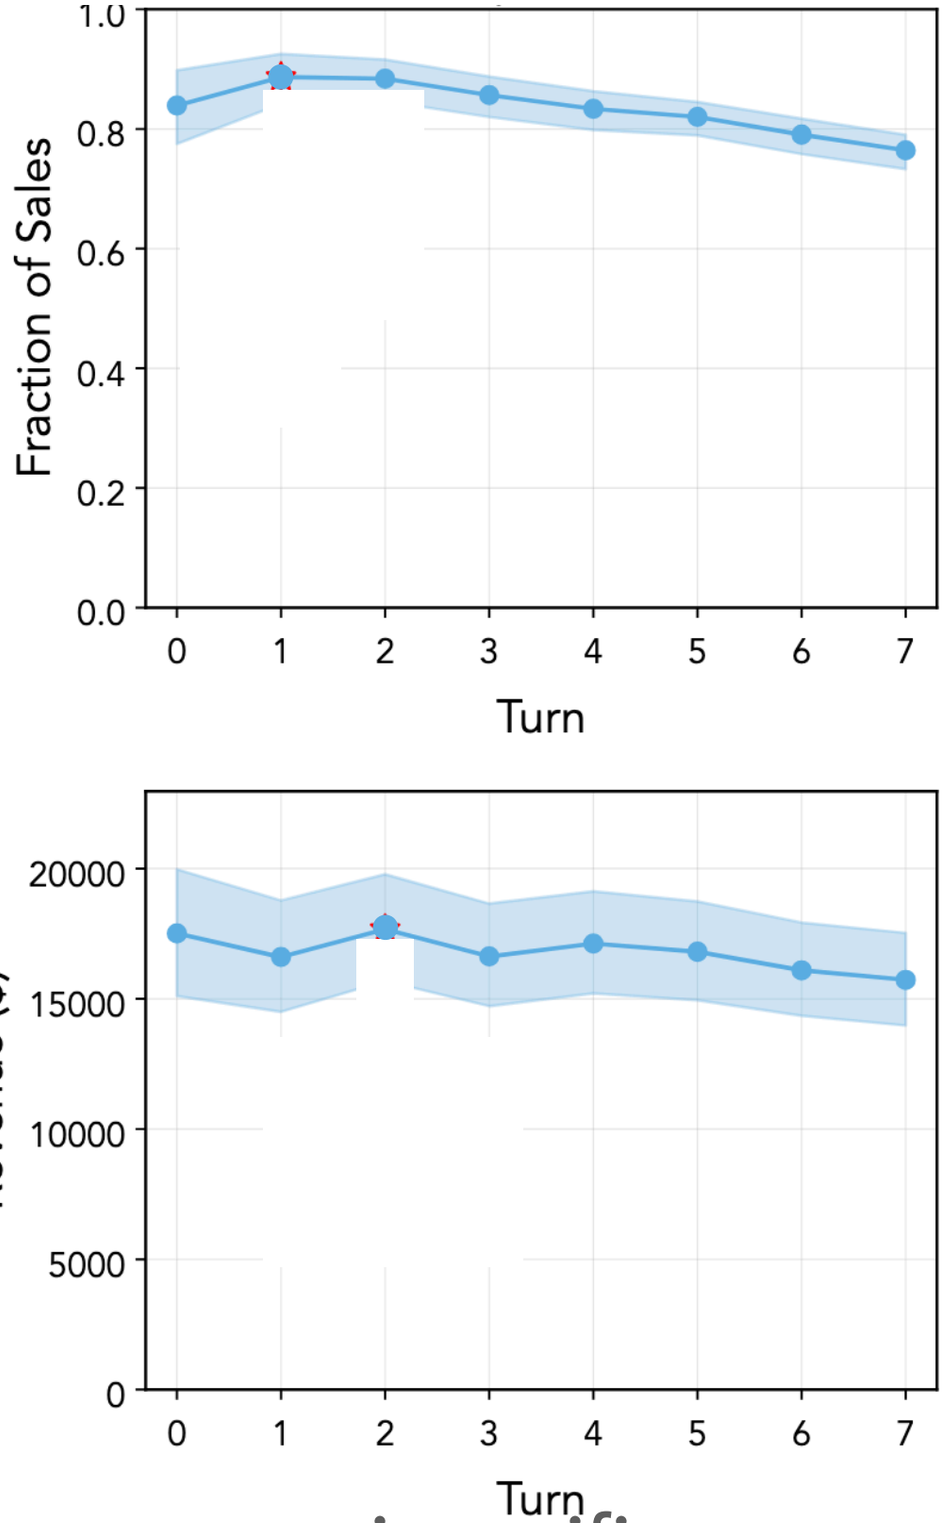
air purifiers

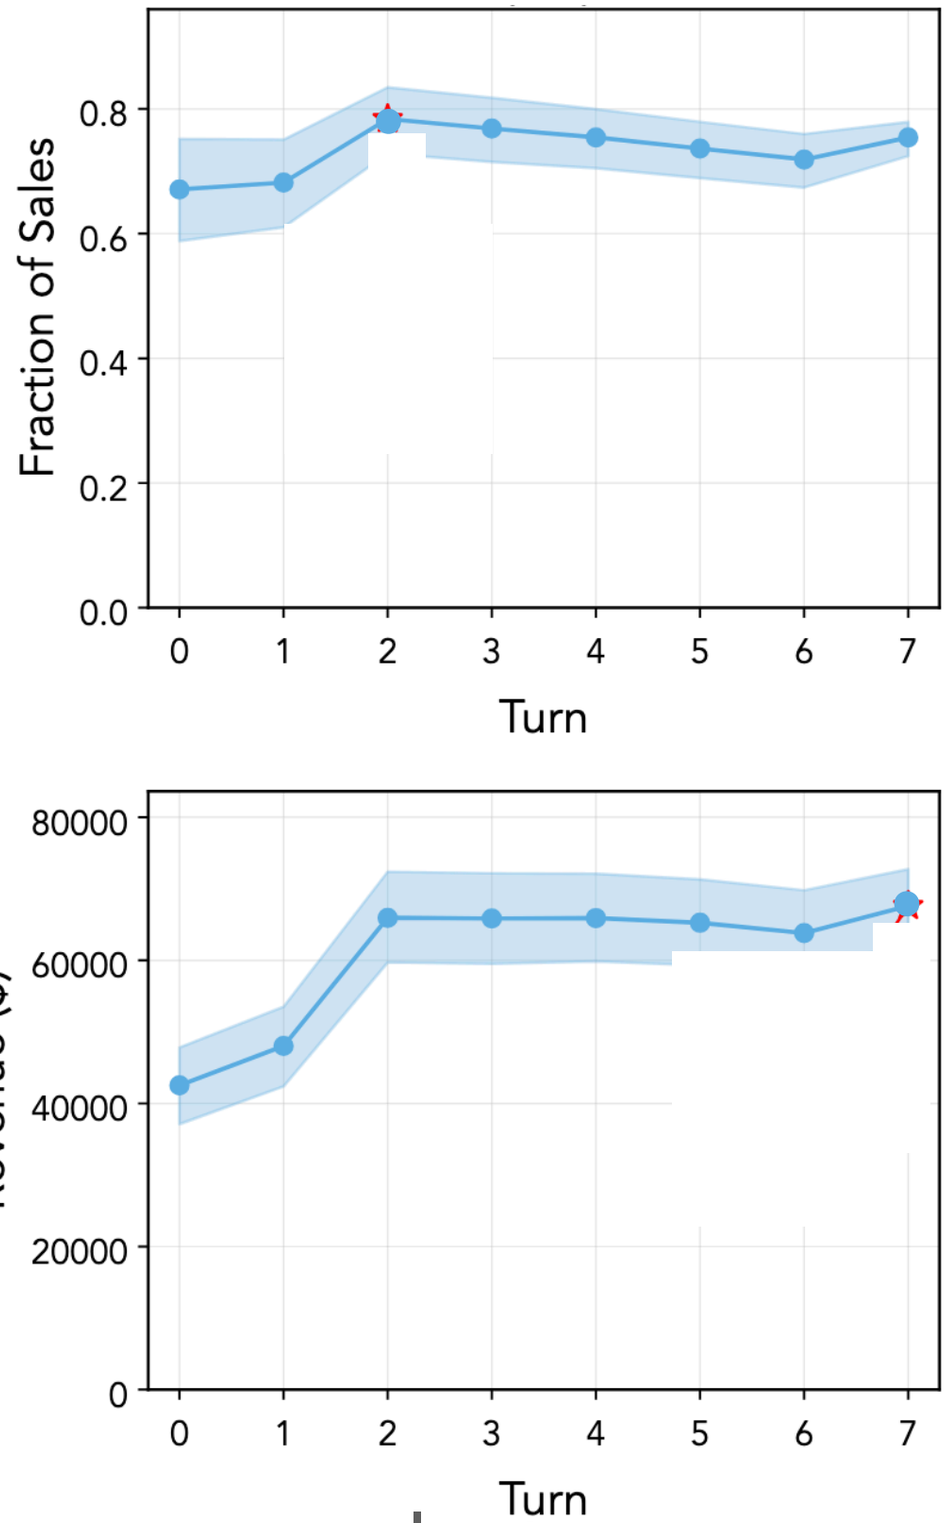
laptops

Testbed with simulated user abandonment

