Managing Congestion in Two-Sided Platforms: Online Rentals*

Caterina Calsamiglia† Laura Doval‡ Alejandro Robinson-Cortés§
Matthew Shum¶

November 7, 2026
Preliminary and incomplete: Please do not quote.

Abstract

In one-to-one matching platforms, congestion arises when customer demand outstrips supply for products with highly desired features. Motivated by key facts observed in a room-rental platform characterized by high congestion, we study platform information design when available products differ along vertical (quality) and horizontal (match-specific) dimensions, and each product can be matched to at most one user. When the platform makes the vertical attribute prominent, consumers agree on which product is best, creating congestion on popular items and leaving others unmatched. When the platform makes the horizontal attribute prominent, consumers disagree and self-sort across products, reducing congestion. We characterize how the trade-off between match utility and congestion depends on the relative value of vertical and horizontal attributes in consumer preferences. We then estimate preferences from our data and show that the current ranking algorithm is inefficient. We simulate counterfactual ranking algorithms which show that algorithms personalizing along horizontally-differentiated attributes can improve upon the baseline by raising match utility with less additional congestion, though the optimal design depends on the platform's tolerance for congestion.

*We thank Jeremy Fox, Andrey Fradkin, Yu-Wei Hsieh, Alessandro Iaria, Xuan Teng and seminar participants at UC Irvine, Boston University, Yale, Toulouse, Amazon, the 22nd CEPR/JIE Conference on Applied Industrial Organization, the BDS-Tinbergen Institute "Search and Markets" Conference, the 2023 Annual Workshop on Applied Economics, the UW Optimal Transport and Econometrics Conference, and Caltech's Applied Micro Lunch for helpful comments. Zion Chang provided research assistance.

†ICREA, IPEG and IZA.
‡Columbia Business School and CEPR.
§University of Exeter.
¶Amazon and Caltech.

1 Introduction

A central question in platform design is what information to present to users on search results pages. On a rental platform, should the page emphasize price and location, or should it emphasize whether the listing's landlord prefers tenants of a particular gender, age, or occupation? On a restaurant delivery app, should users first see price range and cuisine type, or whether the restaurant is vegan, halal, or open late? These choices are not neutral. They determine what users click on, which products receive attention, and ultimately who gets matched with whom.

This paper studies the problem for platforms that operate two-sided matching markets with capacity constraints. Each listing can be matched to at most one user, so when many users request the same listing, most requests fail. The platform's information design determines the degree of congestion in the market. Based on motivating facts we observe in the data, we illustrate the key role of customer preferences, namely the value they place on products' horizontal attributes vs. vertical attributes, for informing the platform's information design.

We study a room-rental platform in Barcelona in which user activity is highly congested. The platform ranks rooms identically for all users, concentrating attention on a small set of top-ranked rooms. We find that the attributes visible on the search page are primarily vertical, while horizontally-differentiated attributes (such as landlord preferences for tenant demographics) are hidden behind the click. As a result, the platform generates congestion without delivering the utility benefits of personalization.

Motivated by these facts, we develop a simple model of ranking design. When the platform emphasizes vertical attributes (quality, on which all consumers agree), consumers flock to the same high-quality products, overwhelming their capacity and leaving others unmatched. When it emphasizes horizontal attributes (match-specific fit, on which consumers disagree), consumers self-sort according to their idiosyncratic preferences, dispersing demand and reducing congestion. A *maximal contrast principle* emerges, whereby the platform maximizes matches by revealing the attribute along which consumer preferences are most opposed. The strength of the trade-off between match utility and congestion depends on whether preferences are primarily vertical or horizontal.

We estimate a two-stage model of user behavior to recover consumer preferences, and simulate counterfactual ranking algorithms that vary the degree of personalization. The platform's actual outcome lies below the frontier, revealing scope for improvement. We then investigate how the shape of this frontier depends on the degree of horizontal differentiation in preferences, as predicted by the theory. By progressively adjusting the weight of horizontal attributes in ranking personalization, we confirm that the utility-congestion trade-off weakens as rankings emphasize attributes on which users disagree rather than attributes on which they agree.

