The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets

Aaron P. Kaye

MIT Sloan & CSAIL

17th Digital Economics Conference, Toulouse School of Economics

January 9, 2025

Recommendation systems determine product rankings

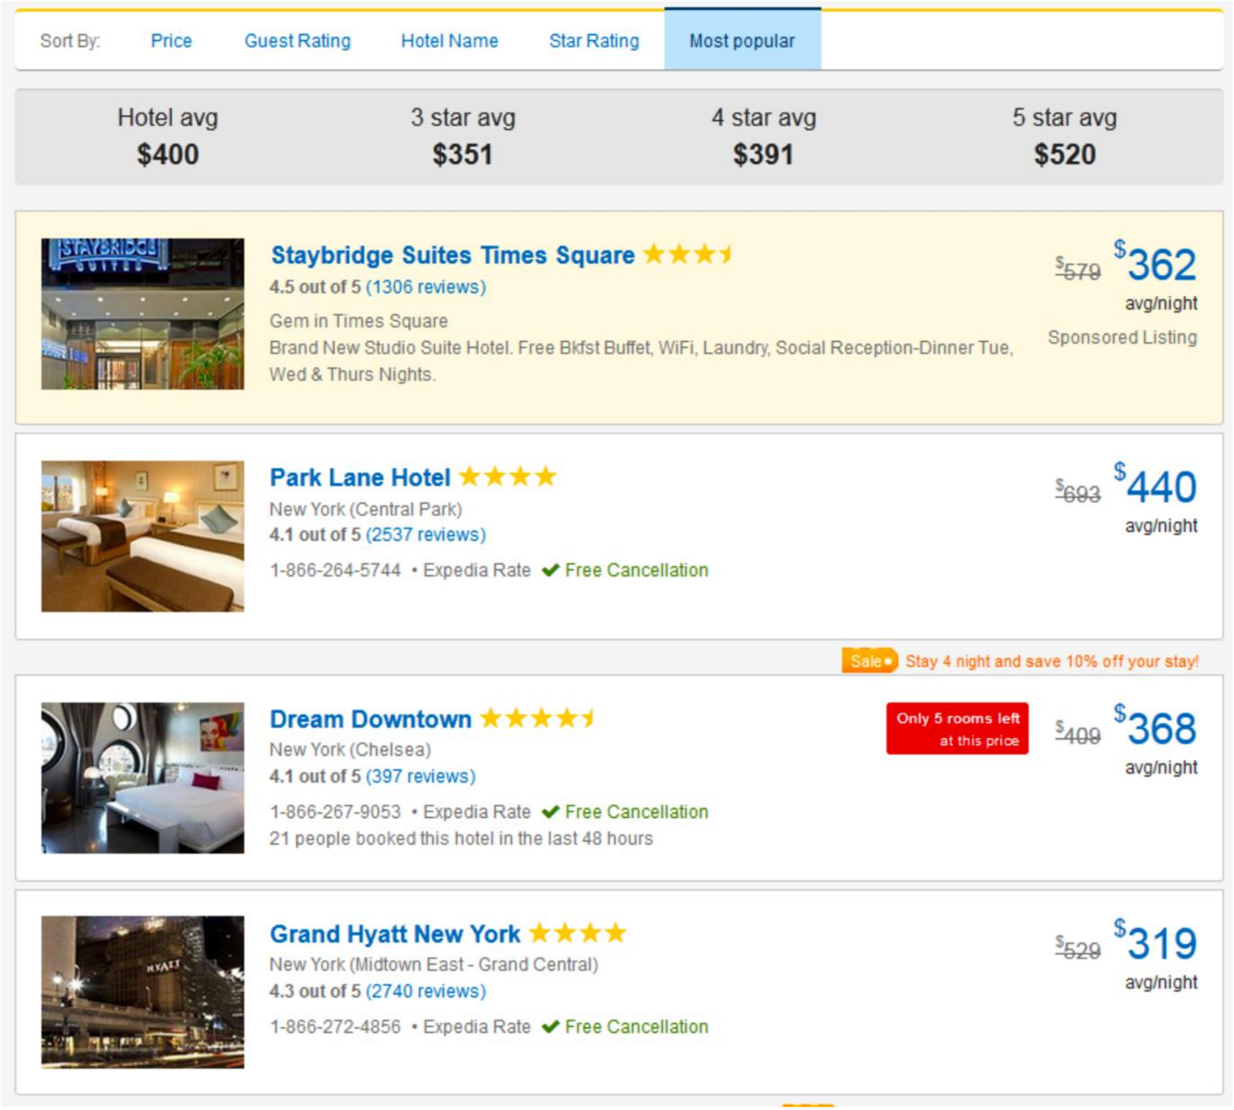

Personalized recommendations tailor product rankings to each consumer

What is a Two-Sided Digital Market?

Consumers | **Online Platform** | **Sellers**

Consumers make search and purchase decisions

Platform Chooses
- Design
  - **Recommendation system**
  - Feature visibility
- Revenue model

Sellers set prices

What is a Two-Sided Digital Market?

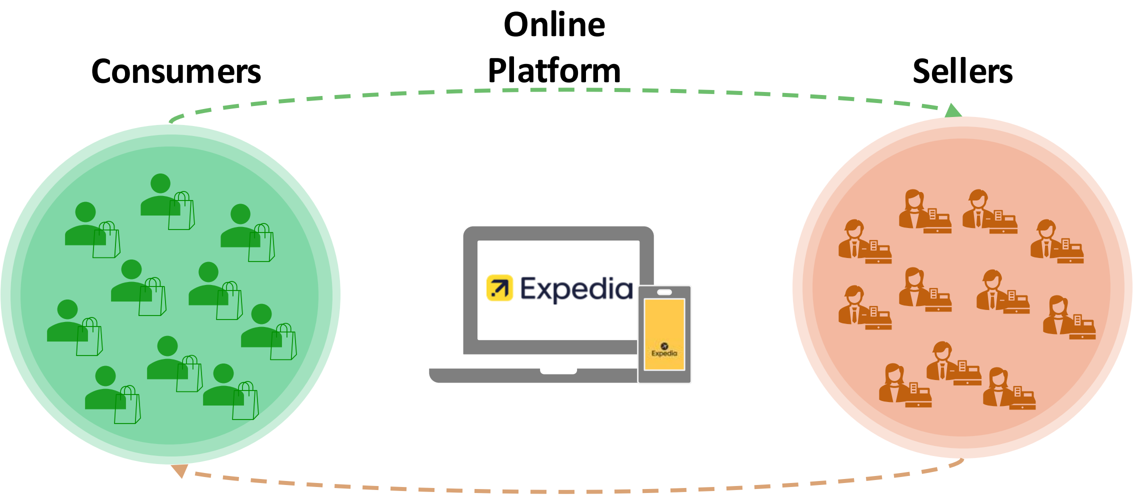

Restaurant and Grocery

instacart

Postmates

GRUBHUB

DOORDASH

Event Tickets

StubHub

ticketmaster

Accommodation

$475 Billion (2022)

Expedia

Booking.com

airbnb

Retail

amazon.com

60% of units sold from third parties

Introduction

Research Question: What are the welfare effects of personalized recommendations when sellers can adjust prices, and consumers update beliefs?

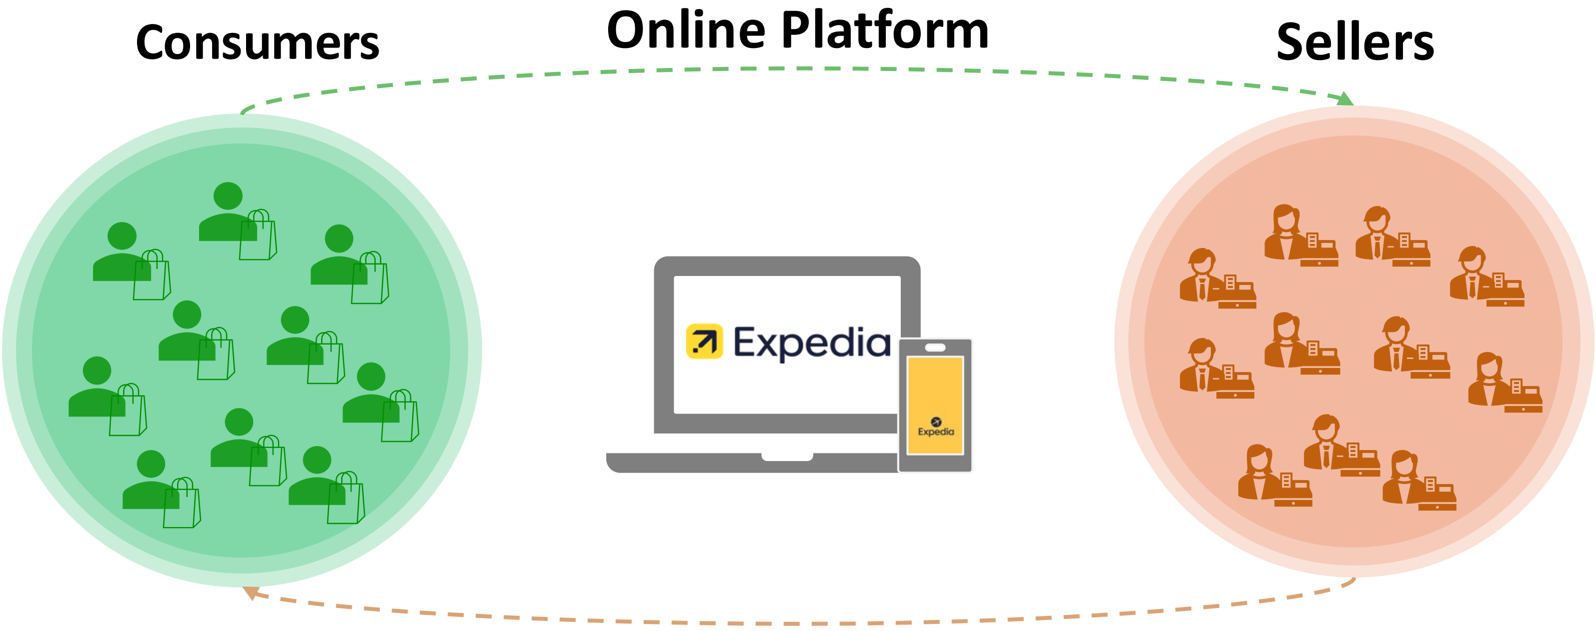

Example of Personalized Recommendations

- Some research suggests: Personalization ⟶↑match quality and ↓ search effort ⟶ ↑Consumer Welfare
- But what happens to **prices**?

Example: The Elvis Hotel

Consumer Types

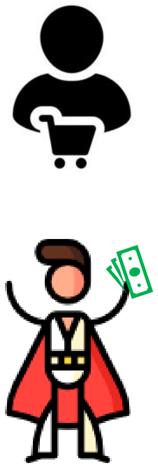

Platform Mediates:
who typically sees the hotel?