at each turn 3% of users who haven't purchased drop-out

Conversion





Commission

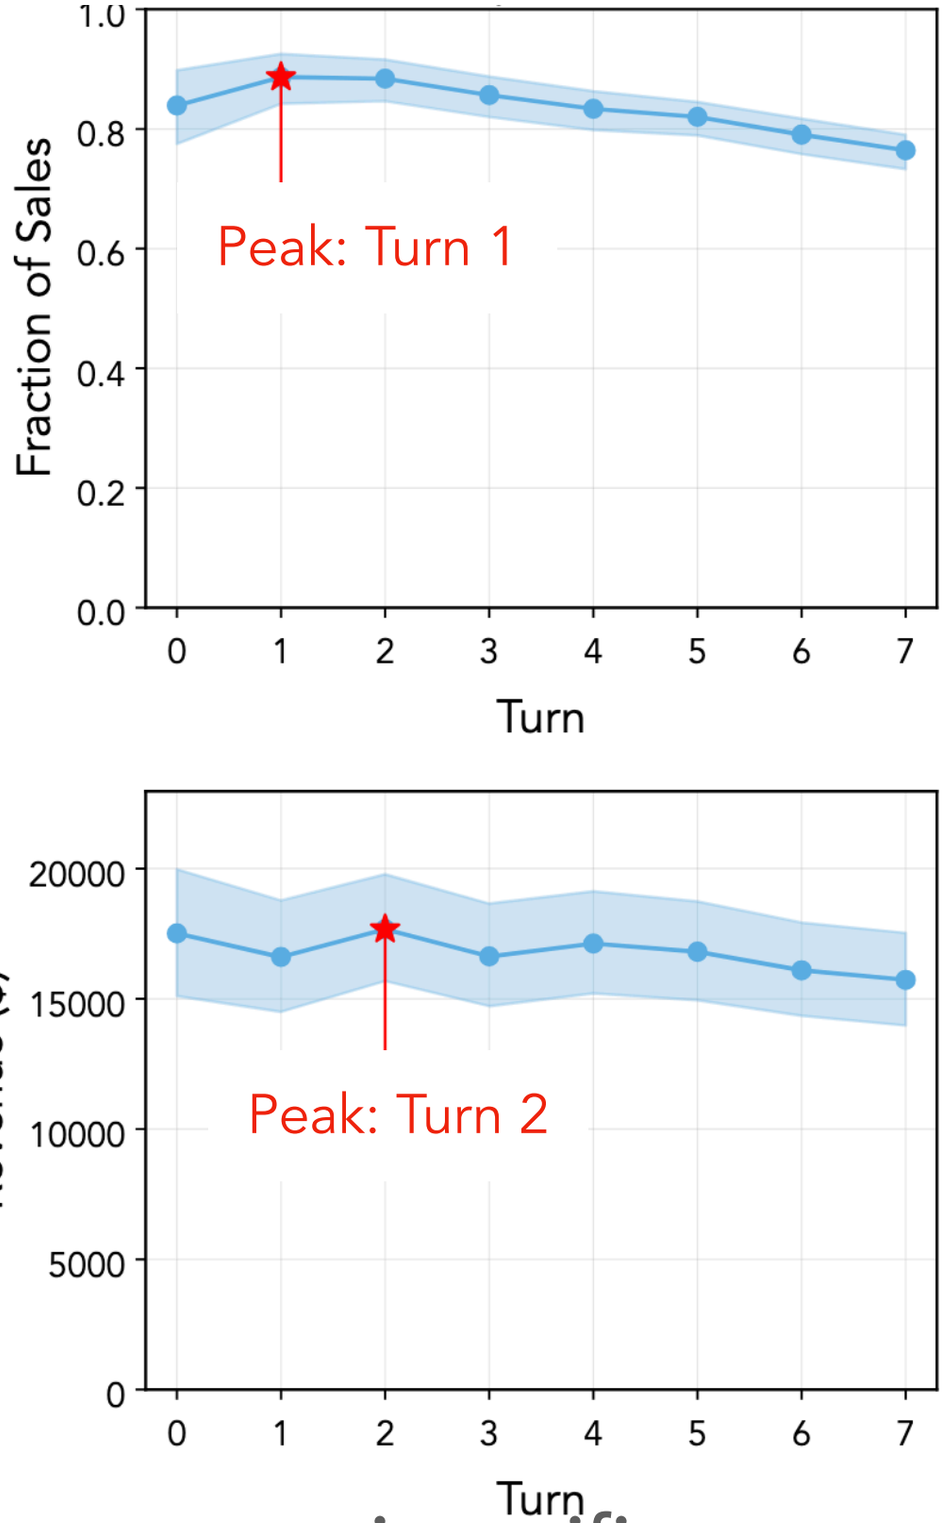
air purifiers

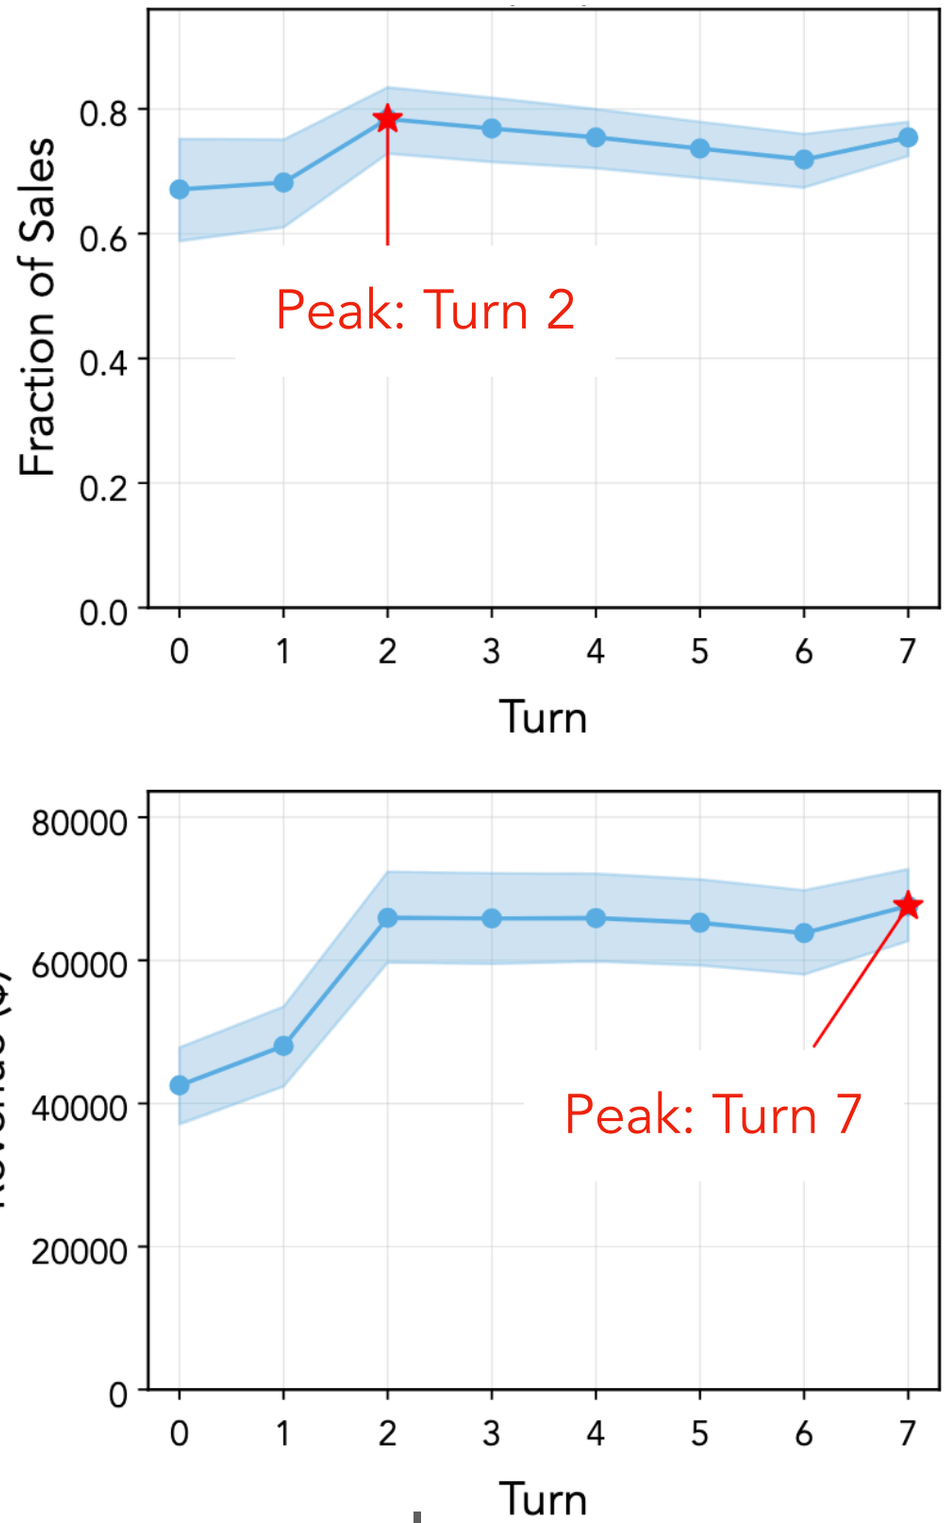