The paper proceeds as follows. Section 2 discusses related literature. Section 3 introduces the room rental platform we study and presents stylized facts and their implications for ranking design. Section 4 introduces the empirical model and reports estimation results of consumer preferences. Section 5 presents counterfactual scenarios illustrating how matching outcomes depend on the platform's personalization approaches. Section 6 concludes.

2 Related Literature

Our paper connects to three strands of the literature: the theoretical literature on information design and congestion in markets, empirical work on congestion in online platforms, and the literature on rankings and consumer search.

On the theoretical side, alleviating inefficiencies caused by congestion in two-sided markets has long been recognized as a goal of effective market design (Roth and Xing, 1997; Roth, 2008). Arnosti, Johari, and Kanoria (2021) study welfare losses from inefficient equilibrium search and screening in congested markets. Immorlica, Lucier, Manshadi, and Wei (2021) study how directing consumer search by limiting choice reduces congestion and the cannibalization of demand from low-type sellers by high-type ones. Ashlagi, Krishnaswamy, Makhijani, Saban, and Shiragur (2022) study the trade-off between expanding choice and increasing matching when designing recommendations in two-sided platforms. Mekonnen (2023) shows that congestion is worse when preferences are common and quality is revealed through directed search than in random matching. Closest related to our analysis is Armstrong and Zhou (2022), who study optimal information structures in a duopoly where consumers observe a private signal of their preferences. They show that the firm-optimal signal amplifies product differentiation, relaxing price competition, while the consumer-optimal signal dampens differentiation. [1] Our paper extends this logic to matching markets with capacity constraints, where the benefit of revealing horizontal attributes is even stronger because it directly reduces losses from congestion.

On the empirical side, inefficiencies arising from congestion have been documented in several online platforms. Fradkin (2017) studies search engine design and rejection rates in Airbnb. Horton (2019) studies how signaling features affect matching rates in online labor markets. Chen, Hsieh, and Lin (2023) study matching patterns in an online dating platform and, like us, point out that recommendation algorithms may lead to congestion. They prescribe an alternative algorithm based on the deferred acceptance mechanism and demonstrate via field experiments that it alleviates congestion. Bai, Chen, Liu, Mu, and Xu (2022) study how congestion and information frictions affect firm dynamics

[1] Choi and Shum (2024) derived a maximal contrast principle for platform design in a one-sided market without capacity constraints, showing that platforms prefer to reveal horizontally-differentiated attributes to reduce price competition.

and market efficiency on AliExpress. In dynamic settings, Lucier, Immorlica, Mobius, Slivkins, Goldstein, Hofman, Jaffe, and Rothschild (2026) study how improving consumer search technology affects market-level learning and welfare, showing that more informative search can degrade learning when the platform cannot distinguish idiosyncratic mismatches from quality deficiencies. Our paper contributes to this empirical literature by providing a structural empirical analysis linking the information shown to users to the degree of congestion on the platform.

The third strand studies the effects of algorithms on consumer search in online marketplaces (e.g., Koulayev, 2014; Hodgson and Lewis, 2020). The importance of rankings in online search has been widely documented (e.g., Ghose, Ipeirotis, and Li, 2014; Chen and Yao, 2017; Dang, Ursu, and Chintagunta, 2022; Farronato, Fradkin, and MacKay, 2023). A common finding in hotel booking platforms is that presenting the best alternatives at the top reduces search costs (De los Santos and Koulayev, 2017; Ursu, 2018). This literature primarily considers many-to-one markets in which product capacity is large enough for any number of users to be matched to the same product. In those markets congestion is not a first-order issue, leading to design implications that differ from ours. While we do not model optimal consumer search, our empirical findings echo this literature, especially regarding the importance of position effects.[2]

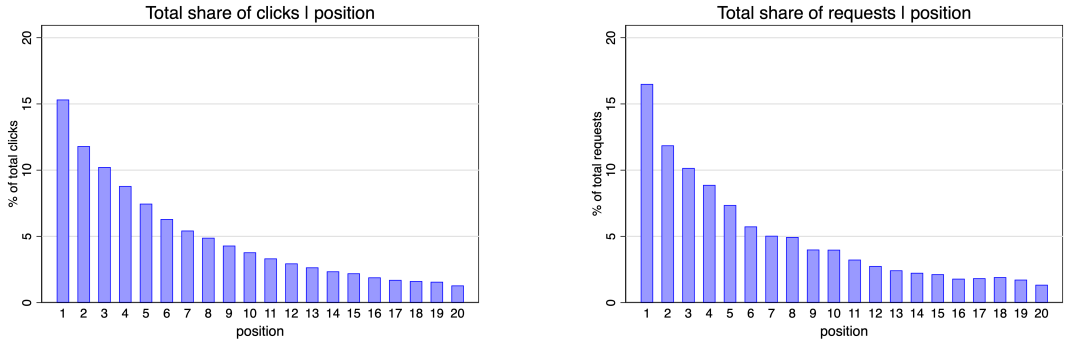
3 Stylized Facts from a Room Rental Platform and Implications for Ranking Design