Default Recommendations

Personalized Recommendations

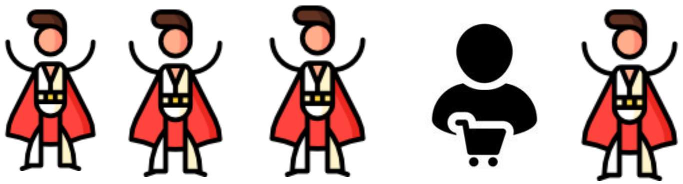

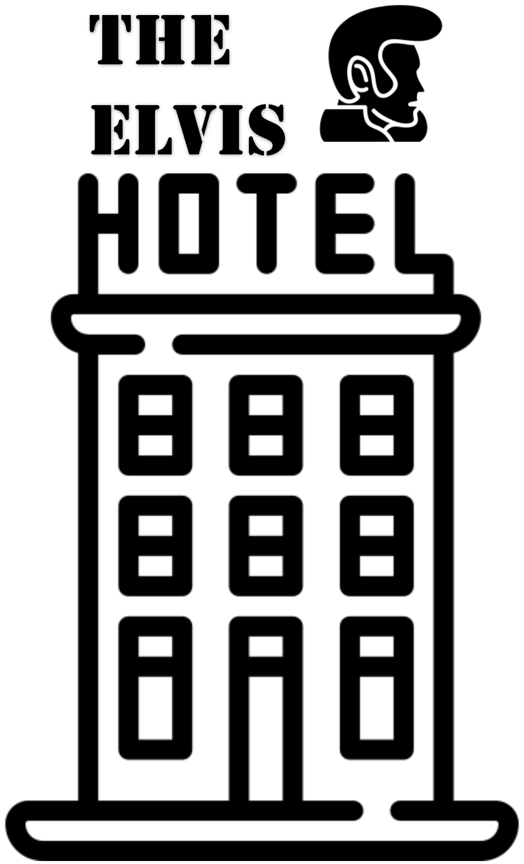

Hotel Sets Prices

Default Recommendations

$120/Night

Personalized Recommendations

↑ **$160/Night**

Welfare Tradeoff: better match but higher prices

More Examples

Price Competition for Product Rankings

Product rankings depend on price and features

- Sellers can improve product ranking by lowering price
- Rec system impacts equilibrium prices

Different recommendation systems change relationship between price and ranking

- ↑ price sensitivity ⟹ ↑ price competition
- Personalization changes competition for product rankings
  - Changes equilibrium prices, but direction unclear

Change in ranking for fixed $ change in price

Back

Co-Ranking of Close Substitutes

Demand depends on availability of substitutes

Example: Two Elvis-themed hotels are close substitutes

Recommendation systems can rank close substitutes similarly or spread them out

- Co-ranking substitutes
  - ↑ seller price competition
  - ↓ likelihood of a purchase on the platform

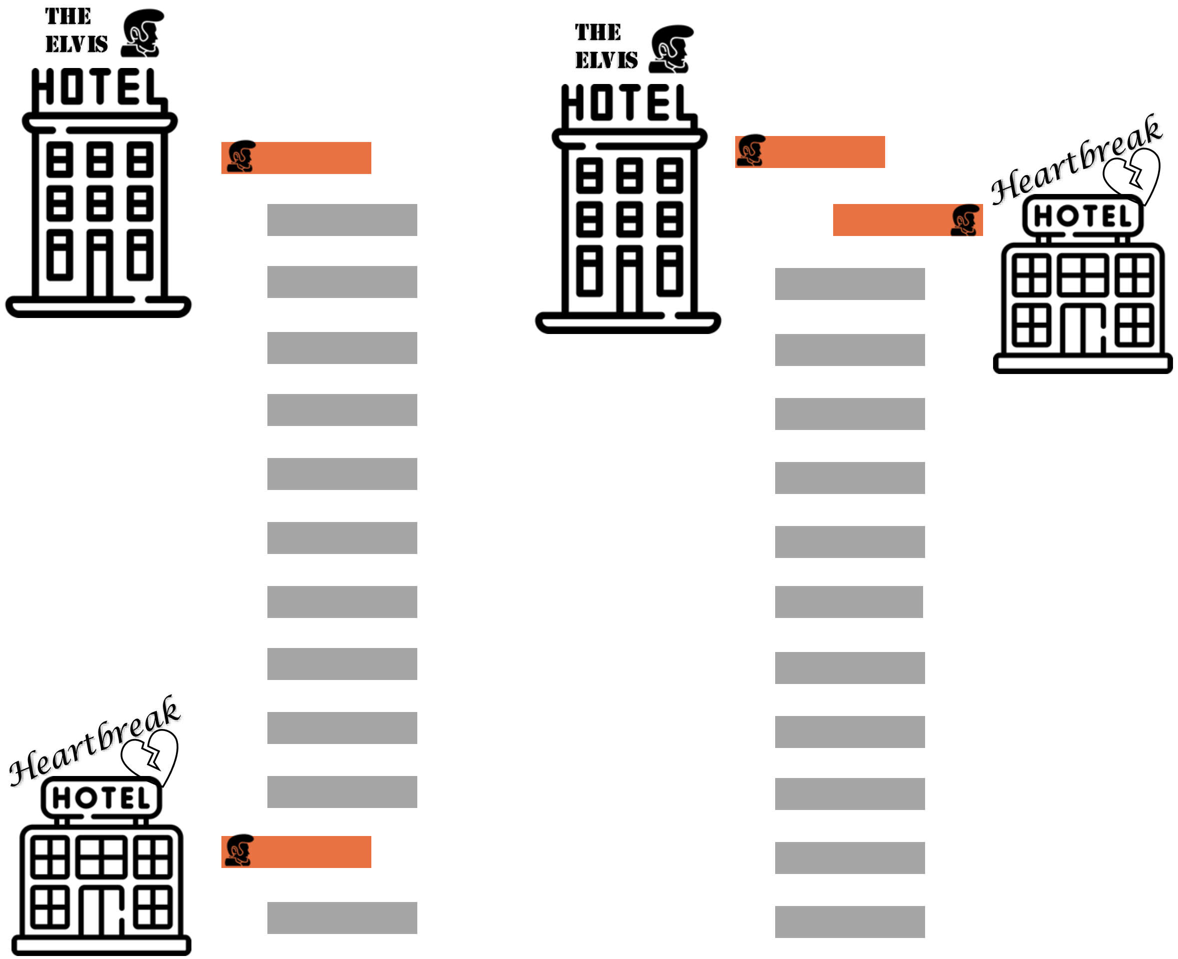

Back

This paper

Research Question: What are the welfare effects of personalized recommendations when sellers adjust prices, and consumers update beliefs?

Data: Click-stream data on hotel search and purchases from Expedia Group

What does this paper do?

1) Show evidence that both search costs and consumer beliefs drive position effects

2) Develops a structural model of demand, platform recommendations, and hotel pricing behavior

3) Trains increasingly personalized recommendation systems using data from an A/B test (RCT)

4) Uses the structural model to evaluate welfare effects of personalized recommendations
   - Holding prices fixed and allowing prices to change

Slot impacts demand even when recommendations are random

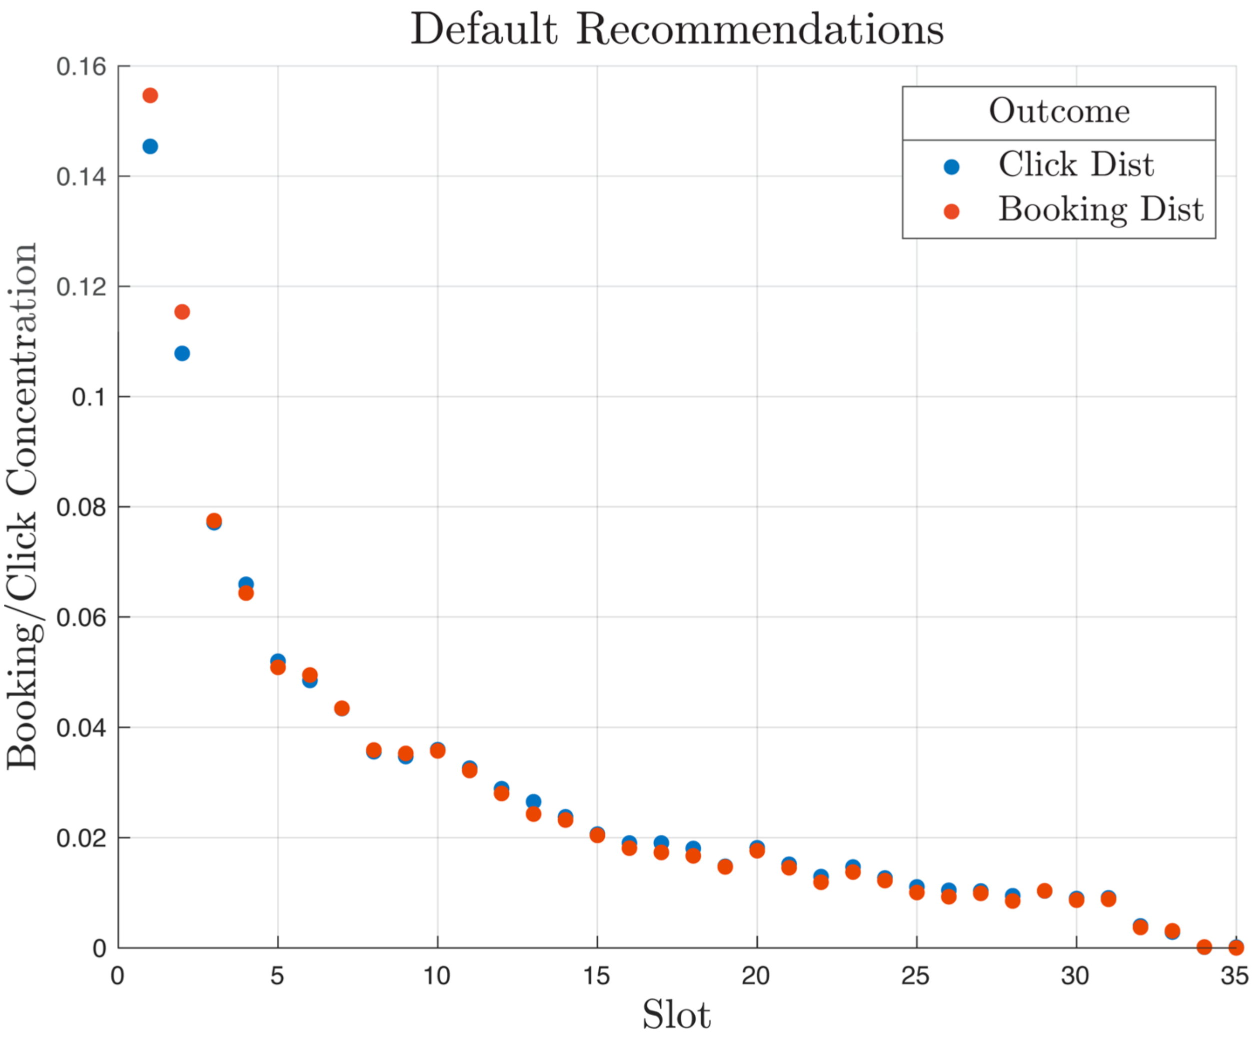

Count Data

Back

Structural Model Outline

Consumer Individual Demand – Sequential Search
- Slot impacts demand via search cost and beliefs
- Estimated via maximum simulated likelihood
  - Inner-loop solves reservation utilities