laptops

Testbed with simulated user abandonment

at each turn 4% of users who haven't purchased drop-out

Conversion





Commission

air purifiers

laptops

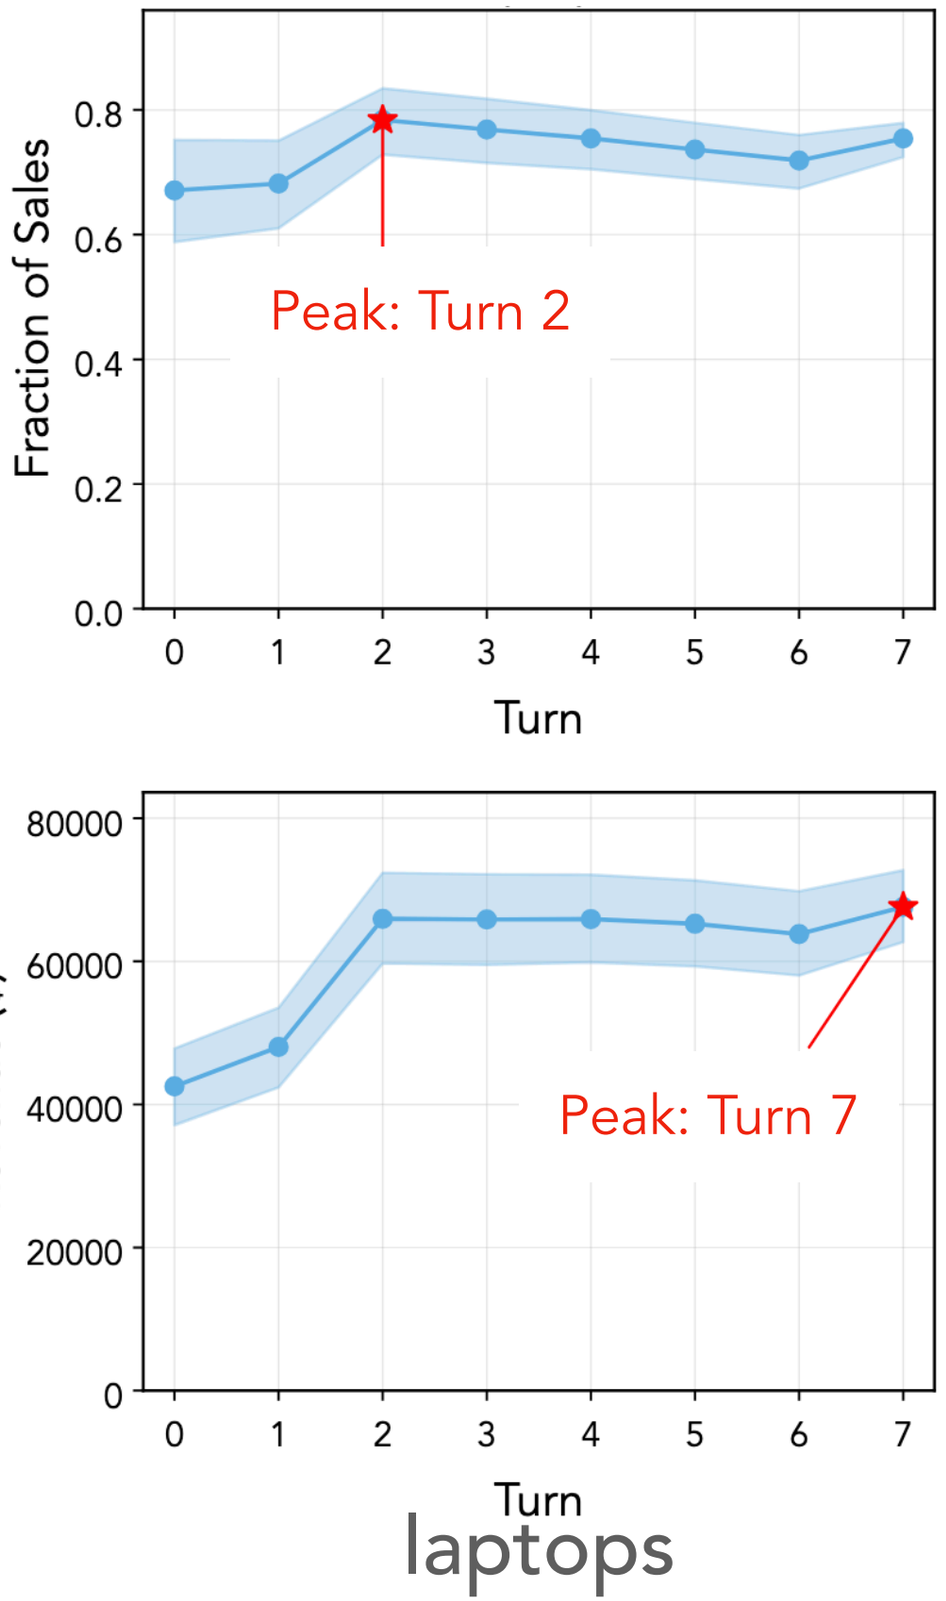
Testbed with simulated user abandonment

at each turn 4% of users who haven't purchased drop-out

Conversion

Commission