Our study platform is one of the largest peer-to-peer rental markets in the world for rooms and apartments, active in major cities across Europe, the Americas, and Asia. During our study period (2018-2020), the platform targeted exchange students and workers seeking rentals of at least one month, differentiating itself from short-term platforms such as Airbnb.

The platform operates via a website or mobile app. Users submit searches and are shown available rooms listed by landlords. We observe the complete sequence of searches, clicks, and requests submitted by each user. The platform brings together renters and landlords but does not participate in the booking process itself. We have limited information on post-request outcomes, so our analysis focuses on clicking and requesting behavior.

In presenting room listings, the platform uses a common ranking across users rather than personalizing rankings to individual preferences. The platform ranks rooms using an algorithm based on a small number of ranking features, including the date the room was posted and whether the landlord has linked a bank account. These features are

[2]We model consumer clicking and requesting conditional on search. While this implies assuming that search behavior is invariant to the ranking algorithm in counterfactuals, we view this favorably compared to making stronger structural assumptions about consumer search behavior.



(a) Share of clicks

(b) Share of requests

Figure 1: Concentration of clicks and requests by position

identical across all potential consumers, implying that every user sees rooms in the same order at any given time.[3]

Our dataset consists of all searches conducted by a random sample of 10% of users from January 2018 to February 2020. We keep users who predominantly search in Barcelona and submitted at least one request during the sample period. The estimation sample contains 1,202 users, 45,462 rooms, and over 2 million search result impressions.

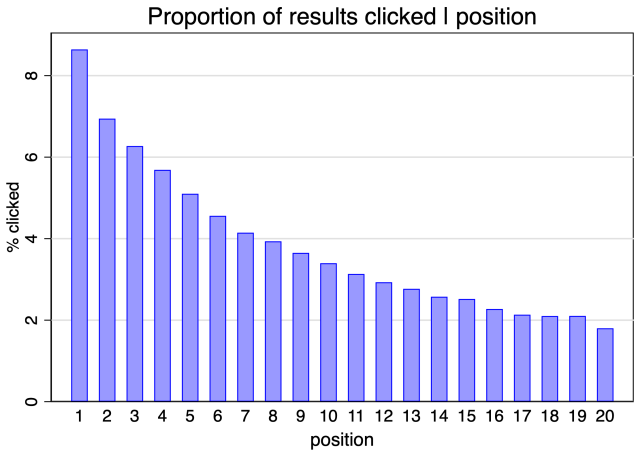
3.1 Evidence of Congestion and Platform Design

The data exhibit substantial congestion. Figure 1 shows that clicks and room requests are concentrated on rooms shown at top positions. Over 15% of all clicks and requests are made to rooms at the first position. This fraction falls to around 4% for the tenth position and asymptotes to around 2% for the twentieth position. Consumers exhibit "top-down" click behavior (favoring the top-ranked room listings).[4]