Needed for welfare

Combine results to get elasticity of demand

Platform – Recommendation Algorithm
- Reverse engineer recommendation system
- Estimated with machine learned ranker and sequential logits

Supply Side – Hotels Choose Prices
- Marginal cost is opportunity cost
- Includes economies of scale and soft-capacity constraints
- BLP type instruments (features of rivals)

Recommendation Systems (Ensemble of LambdaMARTs)

Less Personalized ⟷ More Personalized

Recommender	"Features"	"Query"	"Personalized"	"Most Personalized"
Input Data	Product features	Product features	Product features	Product features
		Query features	Query features	Query features
			Consumer features	Consumer features
				Purchase history

Results

Position Effects: Both search cost and consumer beliefs drive position effects

Without price adjustments, personalization improves welfare
- Hotels: minimal change in quantity and profits
- Platform: minimal change in revenue
- Consumers Surplus: ↑ 2.3% of total booking revenue (~$0.9 Billion gain in 2013)

Primary Results: **Welfare loss** once sellers update prices
- Hotels: ↓ 4.5% decrease in quantity and ↑ 4.9% increase in profits
- Platform: minimal change in revenue
- Consumers Surplus: ↓5% of total booking revenue (~$2 Billion loss in 2013)
  - 190% of the increase in hotel profit

Personalized recs. With Star-level economies of scale and soft capacity constraints
Welfare Loss

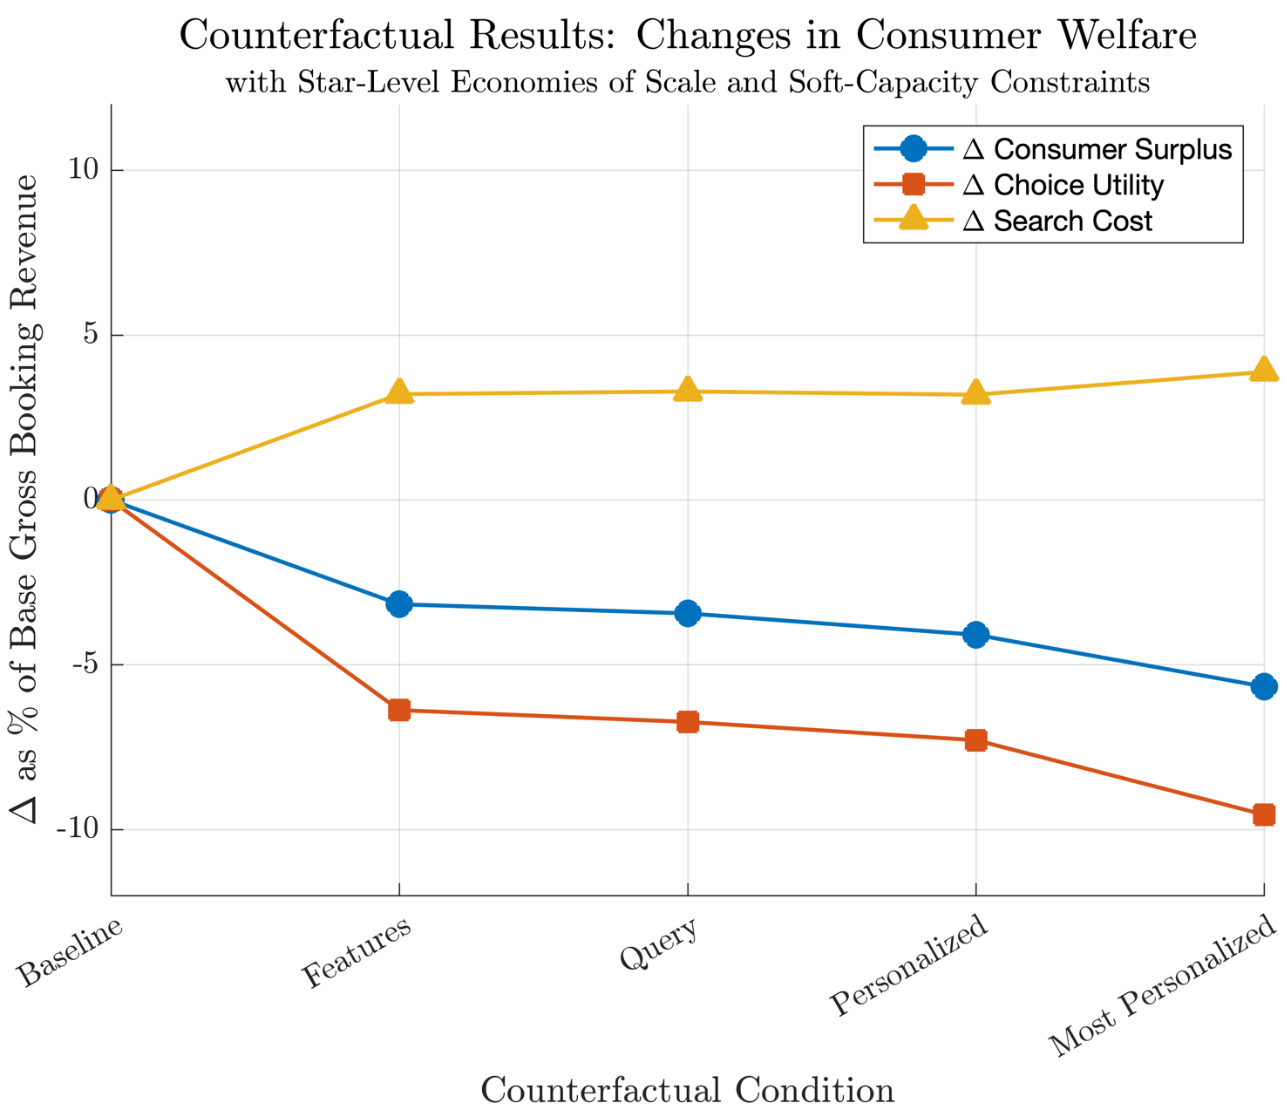

Back

Price Changes

Personalized recs. With Star-level economies of scale and soft capacity constraints
Welfare Loss

Table 5: Counterfactuals with Star-Level Economies of Scale and Soft-Capacity Constraints

		Recommendation System			
Outcomes	**Baseline**	**Features**	**Query**	**Personalized**	**Most Personalized**
Quantity	517.6	495.2	494.8	494.2	494.3
Gross Booking Revenue ($100s)	1,830.09	1,825.62	1,829.00	1,827.90	1,829.79
Hotel Profits ($100s)	974.02	1,020.00	1,021.20	1,021.32	1,022.03
Approx Platform Revenue ($100s)	183.01	182.56	182.90	182.79	182.98
Consumer Welfare					
Δ Consumer Surplus ($100s)	0	-27.37	-62.97	-66.19	-92.02
Δ Choice Utility ($100s)	0	-75.16	-124.19	-118.06	-158.50
Δ Search Cost ($100s)	0	47.79	61.22	51.88	66.48

fixed Prices | figure | w/ fixed mc | w/ common scale

Counterfactual Results Continued

Primary Results

- Welfare loss once sellers update prices.
- Loss increases with level of personalization

Highlights overlooked concern in ecommerce platform research and regulation

- Better recommendation systems may reduce competition and harm consumer welfare

Are there policy alternatives that are welfare improving?

- Next steps: Increase recommendation systems price sensitivity
- Revise recommendations from $f\left(\mathrm{Price}_{\mathrm{jt}}, X_{\mathrm{ijt}}\right)$ with $f\left(\boldsymbol{\alpha}\mathrm{Price}_{\mathrm{jt}}, X_{\mathrm{ijt}}\right)$

Policy Counterfactual: Price Tuned Recommendations (Next Steps)

