Xiao Liu, New York University

Livestream Shopping and Dynamic Customer Interactions

Livestream Shopping and Dynamic Customer Interactions

Xiao Liu

Stern School of Business, New York University

May 2023 @MSI Wharton

Industry Background: Livestream Shopping

Taobao Live:

Total sales transactions	500 billion
Influencers	2 million
Industries	Jewelry, fashion, food
Engagement tool	Coupons

Source: http://www.chuangyejia.com/article-12355989.html

https://www.forbes.com/sites/laurenhallanan/2019/03/15/amazon-live-is-alibabas-live-streaming-without-the-good-bits/#4512f5bf94ab

Industry Background: Livestream Shopping

Livestream Shopping Platforms

	Taobao	Kuaishou	TikTok	Amazon	Wayfair	Facebook	Google
	淘宝直播 LIVE		TikTok	amazon live	wayfair	facebook LIVE	Google Shoploop
Start	2016	2018	2018	2019	2019	2018	2020
Source	Taobao	Taobao	Taobao	Amazon	Wayfair	Multiple	Multiple
DAU	30 mil	100 mil	100 mil	~	~	2 bil	~
Product	Apparel Cosmetics Jewelry	Deals	Cosmetics	Deals Fashion Beauty	Furniture	Used goods	Cosmetics

(1) http://pg.jrj.com.cn/acc/Res/CN_RES/INDUS/2019/7/31/8b2ce355-3ae6-4913-9983-3dcb74f af04c.pdf; (2) https://www.retaildive.com/news/wayfair-to-launch-livestreaming-service-on-w ay-day/552240/ https://99firms.com/blog/f acebook-live-stats/#gref; (3) https://www.theverge.com/2018/12/6/18129201/f acebook-live-shopping-mode-test; (4) https://fashionista.com/2017/07/shop-shops-chinese-app

Industry Background: Livestream Shopping with Discount Coupons

Dynamic Targeting: Who? How Much Discount? What Sequence?

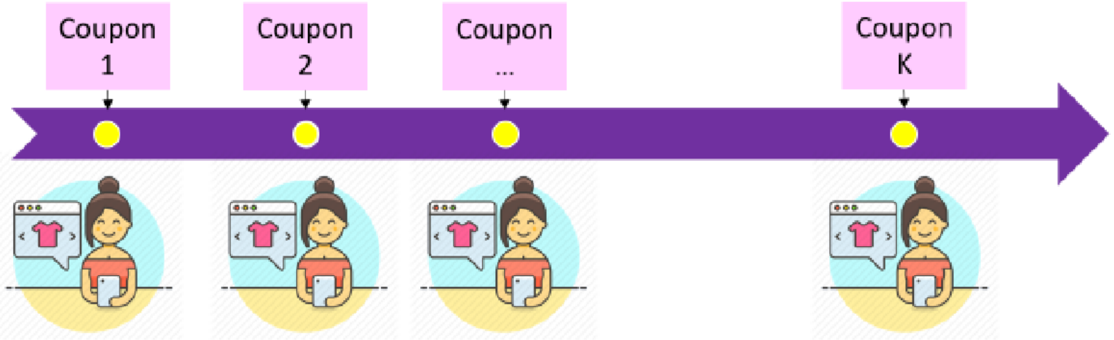

Examples of Dynamic Personalized Coupons

Executive Summary of Questions and Findings

Research Questions

1. How can we develop a theoretical framework that incorporates the intertemporal tradeoffs in dynamic personalized pricing and designs a policy that maximizes revenue?
2. How can we empirically evaluate the performance?
3. What are the gains and what mechanism can explain the gains?

Takeaways:

- Batch Deep Reinforcement Learning (BDRL) increases GMV by 63% in a field experiment compared to the status quo (random allocation)
- BDRL is 34% more effective than static targeting policies, because it incorporates intertemporal tradeoffs, especially the reference price effect
- BDRL is 20% more effective than an econometrics model, because it mitigates model bias
- Recommend a small discount for more attractive hosts and price skimming over time

Executive Summary of Questions and Findings

Research Questions

1. How can we develop a theoretical framework that incorporates the intertemporal tradeoffs in dynamic personalized pricing and designs a policy that maximizes revenue?
2. How can we empirically evaluate the performance?
3. What are the gains and what mechanism can explain the gains?