air purifiers

laptops

Testbed with simulated user abandonment

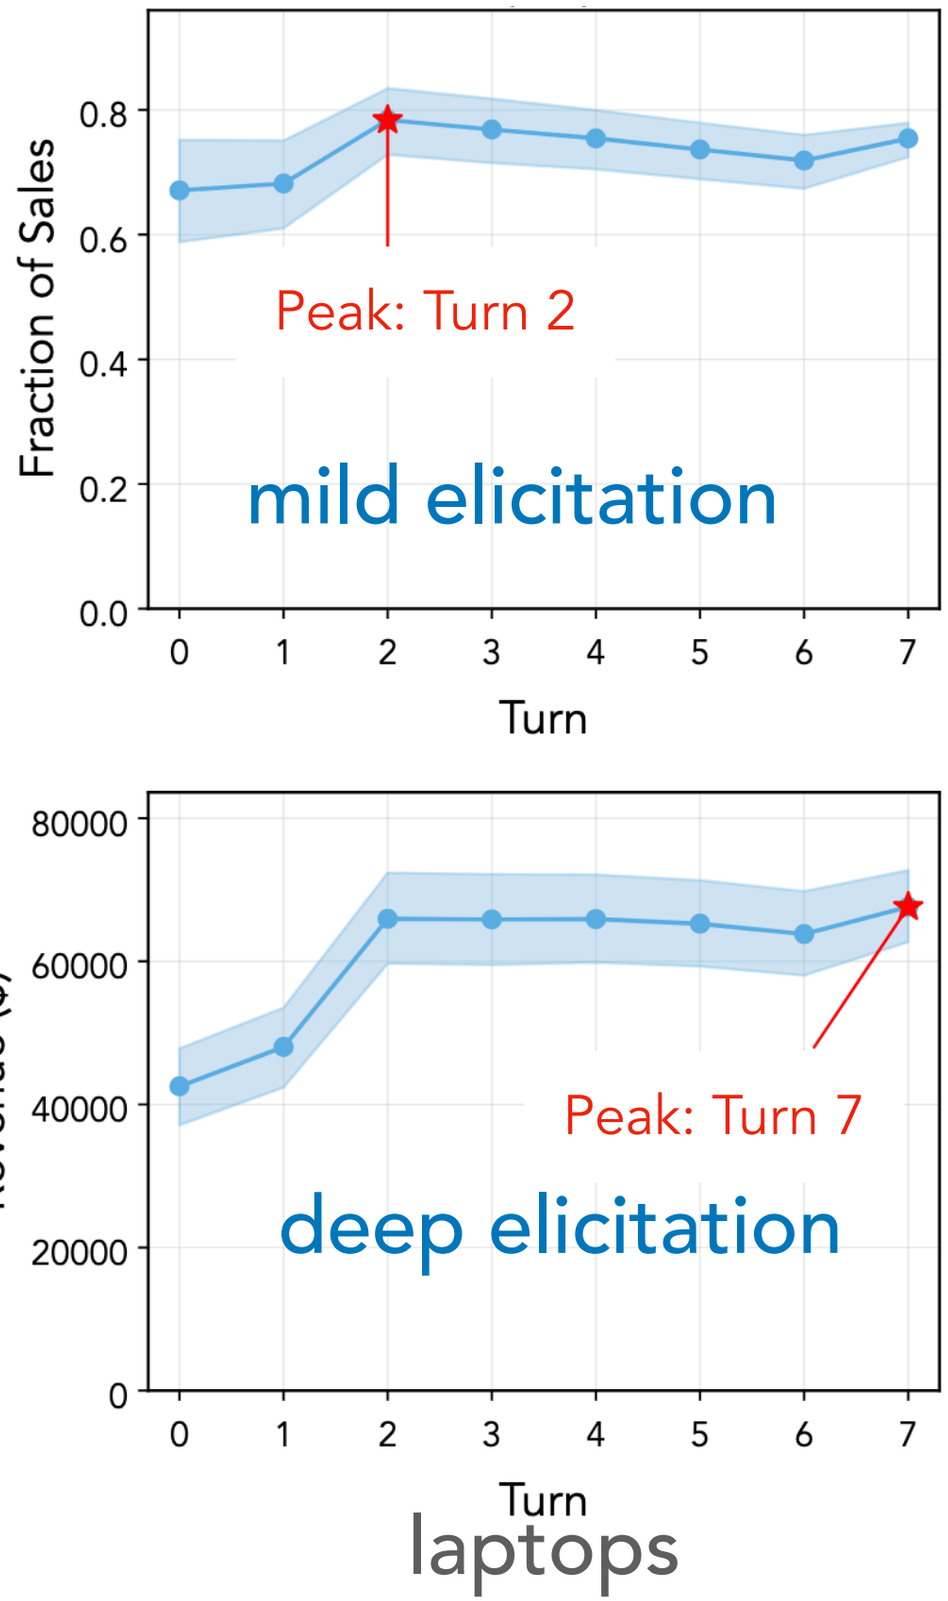
at each turn 4% of users who haven't purchased drop-out