$f(\boldsymbol{\alpha}\text{Price}_{\text{jt}}, X_{\text{ijt}})$

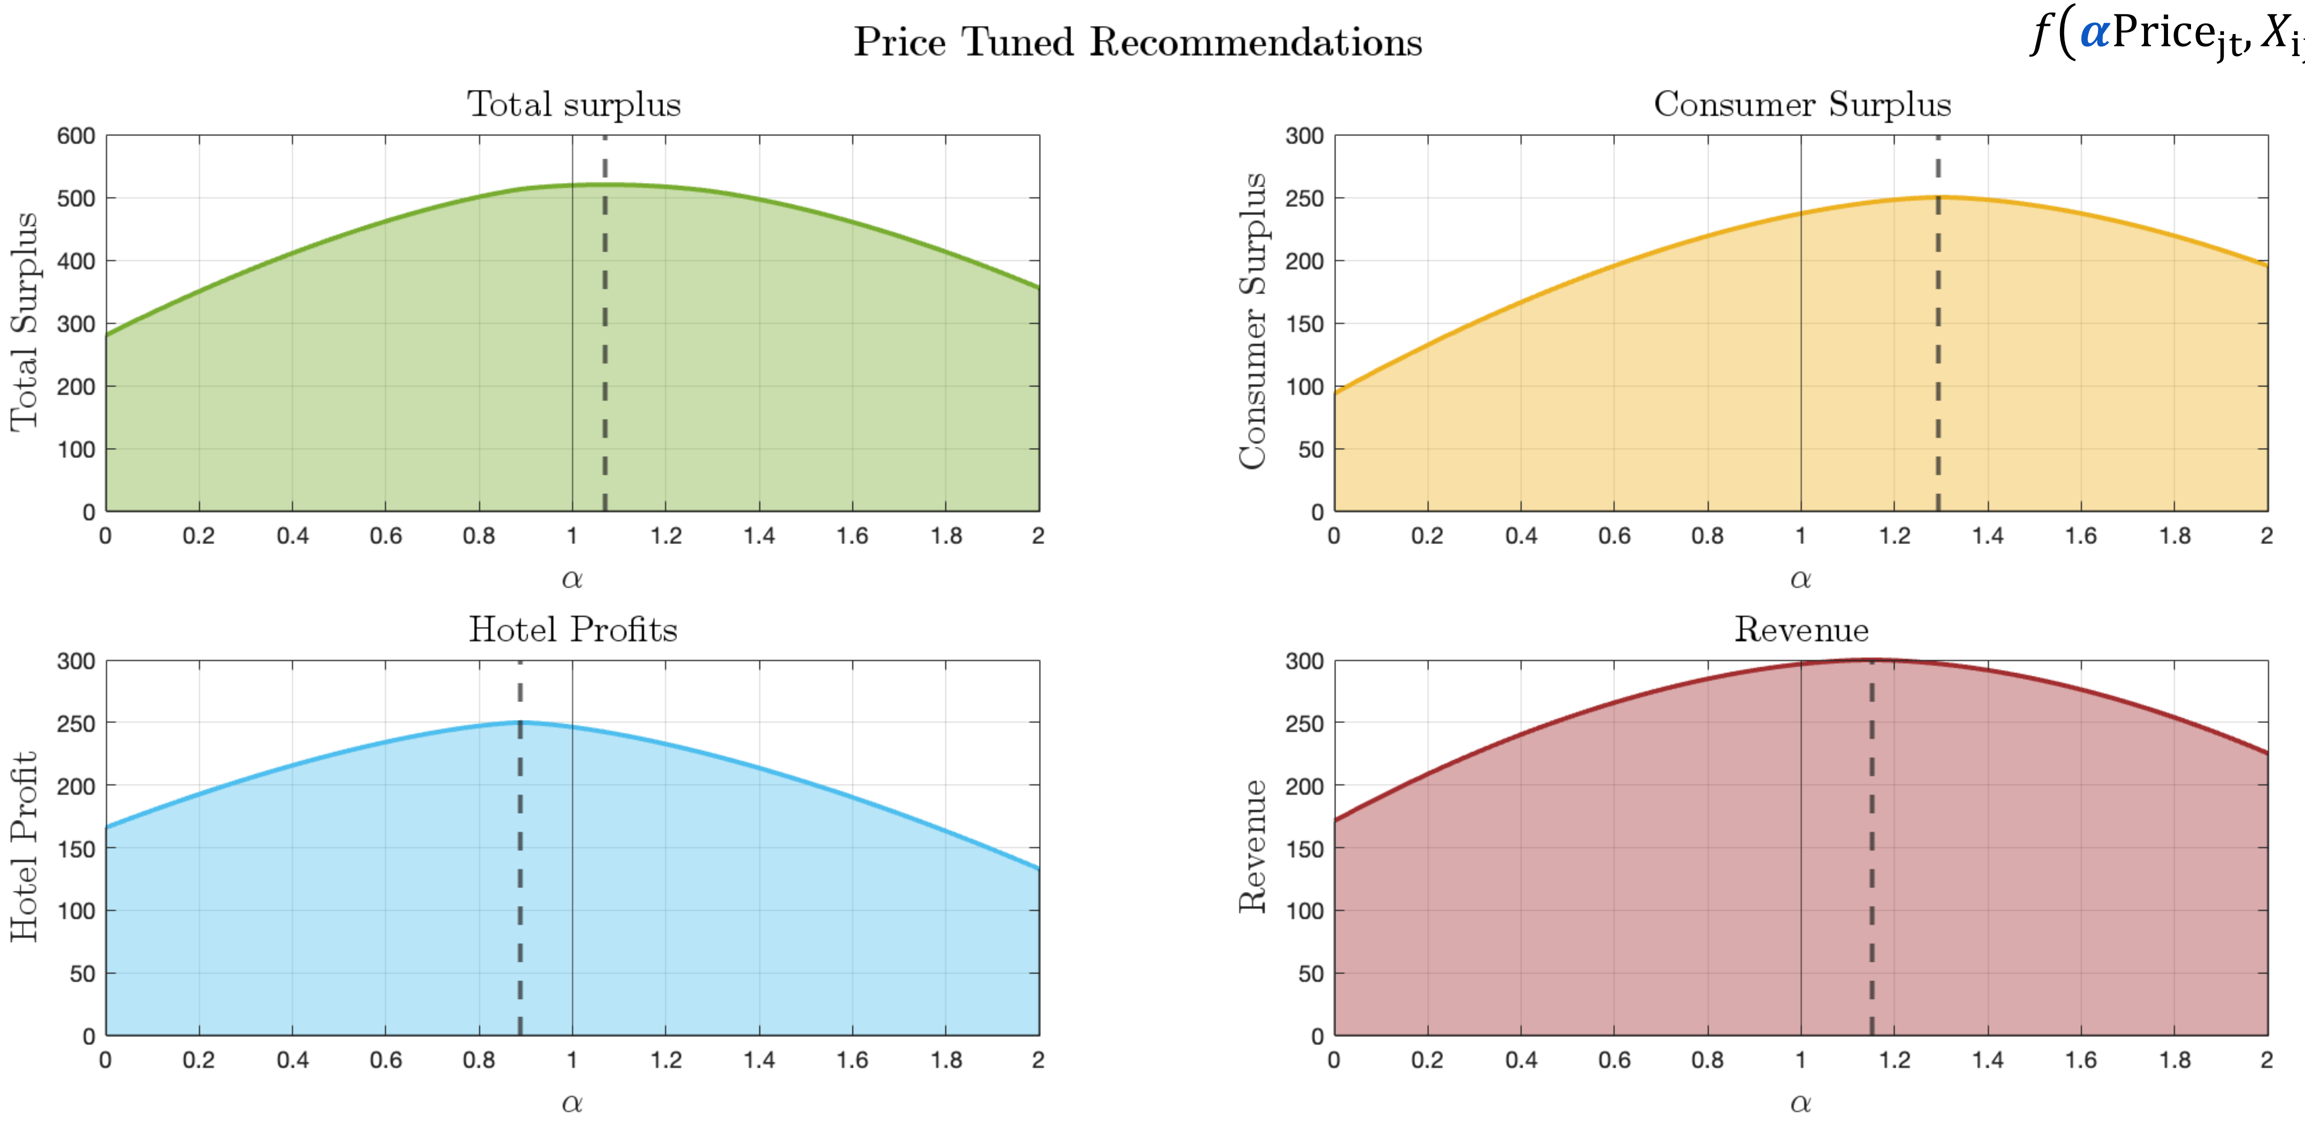

Conclusion

Paper Overview

Structural Model

Data: Clickstream data from Expedia

Demand Model
- Optimal sequential search
- Estimated via maximum simulated likelihood

Platform Model
- Reverse engineer recommendation system

Supply Model
- Hotels choose price
- Marginal cost is opportunity cost of inventory availability
- Economies of scale and soft-capacity constraints

Recommendation Systems

Data: A/B test w/ random slots

Common Recommendations
- Product features

Query Adjusted

+ Query features (ex nights, children)

Personalize on Observables

+ consumer observables (consumer country)

Personalized on Past Purchases

+ past transactions, tracked data

Counterfactuals

Baseline: Default recommendations

Increasingly personalized recs

Ignoring price updates
- Consumer welfare gain

With price updates
- Consumer welfare loss

Without capacity constraints
- Smaller welfare loss

Next Steps

Price transparency

Price transparency w/ personalized

"Price tuned" personalization

Contributions

Feature emphasis Ellison and Ellison (2009), Gardete and Antill (2020), Blake, Moshary, Sweeney, and Tadelis (2021), Abaluck, Compiani, and Zhang (2022)

- Introduces search model where consumers learn about match quality and hidden product features

Self-preferencing Lee and Musolff (2021), Lam (2021), Teng (2022), Farronato, Fradkin, and MacKay (2023), Reimers and Waldfogel (2023)

- "Model of a model" machine learning approach to reverse engineer recommendation systems

Position effects, personalization, recommendations, and platform design Dinerstein, Einav, Levin and Sundaresan (2018) ,Ursu (2018), Compiani, Lewis, Peng and Wang (2021), Agrawal, Athey, Kanodia, and Palikot (2022), Donnelly, Kanodia, Morozov (2023), Moerhing (2023)

- Provides evidence that position effects depend on search cost and consumer beliefs
- Evaluates an industry standard approach to personalizing recommendations
- Structural model that endogenizes seller pricing behavior

Conclusion

Personalization Paradox: $\uparrow$ Personalization of recommendations $\Longrightarrow \downarrow$ Consumer Welfare

- Improve welfare by steering consumers to products that match their tastes
- Worsen welfare since sellers increase prices to profit from less price sensitive demand

Highlights the importance of considering how prices change with platform design policies

- Develops structural model suitable for such counterfactuals

Next Step

- Price tuned recommendation systems

The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets

Aaron P. Kaye

MIT Sloan & CSAIL

17th Digital Economics Conference, Toulouse School of Economics

January 9, 2025

Appendices

Structural Model

DEMAND

PLATFORM RECOMMENDATIONS (EXPEDIA)

SUPPLY (HOTELS)

Demand Model

Indirect per-night utility function

$$u_{ijt} = \underbrace{\delta_{ijt}^{v} + \varepsilon_{ijt}^{v}}_{Visible} + \underbrace{\delta_{ijt}^{h} + \varepsilon_{ijt}^{h}}_{Hidden}$$

Cost c_{ijt} to learn

- $\delta_{ijt}^{v/h}$: utility from consumer and product observables
- $\varepsilon_{ijt}^{v/h}$: match quality

Search cost: Must pay c_{ijt} to learn hidden utility δ_{ijt}^{h} and ε_{ijt}^{h}

- c_{ijt} depends on slot

Demand: Search and purchase decisions depend on $c_{ijt}, \delta_{ijt}^{v}, \varepsilon_{ijt}^{v}, \delta_{ijt}^{h}, \varepsilon_{ijt}^{h}$ and **beliefs** about $\delta_{ijt}^{h}, \varepsilon_{ijt}^{h}$

Platform Design

- Recommendation system orders items into slots, based in part on δ_{ijt}^{v} and δ_{ijt}^{h}, which determines c_{ijt}
- Personalizing recommendations changes c_{ijt} and the relationship between **slot , δ_{ijt}^{v} and δ_{ijt}^{h}**
- Drip pricing shifts a portion of utility from δ_{ijt}^{v} to δ_{ijt}^{h}

Utility | Search Cost | Error Structure | Benchmarks

Demand Model: Model Details

Param	Included Elements
ρ_i	• Per-night transaction price
Visible Features $\delta^v_{ijt}(\boldsymbol{x}^v_{\mathbf{ijt}})$	• Indicators for star rating 1-5 • Brand-star indicators for star rating 2-5 • Property review score (Spline) • Market-Time of Stay Effects • Missing value indicator(s) ○ No star rating, no review score • Consumer segment groups ○ Time ahead of of stay ○ Time of search ○ Search on weekends ○ Number of nights
Hidden Features $\delta^{\boldsymbol{h}}_{ijt}(\mathbf{x}^{\boldsymbol{h}}_{\mathbf{ijt}})$	• Hidden Price • Location desirability score 1 (Spline) • Location desirability score 2 (Spline)

Param	Included Elements
λ	• Determines how much of match quality is learned from search
Random Coefficients	• Inside option • Star-ratings • Price • Search Cost
Correlated Random Coefficients	• Price – Search Cost
Consumer Info $\boldsymbol{\Omega}_{\mathbf{it}}$	• Slot Ranking • Star Rating • Promotions • Covariance with $\mathbf{x}^{\boldsymbol{h}}_{\mathbf{ijt}}$
Consumer Price Info	• Headline price • Mean hidden price rate