Takeaways:

- Batch Deep Reinforcement Learning (BDRL) increases GMV by 63% in a field experiment compared to the status quo (random allocation)
- BDRL is 34% more effective than static targeting policies, because it incorporates intertemporal tradeoffs, especially the reference price effect
- BDRL is 20% more effective than an econometrics model, because it mitigates model bias
- Recommend a small discount for more attractive hosts and price skimming over time

Data

- Consumers: 1 million
- Livestreams: 200.6 K
- Hosts: 11.9 K
- Coupon receive incidence: 25.9 M
- Time: 3 months

Sequence of Events

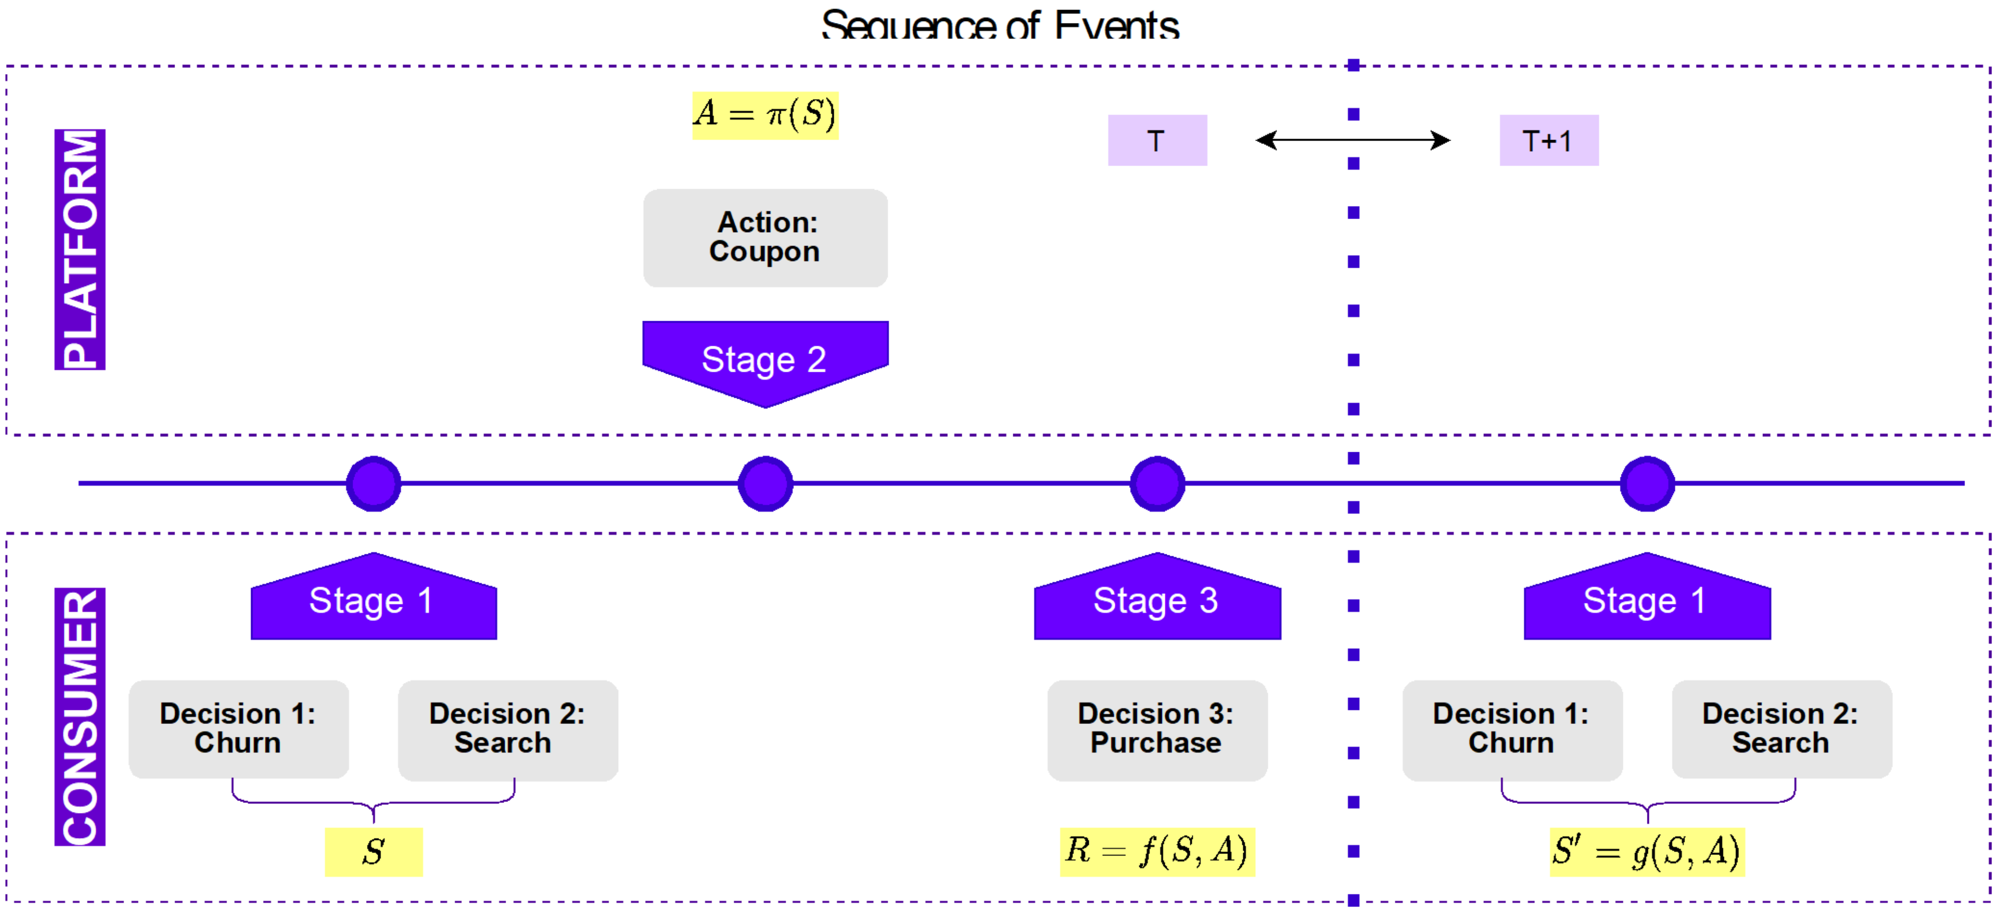

RLFramework

Solution Architecture

	Pricing Strategy	Problem
1	Uniform (Same Price For Everyone)	Ignore heterogeneity
2	Personalized	Ignore intertemporal tradeoffs
3	Personalized + Dynamic + Model-based	Model bias
4	Personalized + Dynamic + Model-free	

Theory of Intertemporal Tradeoffs

1. Reference Price
   - Big discount incentivizes purchase now
   - But lowers consumers' reference price
   - Reduces response to unfavorable coupons tomorrow
   - Optimal sequence: small discount $\Rightarrow$ big discount (Price skimming)
2. State Dependence (Positive) Loyalty
   - Small discount to increase profit margin now
   - But consumers are not locked in (b/c of switching cost)
   - Lower repeated purchase tomorrow
   - Optimal sequence: big discount $\Rightarrow$ small discount (Price penetration)
3. State Dependence (Negative) Variety-seeking
   - Big discount incentivizes purchase now
   - But consumers would like to seek variety
   - Lower purchase tomorrow
   - Optimal sequence: small discount $\Rightarrow$ small discount