Conversion

Commission



air purifiers



laptops

Testbed with simulated user abandonment

at each turn 4% of users who haven't purchased drop-out

Conversion

Commission

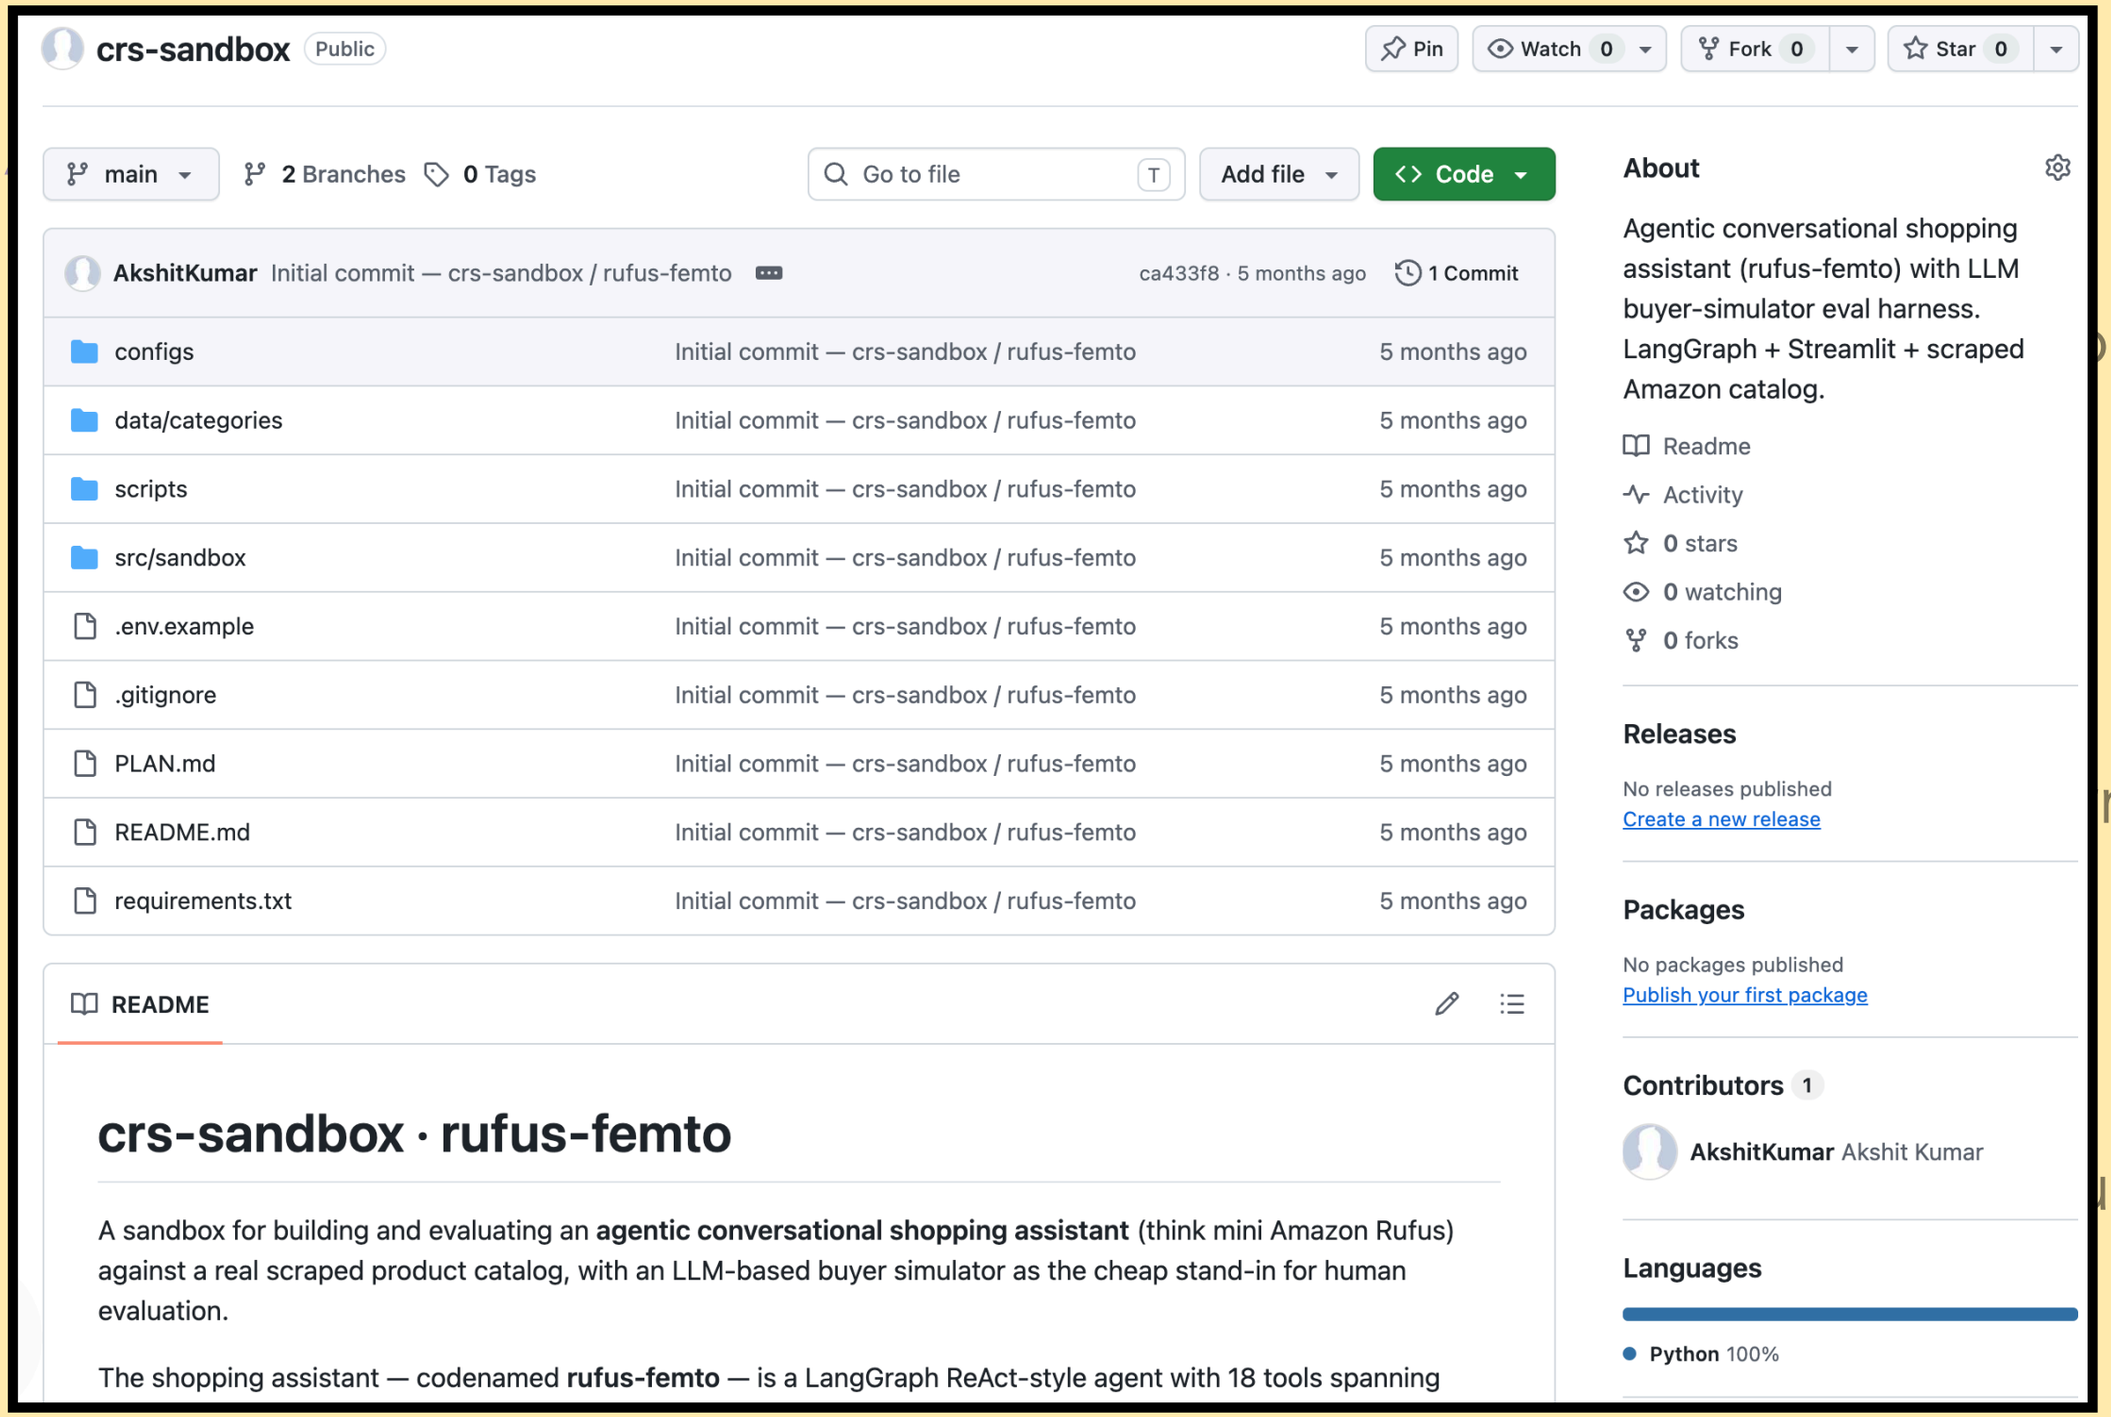




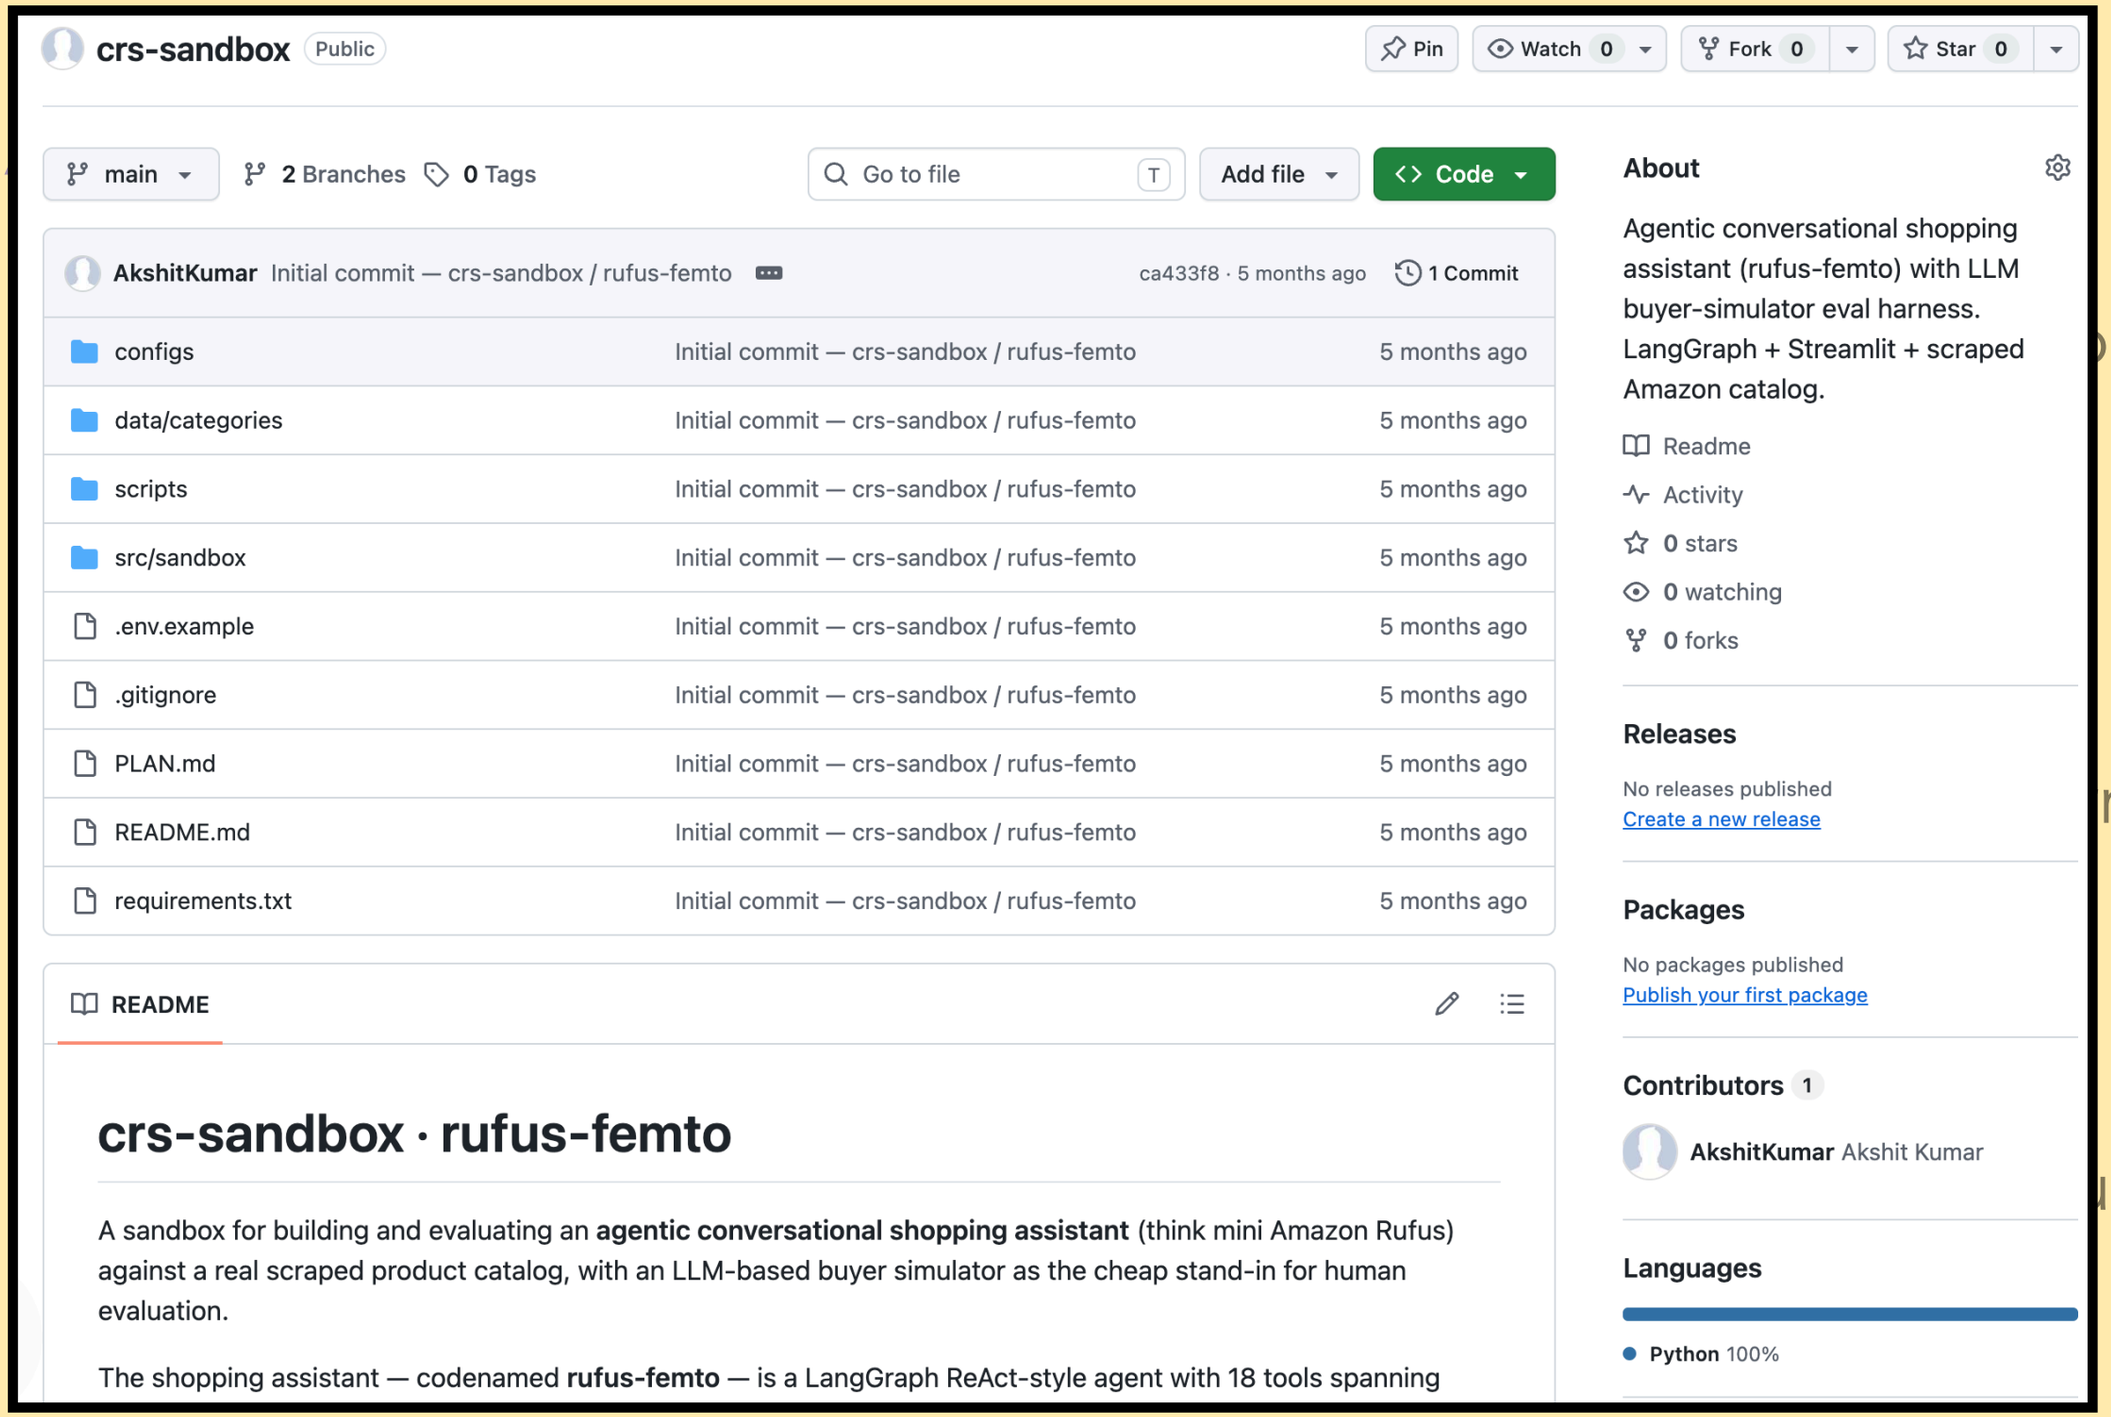
Conclusion

- A parsimonious model of conversational elicitation
  - Conversation as *costly information acquisition* — deeper questions sharpen matching *but* raise abandonment and effort
  - Platform objective drives *stark* differences in how much to ask
    - This has implications for the design of search on different platforms
- LLM-based testbed for the value of elicitation
  - Dataset & Evaluation framework to measure the incremental value of conversation and how it impacts conversion and revenue.

Conclusion