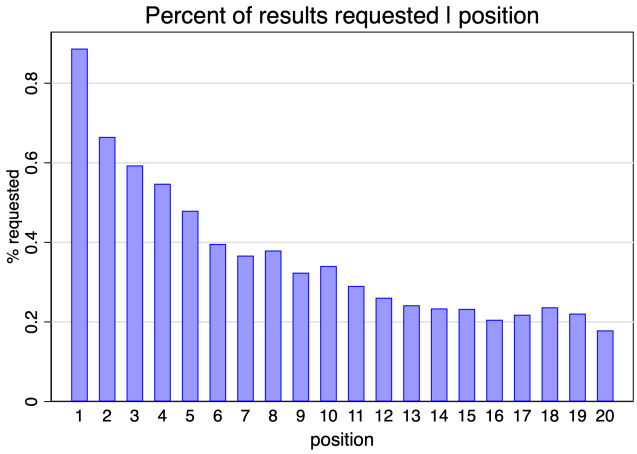
Figure 2 decomposes this further. The top two panels show that rooms in higher positions are more likely to be clicked on (2a) and requested (2b). However, the bottom panel (2c) shows that conditional on being clicked, rooms in higher positions are *not* more likely to be requested. A top-ranked room and the twentieth-ranked room are both requested with probability close to 0.10, conditional on having been clicked. This suggests that the ranking primarily affects whether users click on a room, but not whether they request it conditional on clicking. Since rooms must be clicked before they can be requested, the concentration in Figure 1 is driven by a listing's prominence (or consumers' attention to the top listing) more than preferences.

[3]Precisely, the ranking features determine a partial order, which is then randomized to produce a complete ranking. The *same* random draw is used to break ties across all users, implying that every user faces the same ranking. Users may apply filters (for neighborhoods, price ranges, etc.) which remove rooms from the search results, but the remaining rooms are ranked in the same order.

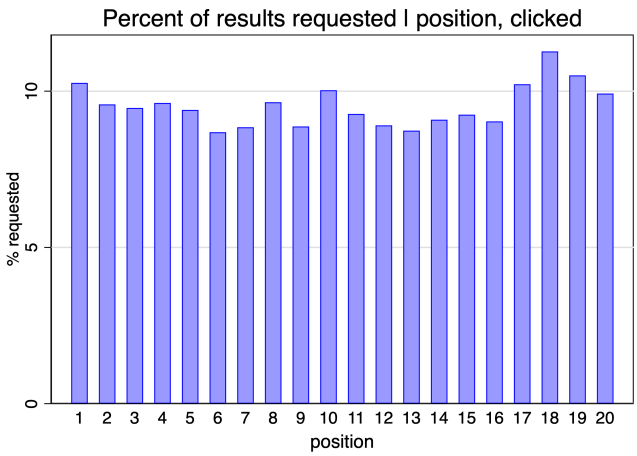
[4]In this paper, we do not study the rationality of consumers' top-down click strategies, but note that such behavior underlies practically all real-world search ranking algorithms.



(a) $\mathbb{P}(\text{Click})$



(b) $\mathbb{P}(\text{Request})$



(c) $\mathbb{P}(\text{Request} \mid \text{Click}{=}1)$

Figure 2: Probability of click and request by position

We now examine how this concentrated attention translates into unequal outcomes across rooms. Figure 3 presents Lorenz curves measuring the concentration of activity across rooms. The top 20% of rooms that appear most often in users' search results accumulate around 65% of total search results. This congestion worsens for clicks: the top 20% of rooms account for almost 70% of total clicks. For requests, the concentration is more pronounced: the top 20% of rooms account for essentially all requests in the data, while 80% of rooms never receive a single request. Since a room can be rented to at most one user, renters are better off requesting rooms that receive fewer requests from other renters, all else equal. The fact that 80% of rooms receive no requests suggests the possibility of allocative gains from reducing congestion.

Two features of the platform's design are relevant for understanding this concentration. First, as discussed above, all users see rooms in the same order at any given time. This makes the observed concentration of attention unsurprising: if all users see the same ordering, they are likely to focus on the same highly-ranked rooms.

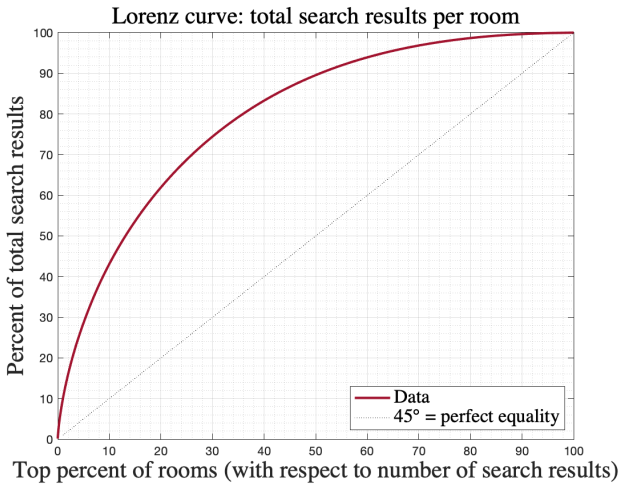
Second, the information available to users differs at different stages of the search process. The search results page displays price, number of flatmates, location, and a thumbnail. After clicking on a room, users see additional information including room amenities, the number of days the listing has been published, and whether the landlord has expressed a preference for tenants of a particular gender, age, or occupation. These latter features generate heterogeneity in preferences: for instance, rooms whose landlords prefer female tenants are likely more attractive to women than to men, and rooms preferring students appeal more to students than to workers.