Theory of Intertemporal Tradeoffs

1. Reference Price
   - Big discount incentivizes purchase now
   - But lowers consumers' reference price
   - Reduces response to unfavorable coupons tomorrow
   - Optimal sequence: small discount $\Rightarrow$ big discount (Price skimming)
2. State Dependence (Positive) Loyalty
   - Small discount to increase profit margin now
   - But consumers are not locked in (b/c of switching cost)
   - Lower repeated purchase tomorrow
   - Optimal sequence: big discount $\Rightarrow$ small discount (Price penetration)
3. State Dependence (Negative) Variety-seeking
   - Big discount incentivizes purchase now
   - But consumers would like to seek variety
   - Lower purchase tomorrow
   - Optimal sequence: small discount $\Rightarrow$ small discount

Theory of Intertemporal Tradeoffs

Reference Price

- Big discount incentivizes purchase now
- But lowers consumers' reference price
- Reduces response to unfavorable coupons tomorrow
- Optimal sequence: small discount $\Rightarrow$ big discount (Price skimming)

State Dependence (Positive) Loyalty

- Small discount to increase profit margin now
- But consumers are not locked in (b/c of switching cost)
- Lower repeated purchase tomorrow
- Optimal sequence: big discount $\Rightarrow$ small discount (Price penetration)

State Dependence (Negative) Variety-seeking

- Big discount incentivizes purchase now
- But consumers would like to seek variety
- Lower purchase tomorrow
- Optimal sequence: small discount $\Rightarrow$ small discount

Solution 4: Model-free Dynamic Personalized

$$\pi^* = \underset{\pi=\{A_{it}\}_{i,t}}{\operatorname{argmax}} E\left[\sum_{i=1}^{I}\sum_{t=0}^{T_i}\delta^t R_{it}(S_{it}, A_{it})\right]$$