Demand Estimation Details (Maximum Simulated Likelihood)

Construct joint likelihood of search and purchase decisions combining Details

1. Sequential search rules Weitzman (1979)
2. Logit-smoothing Train (2002, 2009)

Sample selection adjustments

- Selection on clicks ⟶ condition likelihoods on at least one click Details
- Selection on purchases ⟶ sample weights

Test structural assumptions on position effects Details

- Repeat demand estimation estimation under alternative structural assumptions
  - Position effects depend on search cost and beliefs (primary specification)
  - Position effects depend on only on search cost (benchmark specification)

Demand Estimation: Utility, Search Cost, and Reservation Utility

Per-night utility:

$$u_{ijt}^{[s]} = \delta_{ijt}^{v[s]} + \delta_{ijt}^{h[s]} + \lambda\varepsilon_{ijt}^{v[s]} + \varepsilon_{ijt}^{h[s]}(\lambda)$$

Search Cost:

$$c_{ijt}^{[s]} = \log\left(1 + \exp\left(\kappa_i^{[s]} + \sum_{k\in K} \tau_k \left(\log\left(slot_{ijt}^{appear}\right) - \gamma_k\right)_+\right)\right)$$

Reservation Utility:

$$r_{ijt}^{[s]} = \delta_{ijt}^{v[s]} + \lambda\varepsilon_{ijt}^{v[s]} + E\left[\delta_{ijt}^{h[s]} | \Omega_{\mathrm{it}}\right] + \zeta_{ijt}^{[s]}$$

Reservation utility setup

- Information set, Ω_{it} includes star-rating, base price, slot rank, and promotions
- $E\left[\delta_{ijt}^{h[s]} | \Omega_{\mathrm{it}}\right]$ solved by getting $\mathrm{E}\left[\mathrm{x}_{ijt}^{h} | \Omega_{\mathrm{it}}\right]$ before estimation Details
- $\zeta_{ijt}^{[s]} = V\left(c_{ijt}^{[s]}, \rho_i^{[s]}, \beta_i^{h[s]} | \Omega_{\mathrm{it}}, \theta\right)$ solved in inner loop with grid interpolation Details

Price Details Sequential Search

How does Price Impact Demand?

$\Omega_{\mathrm{i}}: \{Star\ Rating, Price, slot^{rank}, Promotion\}$

Utility:

- Directly through preference for price

$$u_{ijt} = \alpha_i - e^{\rho_i^{[s]}} \underbrace{\left(p_{jt}^{base} + p_{jt}^{h}\right)}_{transaction\ price} \underbrace{+\beta_i^{v} x_j^{v} + \beta_i^{h} \mathrm{x}_j^{h}}_{features} + \underbrace{\delta_{\mathrm{it}}}_{\substack{time,\\ segment\ FE}} + \underbrace{\lambda \varepsilon_{ijt}^{v[s]} + \varepsilon_{ijt}^{h[s]}(\lambda)}_{match\ quality\ \sim EV1}$$

Search Cost:

- Indirectly through slot (slot is a function of price)

$$c_{ijt}^{[s]} = \log\left(1 + \exp\left(\kappa_i^{[s]} + \sum_{k \in K} \tau_k \left(\log\left(slot_{ijt}^{appear}\right) - \gamma_k\right)_+\right)\right)$$

Reservation Utility:

- Directly through expected price
- Indirectly through expected utility of hidden features (via slot and expected price)
- Indirectly through state variables of value function, $\zeta_{ijt}^{[s]}$ (conditional distribution of hidden utility, and search cost)

$$r_{ijt}^{[s]} = \underbrace{\alpha_i^{[s]} - e^{\rho_i^{[s]}} p_{jt}^{base} + \beta_i^{v[s]} x_j^{v} + \delta_{\mathrm{it}} + \lambda \varepsilon_{ijt}^{v[s]}}_{Known\ Pre-search} - \underbrace{e^{\rho_i^{[s]}} E\left[p_{jt}^{h} | \Omega_{\mathrm{i}}\right] + \beta_i^{h[s]} E\left[\mathrm{x}_j^{h} | \Omega_{\mathrm{i}}\right]}_{Expected\ Hidden\ Utility} + \underbrace{\zeta_{ijt}^{[s]}\left(c_{ijt}^{[s]}, \rho_i^{[s]}, x_i^{nights}, E\left[p_{jt}^{h} | \Omega_{\mathrm{i}}\right], slot_{ijt}^{rank}\right)}_{Reservation\ Value\ Func.}$$

Back

Consumer Choice Model Identification

Table 7.1: Informal Identification of Demand Parameters

	Sequential Search Conditions				Notable Variation		
Parameters	Order	Continuation	Stopping	Choice	Nights	Diversion	Displacement
Utility Parameters							
Consumer Segments: δ_{it}		✓	✓	✓†	✓	✓	✓
Time Effects: ξ_{it}^{month},ξ_{it}^{day}		✓	✓	✓†	✓	✓	✓
Mean: ρ, β^v, β^h	✓	✓	✓	✓	✓	✓	✓
Heterogeneous: Σ_u	✓*	✓*	✓*	✓*	✓	✓	✓
Visible Error Scale: λ	✓	✓	✓		✓	✓	✓
Search Cost Parameters							
Mean: κ, τ_k	✓	✓	✓		✓	✓	✓
Heterogeneous: Σ_κ	✓*	✓$*$	✓*		✓	✓	✓

Note: Checkmarks with an asterisk (✓*) indicate parameters that are identified by repeated decisions within consumer (e.g., clicks and purchase). Checkmarks with a dagger (✓†) indicate parameters that are identified by selecting an inside good versus the outside option, but not from the choice of one inside good over another. "Nights" refers to length of stay. "Diversion" refers to substitution patters from variation in product features and availability. "Displacement" refers the variation in positions caused by advertisements/opaque offers.