Under the current design, substantial congestion emerges. Whether alternative ranking designs could reduce this congestion while maintaining or improving match quality is the question we address in the remainder of this paper.

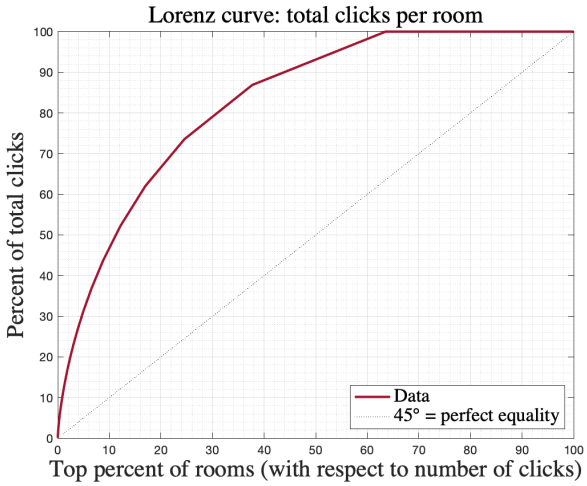
3.2 Implications for Ranking Design: an illustrative model

To understand how ranking design affects congestion, we develop an illustrative model with two key features motivated by the evidence above. First, rankings strongly influence demand, as consumers exhibit top-down click behavior. Second, products differ along vertical attributes that all consumers value similarly (quality) and horizontal attributes where consumer preferences differ (match-specific fit).

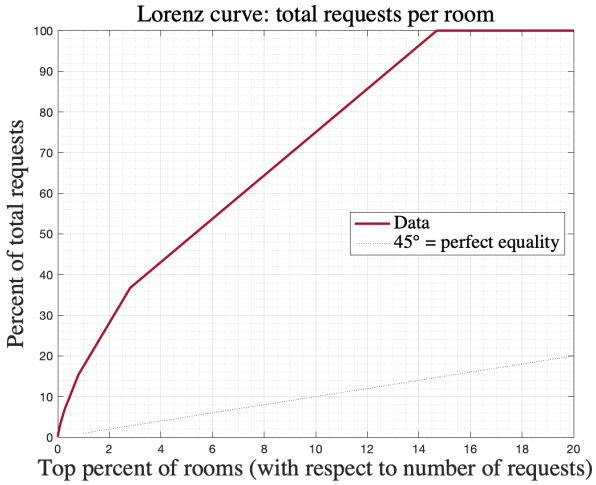
To illustrate the implications, consider the following matching market. There are two products ($j = 1, 2$), each with fixed capacity (supply) $1/2$, and a continuum of consumers with total mass 1. Products differ along two dimensions of attributes, which are unknown to consumers. The vertical dimension is quality v_j, which is equally valued by all consumers. The horizontal dimension is type $h_j \in \{A, B\}$. Consumer i has type $\theta_i \in$



(a) Searches



(b) Clicks



(c) Requests

Figure 3: Concentration of searches, clicks and requests across rooms

$\{A, B\}$ (drawn independently with equal probability) and receives a bonus $b > 0$ when matched to a product of her own type.

Consumer i's utility from product j is:

$$u_{ij} = \lambda \cdot v_j + (1 - \lambda) \cdot b \cdot \mathbf{1}[\theta_i = h_j], \tag{1}$$

where $\lambda \in (0, 1)$ governs the relative importance of vertical versus horizontal fit.

Without loss of generality, assume product 1 has higher vertical quality ($v_1 > v_2$), and has horizontal type A. Product 2 has horizontal type B, so $h_1 = A$ and $h_2 = B$. The quality advantage is $\Delta v = v_1 - v_2 > 0$. Note that consumers do not know this, and rely on the platform's ranking design to guide them to the product with the right features. Ex ante, consumers regard each product as equally likely to be the higher-quality one and equally likely to be of either horizontal type, independently.