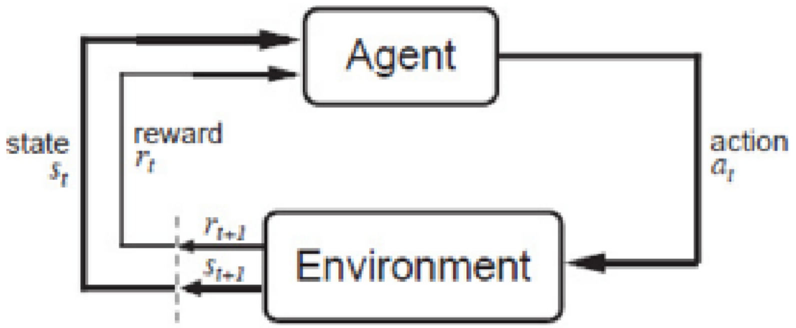

Deep Reinforcement Learning

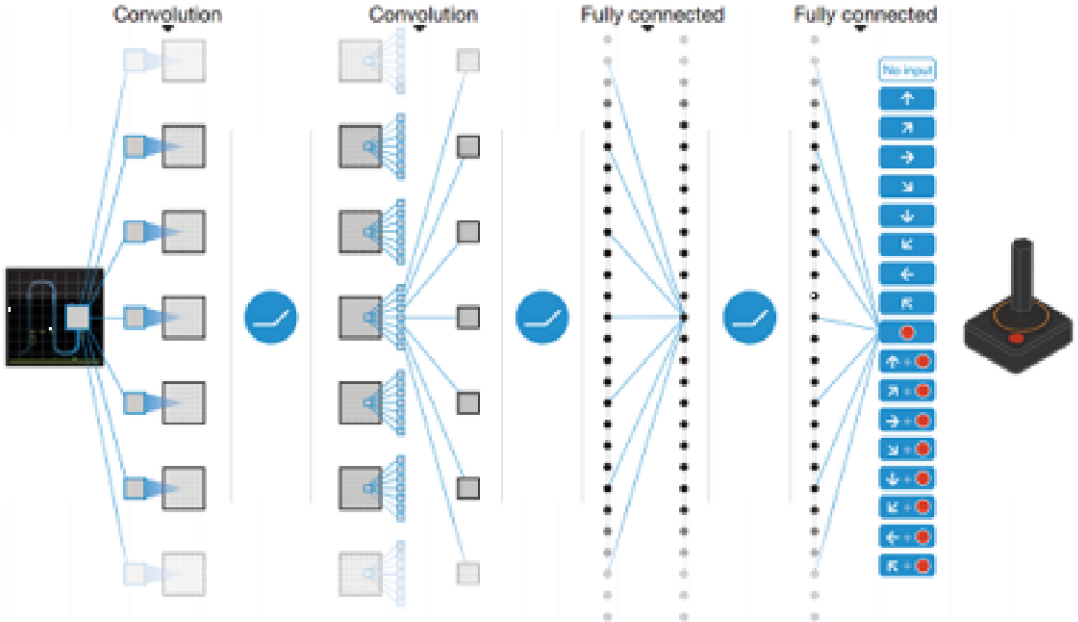

Source: https://adeshpande3.github.io/Deep-Learning-Research-Review-Week-2-Reinforcement-Learning