Demand Results

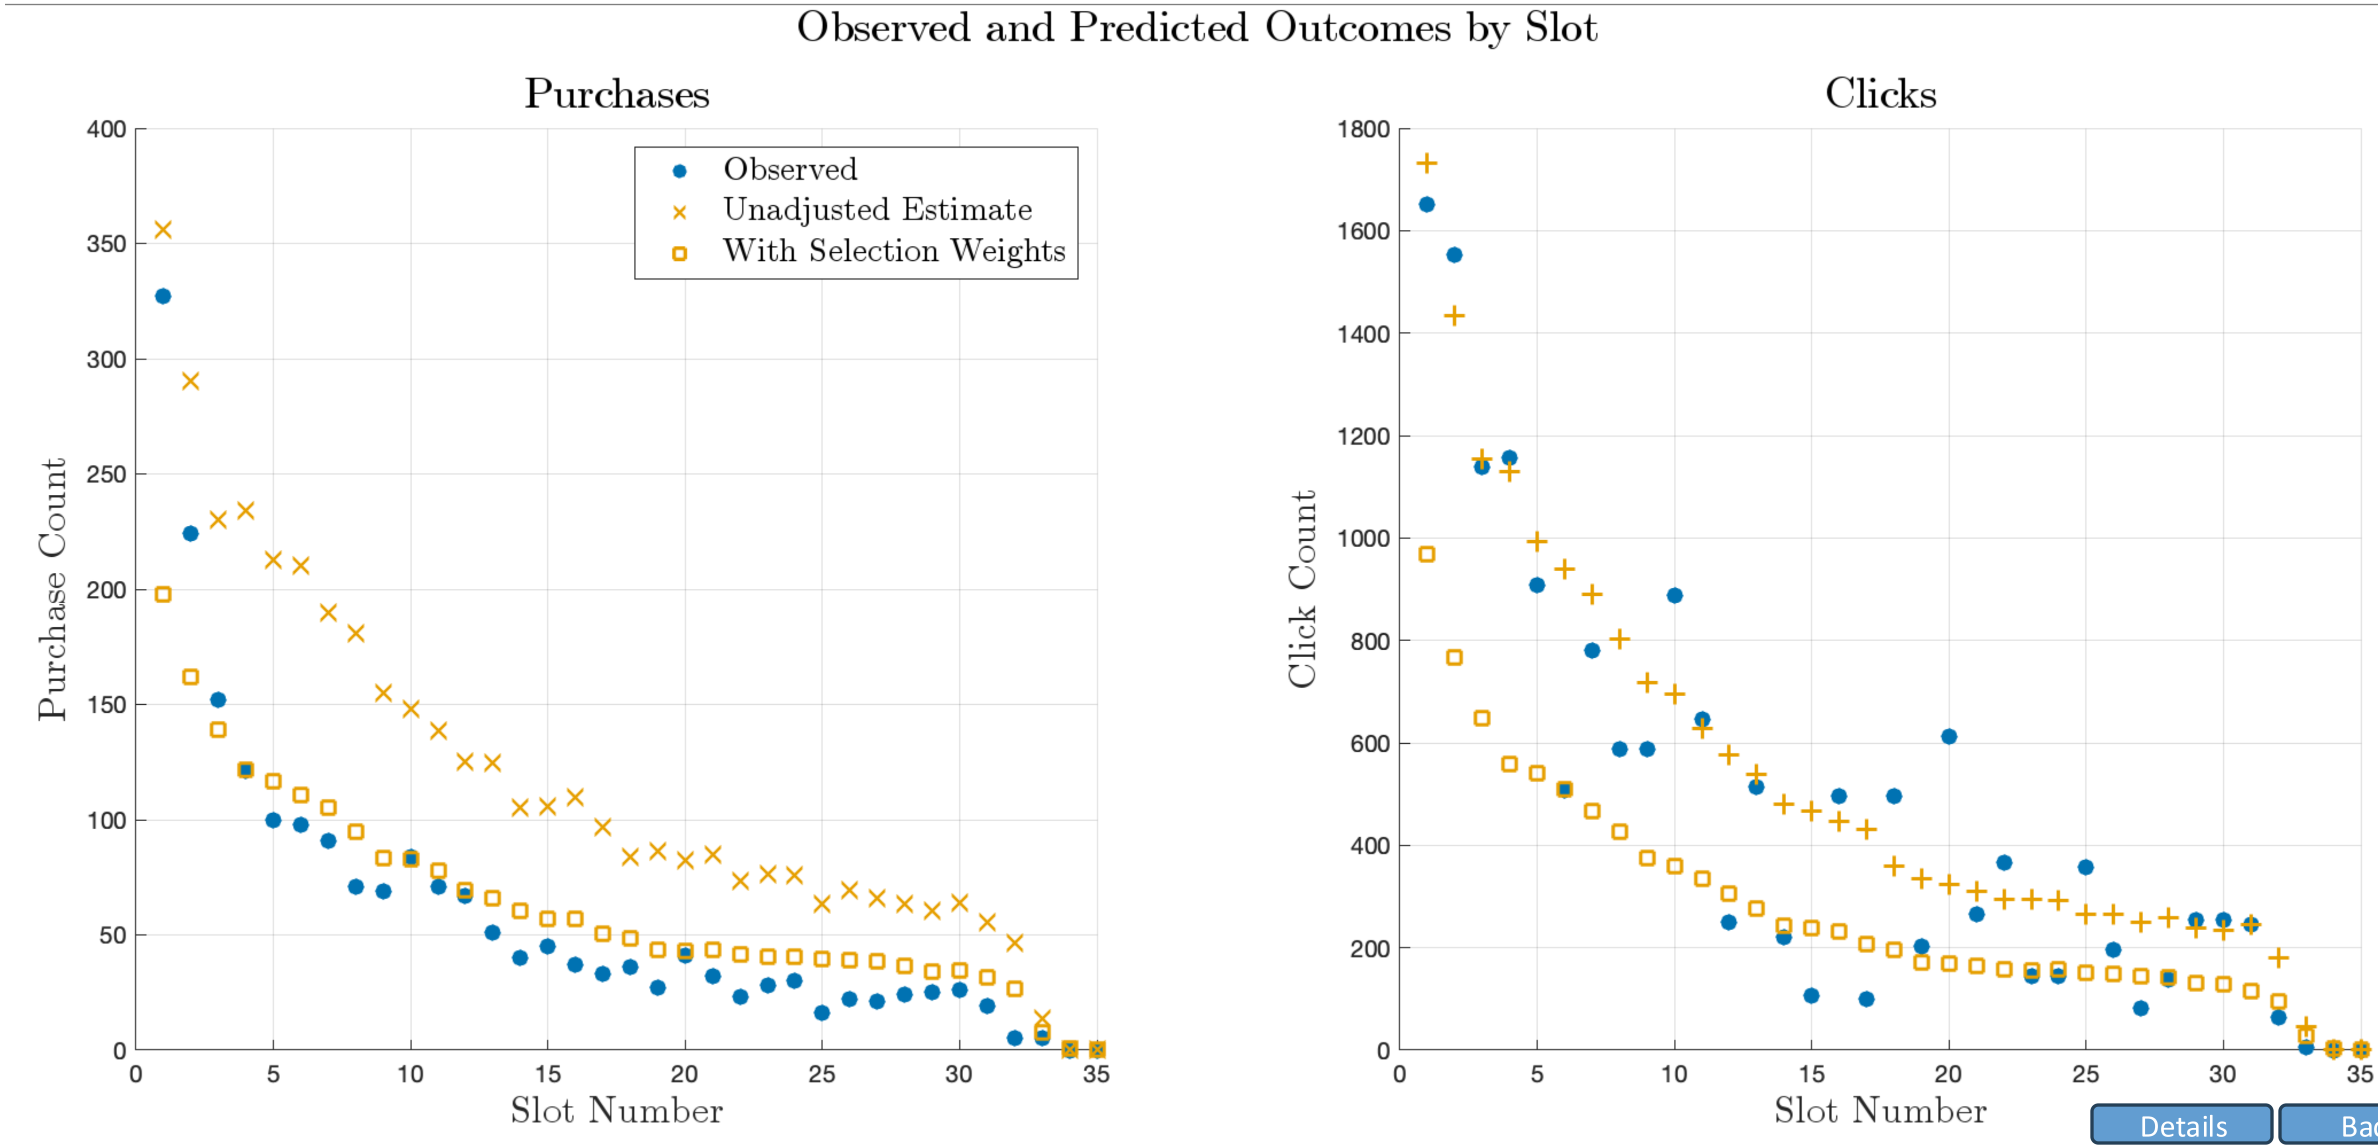

Details Back

Demand Estimates

Table 7.2: Demand Parameter Estimates

Utility Parameters

Variable	(1)
Outside option	1.90
Price ($100s) ρ	-1.76
Match quality split λ	0.28
Visible Features	
3 star	0.30
4 star	0.54
5 star	0.48
Non-star	0.31
2 star brand	-0.16
3 star brand	-0.28
4 star brand	0.03
5 star brand	0.29
Prop. review score	
Spline 1: score 1–3	-0.51
Spline 2: score 3–5	0.04
Mi. dummy	-1.40
Hidden Features	
Location score 1	
Spline 1	0.52
Spline 2	-0.51
Spline 3	0.05
Spline 4	2.61
Location score 2	
Spline 1	0.27
Spline 2	1.50
Spline 3	0.37
Mi. dummy	1.64

Estimation Details

Observations	2,262
Weighted obs.	13,444
Halton draws	400
Smoothing term ω	0.2
Grid points	1,692
Log likelihood	-85,028

Search Cost Parameters

Variable	(1)
Constant	-1.10
Log Slot Appear	
Spline 1	0.11
Spline 2	0.21
Spline 3	0.37
Spline 4	0.08

Random Coefficients

Parameter	(1)
σ^2_{price}	0.729
$\sigma^2_{\text{inside option}}$	0.095
$\sigma^2_{\text{1 \& 2 star}}$	0.080
$\sigma^2_{\text{3 star}}$	0.009
$\sigma^2_{\text{4 star}}$	0.015
$\sigma^2_{\text{5 star}}$	0.000
$\sigma^2_{\text{search cost}}$	0.279
$\sigma^2_{\text{price-search cost}}$	-0.444

Additional Controls

Day of week	✓
Month	✓
Time before stay	✓
Length of stay	✓
Search time	✓
Search on weekends	✓

Notes: Likelihood is the logit-smoothed likelihood for joint search and purchase decisions. Splines are linear B-spline. Variance of random coefficients estimated using Cholesky decomposition.

Back

Structural Model

DEMAND

PLATFORM RECOMMENDATIONS (EXPEDIA)

SUPPLY (HOTELS)

Model: Platform Two Step "Model of the Model" Approach

Platform recommendation system orders results relevance, u^r_{ijt}

$$u^r_{ijtn} = \beta^{slot}_n \underbrace{\psi_{ijt}}_{\text{Deterministic}} + \underbrace{\epsilon_{ijt}}_{\text{Random}}$$

With

$$\psi_{ijt} = f\left(x^r_{ijt}\right)$$

- x^r_{ijt} includes price, product features, consumer observables, and query specific information
- The underlying recommendation systems can be quite complicated

Estimation:

- Approximate $f\left(x^r_{ijt}\right)$using LambdaMART, a machine learned ranker Burges (2010)
- Create out-of-fold predictions of $\hat{\psi}_{ijt}$
- Normalize $\hat{\psi}_{ijt}$
- Fit sequential logit on $\hat{\psi}_{ijt}$ to estimate β^{slot}_n for each slot