crs-sandbox
Public
Pin
Watch 0
Fork 0
Star 0
main
2 Branches
0 Tags
Go to file
T
Add file
Code
AkshitKumar Initial commit — crs-sandbox / rufus-femto
ca433f8 · 5 months ago
1 Commit
configs Initial commit — crs-sandbox / rufus-femto 5 months ago
data/categories Initial commit — crs-sandbox / rufus-femto 5 months ago
scripts Initial commit — crs-sandbox / rufus-femto 5 months ago
src/sandbox Initial commit — crs-sandbox / rufus-femto 5 months ago
.env.example Initial commit — crs-sandbox / rufus-femto 5 months ago
.gitignore Initial commit — crs-sandbox / rufus-femto 5 months ago
PLAN.md Initial commit — crs-sandbox / rufus-femto 5 months ago
README.md Initial commit — crs-sandbox / rufus-femto 5 months ago
requirements.txt Initial commit — crs-sandbox / rufus-femto 5 months ago
README
crs-sandbox · rufus-femto
A sandbox for building and evaluating an agentic conversational shopping assistant (think mini Amazon Rufus) against a real scraped product catalog, with an LLM-based buyer simulator as the cheap stand-in for human evaluation.
The shopping assistant — codenamed rufus-femto — is a LangGraph ReAct-style agent with 18 tools spanning
About
Agentic conversational shopping assistant (rufus-femto) with LLM buyer-simulator eval harness. LangGraph + Streamlit + scraped Amazon catalog.
Readme
Activity
0 stars
0 watching
0 forks
Releases
No releases published
Create a new release
Packages
No packages published
Publish your first package
Contributors 1
AkshitKumar Akshit Kumar
Languages
Python 100%

https://github.com/AkshitKumar/crs-sandbox



Thank You