Model Setup: Reinforcement Learning Framework

Reinforcement Learning Framework

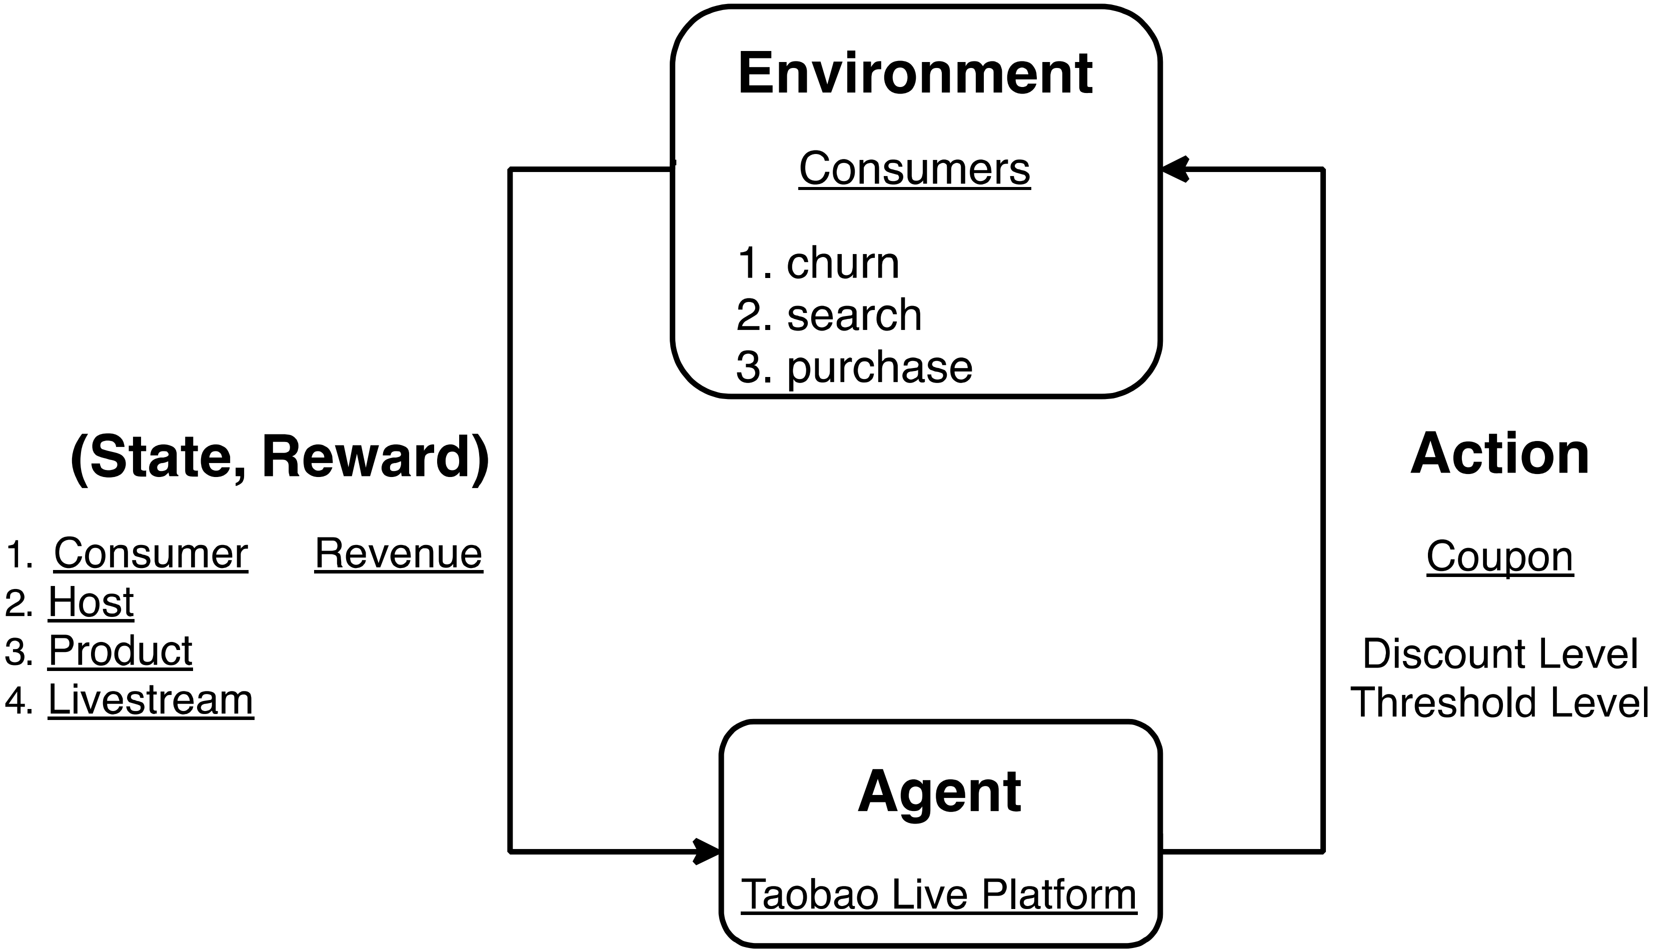

State

Categorization of State Variables

	Consumer	Host	Product	Livestream
Static*	Demographics, Product preference	Demographics, attractiveness	Category, price, penetration, retention	Visit, engagement (comment, like, etc.)
Dynamic	[Recency, frequency, monetary value] of Coupons received $\times$ Coupon category $\times$ Product category; Churn	Frequency of sellers visited	[Recency, frequency, monetary value] of Products purchased $\times$ Product category	[Recency, frequency] of webpages visited

*Static features are defined as those not affected by actions. Static features may come from a stationary distribution instead of being a fixed number.

Dynamic States Incorporate Intertemporal Tradeoffs

Intertemporal Tradeoffs and Dynamic State Variables

Intertemporal Tradeoffs	Dynamic State Variable
Reference Price	monetary_coupon
Loyalty/Inertia	frequency_product
Variety Seeking	frequency_product

Policy Learning-BCQ

Model Intuition

- Model bias? $\Rightarrow$ Model-free: $Q^{new}(\mathbf{S},A) \leftarrow (1-\alpha)\underbrace{Q(\mathbf{S},A)}_{oldvalue} + \alpha * \underbrace{(R + \delta \max_A Q(\mathbf{S}',A))}_{update})$ ▸ Model-free
- Curse of dimensionality? $\Rightarrow$ Deep learning: $Q(\mathbf{S},A) = Q(\mathbf{S},A;\theta)$
- Cost of experimentation $\Rightarrow$ Batch ▸ Batch