Estimation Steps

Details

Platform Results – Out of Sample Fit

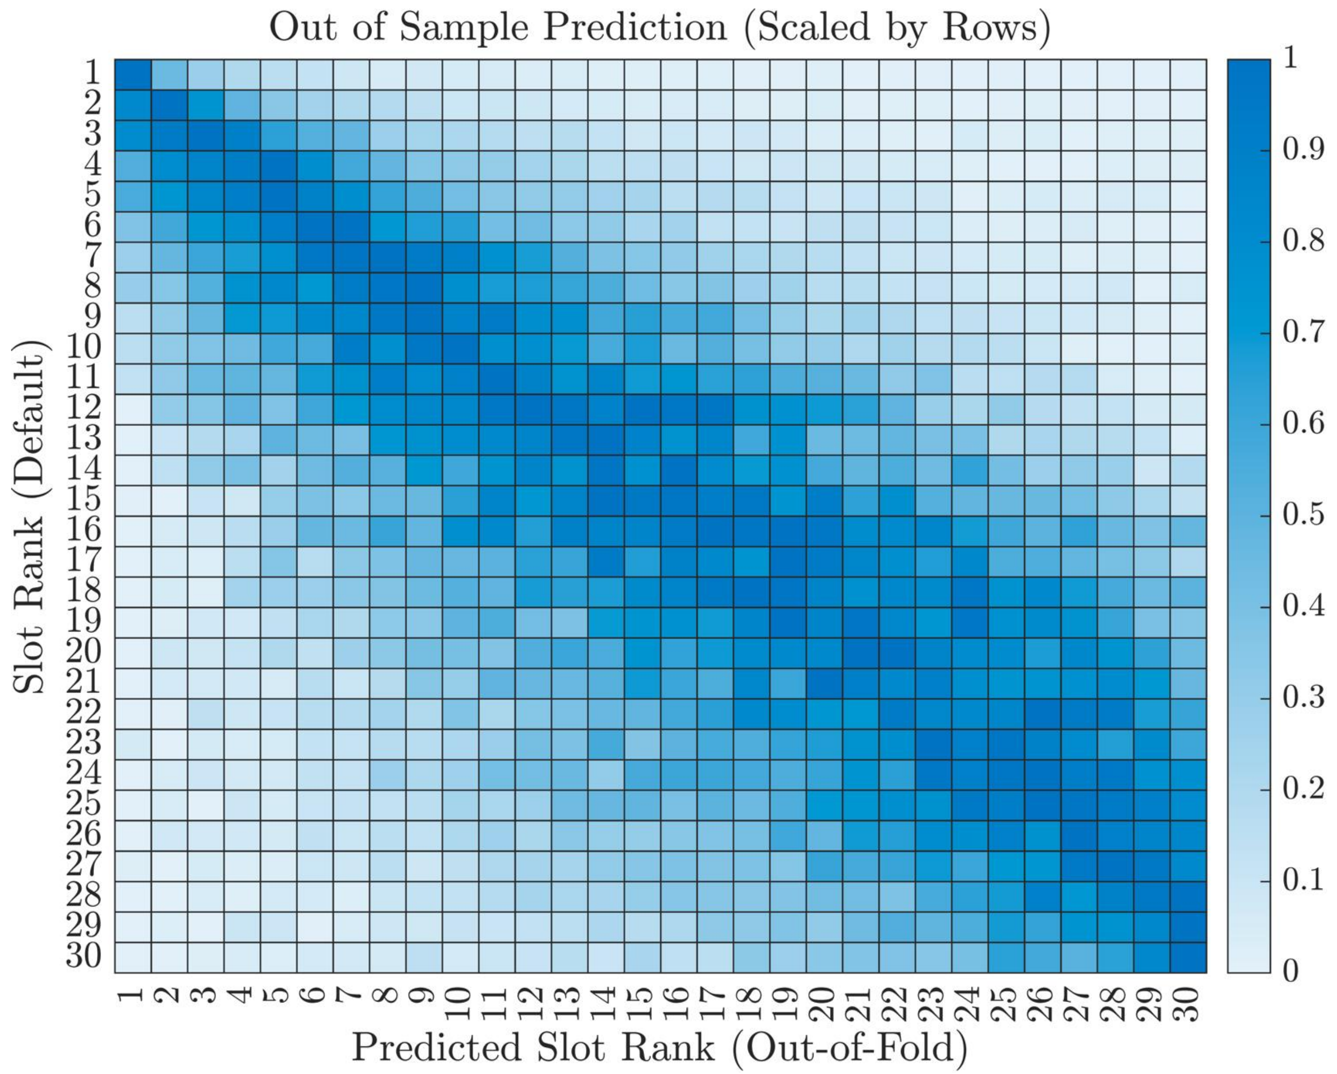

Back

Table

In of Sample

Platform Results – In Sample Fit

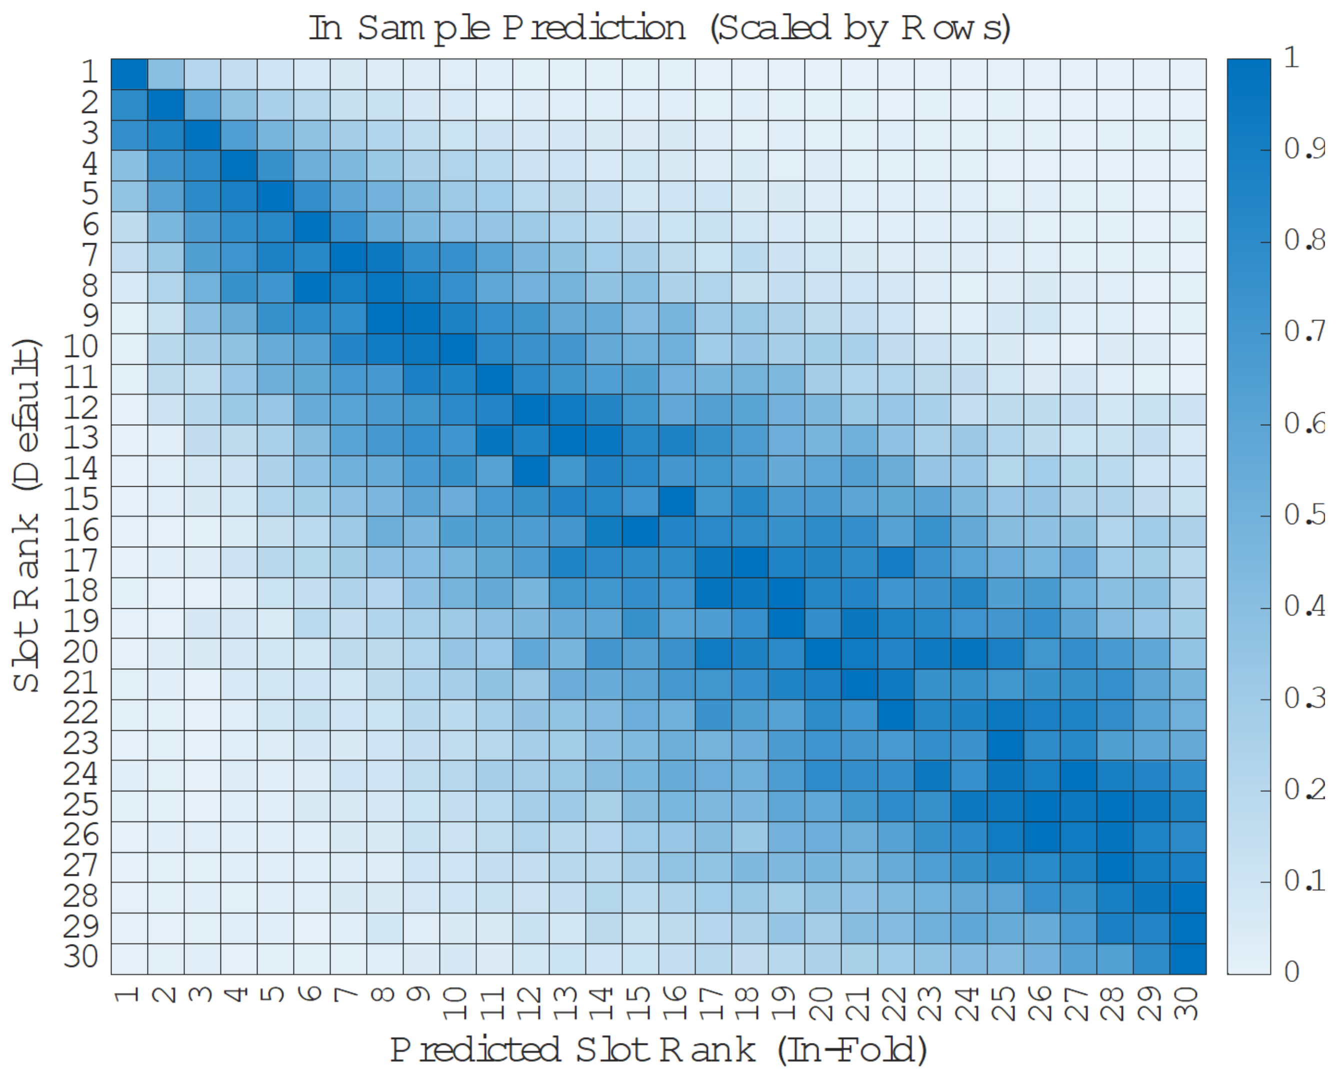

Out of Sample

Platform Results – Out of Sample Fit

Measure	Model	NDCG Loss	Conc Loss
1	Random Benchmark	0.175	0.506
2	LambdaMART (NDCG): Full	0.060	0.276
3	LambdaMART (NDCG): Fold 1	0.061	0.277
4	LambdaMART (NDCG): Fold 2	0.061	0.277
5	LambdaMART (NDCG): Fold 3	0.061	0.276
6	LambdaMART (NDCG): Fold 4	0.061	0.277
7	LambdaMART (NDCG): Fold 5	0.061	0.278
8	LambdaMART (NDCG): Fold 6	0.061	0.276
9	LambdaMART (NDCG): Fold 7	0.060	0.277
10	LambdaMART (NDCG): Fold 8	0.061	0.277
11	LambdaMART (NDCG): Ens	0.059	0.272

vs Benchmark

Platform Model Sequential Logit Results

Estimate β_n^{slot} for each slot:

$$u_{ijt}^r = \beta_n^{slot} \hat{\psi}_{ijt} + \epsilon_{ijt}$$

- β_n^{slot} governs how deterministic each slot assignment is in relevance score (ψ_{ijt}).

Takeaway

Position on the page is more deterministic higher on the page.

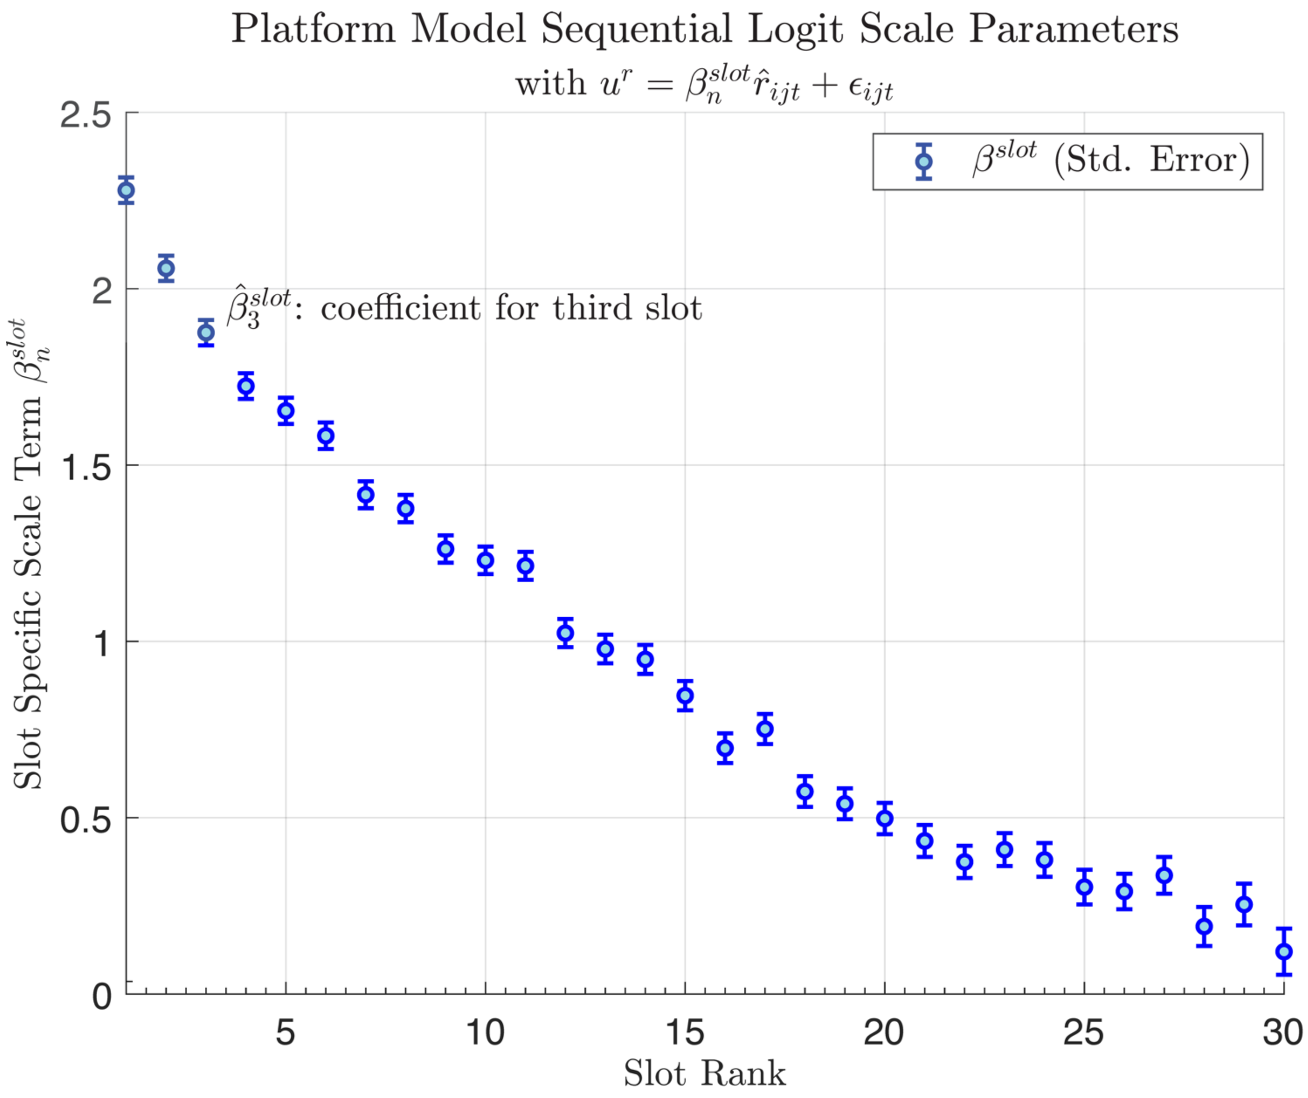

Structural Model

DEMAND

PLATFORM RECOMMENDATIONS (EXPEDIA)