Benchmark

Benchmark Models

	Type	Model
1	Uniform (Same Price for Everyone)	Regression w/o interaction
2	Personalized	GBDT; Deep Neural Networks Orthogonal Random Forest
3	Personalized + Dynamic + Model-based	Econometrics
4	Personalized + Dynamic + Model-free	BDRL

Model Comparison

Model Comparison Using the Doubly Robust Estimator

	1. Static Uniform	2. Static Personalized			3. Model-based Dynamic Personalized	4. Model-free Dynamic Personalized
	A: Regression	B: GBDT	C: DNN	D: ORF	E: Econometrics	F: Proposed BDRL
GMV	6.57	7.57	6.99	7.51	8.39	9.57
Gain	+12%	+29%	+19%	+28%	+43%	+63%

An increased GMV of $13.1 million (91.8 million RMB) for the consumers in our sample

- GBDT: Gradient Boosted Decision Tree
- DNN: Deep Neural Networks
- ORF: Orthogonal Random Forest
- BDRL: Batch Deep Reinforcement Learning

Field Experiment

- 2 weeks
- 80% Random, 10% Econometrics, 10% BDRL

Field Experiment Result

	Random Allocation	Model-based Dynamic Personalized (Econometrics)	Model-free Dynamic Personalized (BDRL)
GMV	6.98	9.70	11.16
Gain	~	+39%	+60%

Policy 1:When to Target?

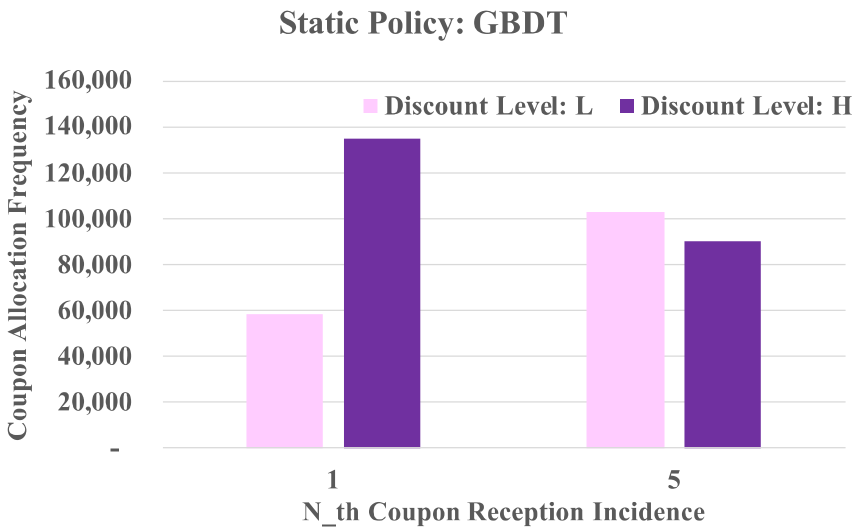

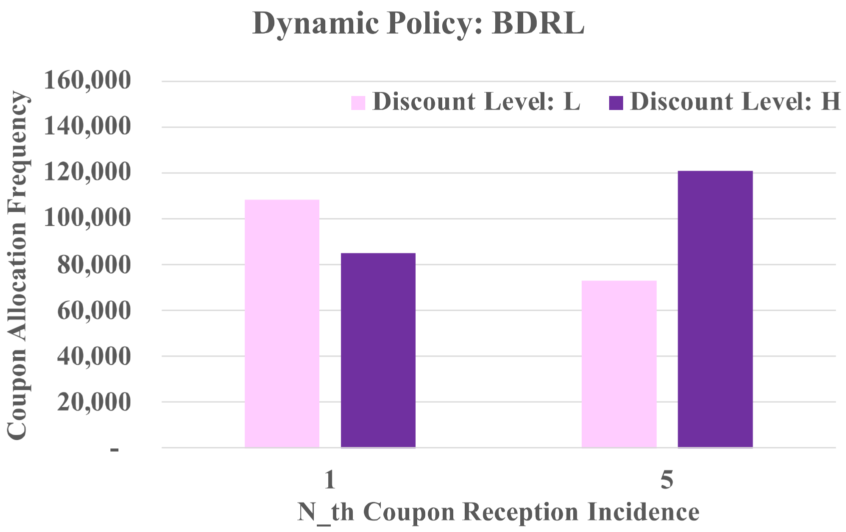

The dynamic policy gradually increases the high discount level coupons to avoid the reference price effect

Policy 2: Who to Target?

Targeting Rule Under BCQ: Host Attractiveness

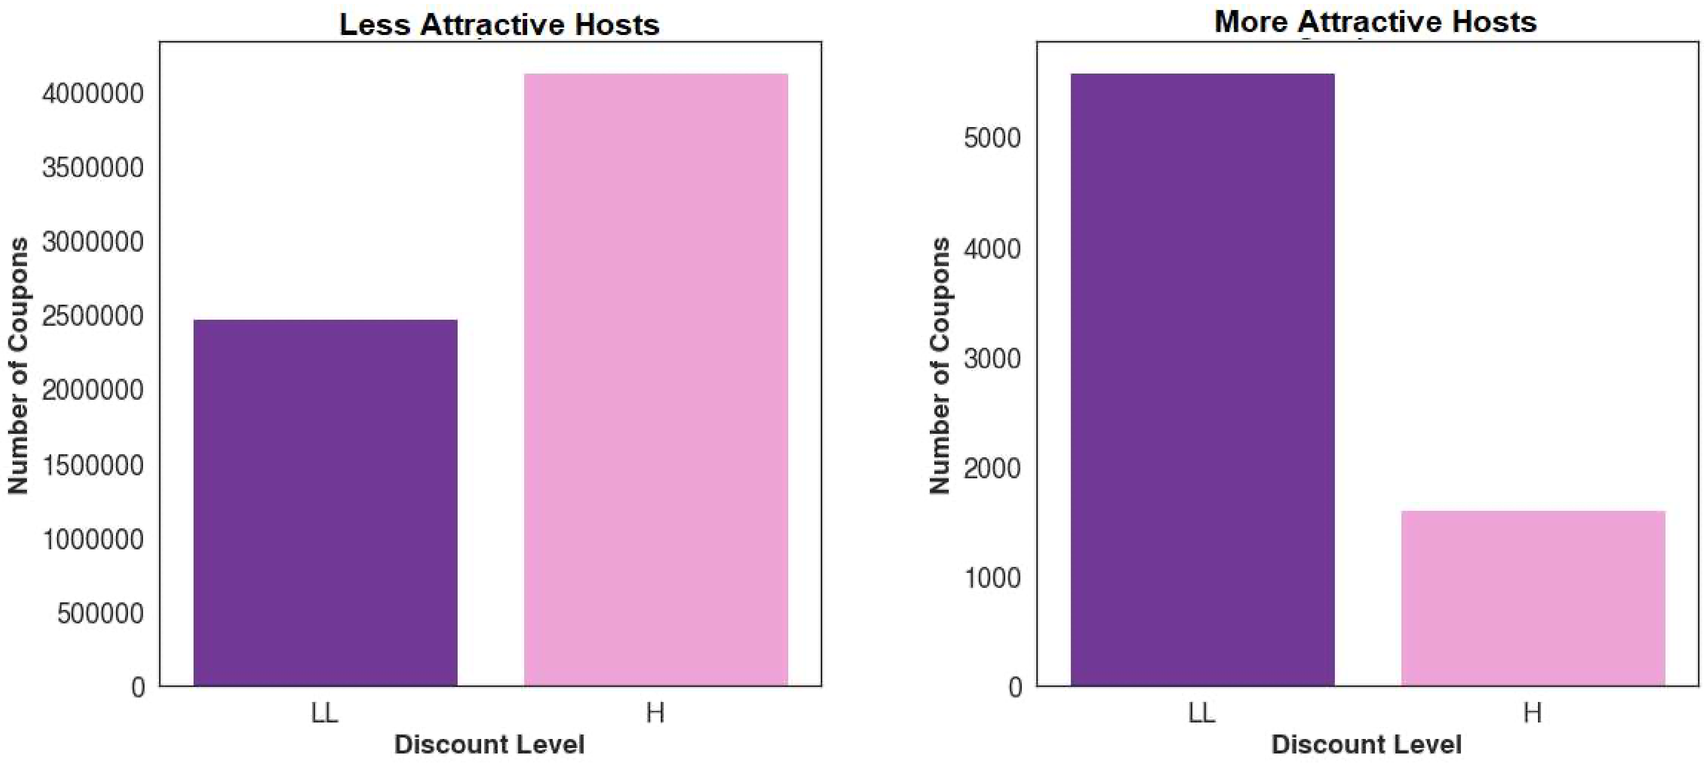

Better deals for less attractive hosts

Policy 3: When + Who to Target?