SUPPLY (HOTELS)

Model: Supply Side

Sellers expected profits solving

Platform fee and tax rate

Avg. Variable Cost

$$\underset{p_{jtt'}}{\operatorname{argmax}} E\left[\left((1-\varphi)p_{jtt'} - c_{jtt'}\right) q_{jtt'} \,|\Omega_{jtt'}\right]$$

Quantity

Seller foc

$$\frac{mc_{jtt'}}{(1-\varphi)} = p_{jtt'} + \left(\frac{\partial q_{jtt'}}{\partial p_{jtt'}}\right)^{-1} q_{jtt'}$$

Depends on platform rec system and preferences

Price Schedule ($p_{jtt'}$)

- $p_{jtt'}$ is price for room-night *j*, staying period *t*, and searching period *t'*

Marginal Cost ($mc_{jtt'}$)

- Opportunity cost of the unit available to sell the next period
- Can include additional expected profits conditional on purchase (room service, dining, gambling)

Sellers know $(\Omega_{jtt'})$

- Own costs, elasticity of demand, competing product features and availability

Back

Supply Side Estimation: Three Stage Least Squares

Hotels face **economies of scale and capacity constraints**

- Known from data or platform/demand model

Not separable

$$\frac{mc_{base} + \frac{\partial c(q)}{\partial q} q(\theta, p)}{(1 - \varphi)} = p + \left(\frac{\partial q}{\partial p}\right)^{-1} q(\theta, p)$$

First stage: IV for q_{jt}

$$q_{jtt'} = \alpha_1 x_{jtt'} + \alpha_2 \mathrm{z}_{\mathrm{jtt'}} + \varepsilon_{\mathrm{jtt'}}$$

- $x_{jtt'}$: product features, market-subperiod effects
- z_{jt}: product features and availability of other products in same market, own-star rating interactions.

Second Stage: IV for q_{jt}^2

$$q_{jtt'}^2 = \alpha_3 \left(\hat{q}_{jtt'}^{step\,1}\right)^2 + \epsilon_{\mathrm{jtt'}}$$

Third Stage

$$mc_{\mathrm{jtt'}} = \beta x_{jtt'} + \gamma_1 \hat{q}_{jtt'}^{step\,1} + \gamma_2 \widehat{q_{jtt'}^2}^{step\,2} + \omega_{jtt'}$$

Estimates

Distribution of Marginal Costs by Star-Rating

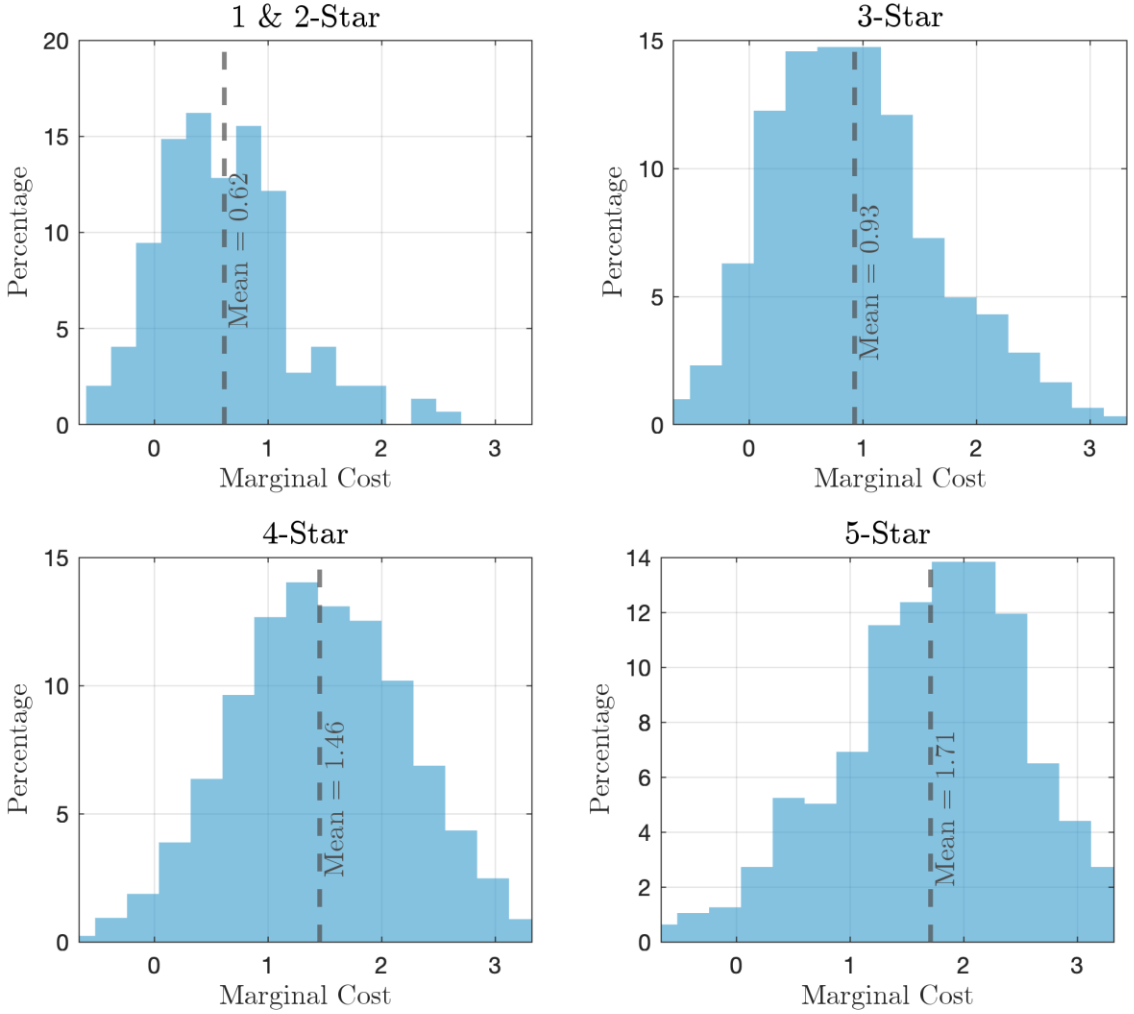

Note: Marginal cost in $100 per room-night.
Each series is truncated at its respective 5th and 95th percentile.

Table 7.5: Supply Side IV Regression Analysis Results

	Pooled	Star Rating Specific			
Variable	**All Ratings**	**Two-Star**	**Three-Star**	**Four-Star**	**Five-Star**
Intercept	-0.307	0.527	1.022***	1.484***	3.685***
	(0.503)	(0.614)	(0.124)	(0.114)	(0.134)
$\hat{q}_j^{(1)}$	-0.198***	-0.091	-0.036	-0.091	-1.287***
	(0.037)	(0.072)	(0.094)	(0.092)	(0.149)
$\hat{q^2}_j^{(2)}$	0.033***	-0.016	0.033*	0.026**	0.199***
	(0.008)	(0.012)	(0.032)	(0.028)	(0.051)
Two/Three-Star	0.577***	–	–	–	–
	(0.088)				
Four-Star	1.010***	–	–	–	–
	(0.086)				
Five-Star	2.688***	–	–	–	–
	(0.111)				
Additional Controls					
Product Features	✓	✓			
Location Desirability	✓	✓			
Month–Weekend–Subgroup	✓	✓			
Observations	3492	3492			
Degrees of Freedom	3437	3429			
RMSE	0.783	0.761			
R^2	0.638	0.660			
Adjusted R^2	0.632	0.654			
First-stage F-statistic	103	103			

Personalized Recommendation Systems Training for Counterfactuals

Recommendation Systems

Train ranking systems using data from Expedia's RCT

Details

- Data from RCT were displayed in random order
- Relevance scores: Booking = 5, Click = 1, Impression = 0
- Model training approach based on winning entry
  - Ensemble of LambdaMARTs with NDCG Loss (170 models)

Use increasing levels of personalization

Details

- Common Recommendations: Product features, competitive info
- Query Adjusted: + query features (ex nights, children)
- Personalize: + consumer observables (ex: consumer country)
- Most Personalized: + past transactions, tracked navigation data

Evaluate out of sample performance

- Out of sample fit should improve with personalization

Details

Recommendation Systems (Ensemble of LambdaMARTs)

Less Personalized ⟷ More Personalized

Recommender	“Features”	“Query”	“Personalized”	“Most Personalized”
Input Data	Product features	Product features Query features	Product features Query features Consumer features	Product features Query features Consumer features Purchase history

Recommendation System Performance

Out of sample performance improves with level of personalization

Measure	Model	NDCG Loss	Conc Loss	MAP	MRR
1	Random Benchmark	0.673	0.480	0.850	0.846
6	LambdaMART (Ensemble): Base Info	0.544	0.302	0.699	0.692
7	LambdaMART (Ensemble): with Query Info	0.540	0.301	0.695	0.686
8	LambdaMART (Ensemble): Personalized Basic	0.537	0.299	0.692	0.681
9	LambdaMART (Ensemble): Personalized Full	0.533	0.300	0.686	0.676

Table 2: Comparison of Model Results

Back

Counterfactuals

Simulation

Counterfactual Setup

Counterfactual Timing

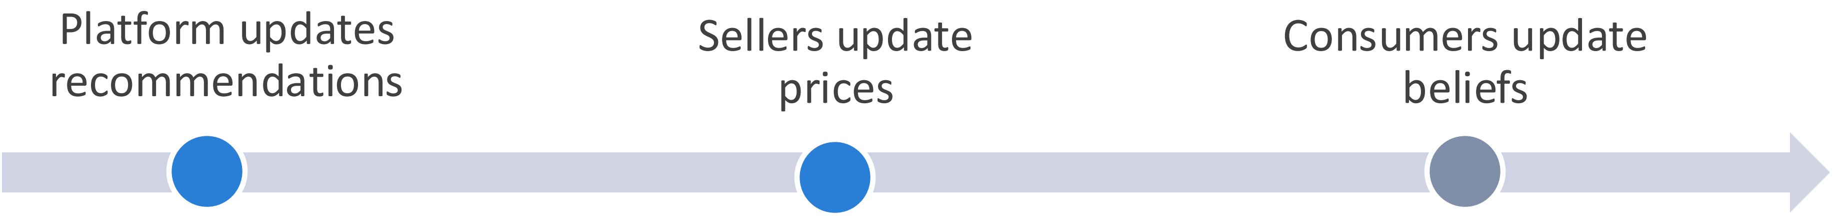

Baseline

- Subperiod uniform pricing (month, weekend-weekday, time before stay)

Use increasing levels of personalization

- Common Recommendations: Product features, competitive info
- Query Adjusted: + query features (ex nights, children)
- Personalize: + consumer observables (ex consumer country)
- Personalized Plus: + past transactions, tracked navigation data