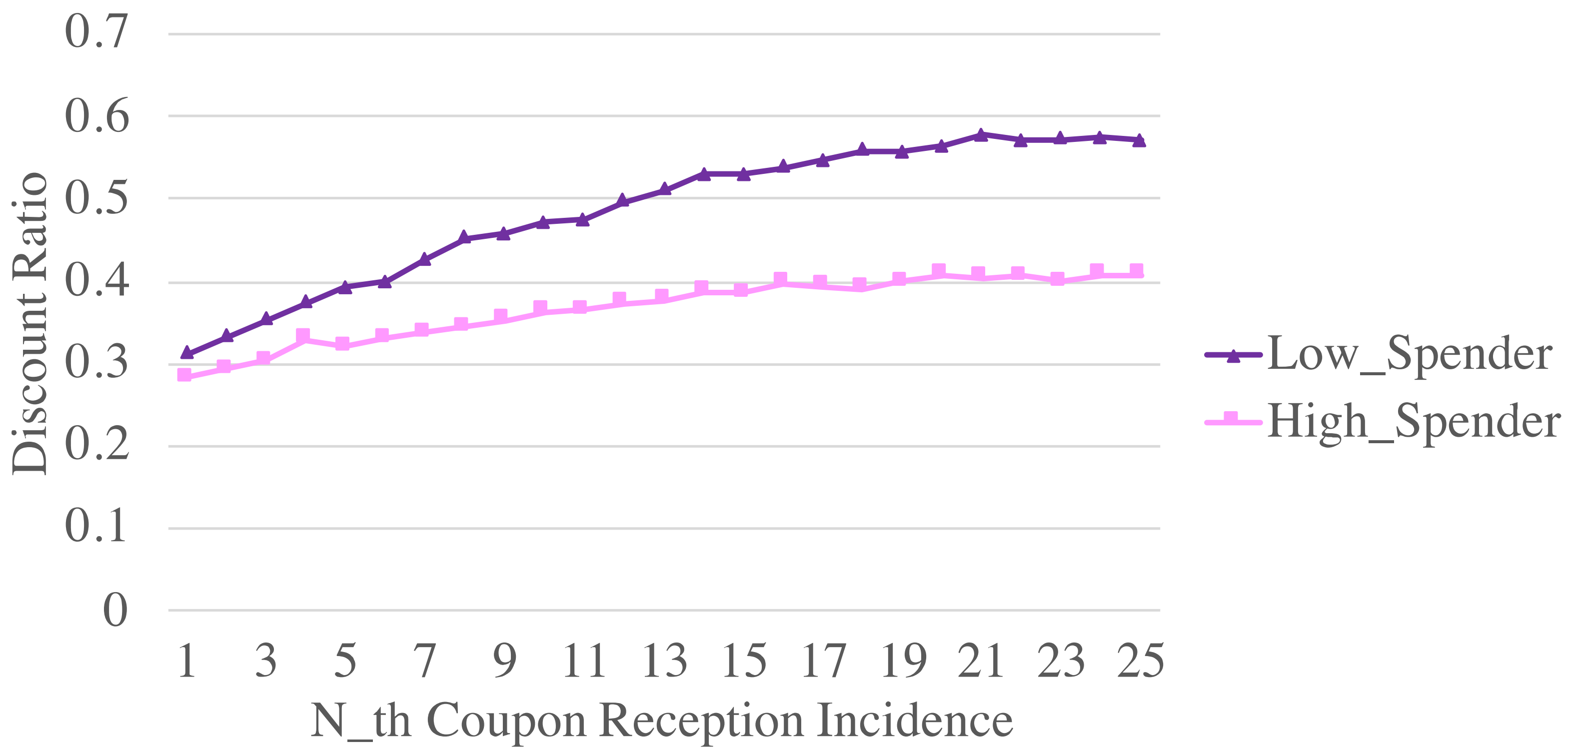

Low spenders: More price sensitive, bigger discounts and faster increase

High spenders: Less price sensitive, smaller discounts and slower increase

How Can You Apply It?

Generalizable Framework, Can Be Applied To

- Dynamic, high-frequency interventions
- Customer Lifetime Value Optimization

Model Requirements

- Individual-level Panel Data
- Real-time interaction digital infrustructure

Lunch with Speakers

Upcoming Events:

Webinar:
- **5/9** – **Ensembling Experiments to Optimize Interventions Along the Customer Journey** | Yicheng Song, University of Minnesota

Workshop:
- **5/16** – **The Customer-Base Audit** | Peter Fader, University of Pennsylvania

Workshop:
- **5/23 – Digital Customer Engagement** | Wendy Moe, University of Maryland

Book Series Webinar:
- **5/30 – Power and Prediction: The Disruptive Economics of Artificial Intelligence** | Avi Goldfarb, University of Toronto

Webinar:
- **6/27 – Regulating Privacy Online: The Economic Impact of the GDPR** | Samuel Goldberg, Stanford University

In-Person Event:
- **Fall 2023** – **MSI Accelerator** | New York, NY

Register Now at msi.org/2023-calendar-of-events/