Consumers observe their own type θ_i but not the products' characteristics (v_j, h_j). For convenience, we assume that the platform observes product attributes and can infer consumers' types θ_i $\forall i$ from observable characteristics, and can use this information to construct a (personalized) ranking to each consumer. Motivated by the strong top-down effects documented above, we assume every consumer requests the product shown at the top of their ranking, and requests only one product. If a mass d_j of consumers request product j, each matches with probability $\min\{1, (1/2)/d_j\}$. When demand is balanced ($d_1 = d_2 = 1/2$), everyone matches. When demand is imbalanced, the over-demanded product is rationed and consumers who fail to match receive utility zero.

The platform can personalize rankings along different dimensions. We consider three policies that represent distinct approaches to information design.

Policy V (personalize by vertical quality). The platform ranks product 1 first for all consumers because $v_1 > v_2$. The resulting demand is $d_1 = 1$, $d_2 = 0$, implying half of the consumers end up matched.

Policy H (personalize by horizontal fit). The platform ranks product 1 first for type-A consumers (since $h_1 = A$) and product 2 first for type-B consumers. Demand is $d_1 = 1/2$, $d_2 = 1/2$, so that all consumers are matched.

Policy VH (personalize by total utility). The platform ranks products according to each consumer's utility u_{ij}. Type-A consumers all see product 1 first, as both v and h favor it. Type-B consumers see product 2 first iff $(1 - \lambda)b > \lambda \Delta v$. Define $r = \lambda \Delta v / [(1 - \lambda)b]$. When $r > 1$, the vertical attribute dominates and all consumers see product 1 first. When $r < 1$, horizontal fit dominates and demand splits by type. The outcome therefore depends on which force is stronger. The fraction of matched

consumers is 1/2 if $r > 1$, and 1 if $r < 1$. We set aside the knife-edge case $r = 1$, in which type-B consumers are indifferent.

Under Policy H, demand always splits 50-50 because consumers sort on horizontal type. Under Policy V, all consumers request the high-quality product. Under Policy VH, matches equal 1/2 when $r > 1$ and 1 when $r < 1$. This illustrates how the platform should personalize along the attributes on which consumers most disagree to minimize congestion:

Proposition 1 (Maximal Contrast Principle)**.** *The platform maximizes the number of successful matches by personalizing along the horizontal attribute:*

$$\textit{Matches under } H = 1 \;\geq\; \textit{Matches under } VH \;\geq\; \textit{Matches under } V = 1/2.$$

Policy H is weakly better than VH, and strictly better when $r > 1$ (vertical dominates).[5]

3.2.1 Personalization and the Utility-Congestion Frontier

The three policies above represent extreme cases of full personalization. In practice, platforms may have imperfect information, face computational constraints, or deliberately inject randomness or diversity into results. We now consider a more general ranking algorithm indexed by $\alpha \in [0, 1]$, where:

- At $\alpha = 1$: the platform implements full personalization (each consumer sees her utility-maximizing product first)
- At $\alpha = 0$: the platform uses a fully random ranking (each consumer is equally likely to see either product first)
- For $0 < \alpha < 1$: with probability α each consumer sees her utility-maximizing product first; with probability $1 - \alpha$ she sees a randomly selected product first

For convenience, we focus on Policy VH (ranking by total utility), noting that the parameter $r = \lambda\Delta v/[(1-\lambda)b]$ nests the other two policies as limiting cases.[6] We examine how the degree of personalization α affects utility and congestion in each regime.

Case 1: $r > 1$ **(vertical dominates).** All consumers prefer product 1 regardless of type. Under full personalization ($\alpha = 1$), all consumers are directed to product 1: $d_1 = 1$, matches $= 1/2$. Under random ranking ($\alpha = 0$), demand splits 50-50 and everyone matches. For intermediate α, demand for product 1 consists of two groups: the fraction α who are personalized (all directed to product 1 since everyone prefers it), and half of

[5] This ranking of policies is unchanged if, instead of following the ranking mechanically, consumers choose rationally given the information the ranking conveys, with or without an independent preference for the first-ranked product.

[6] When r is large (quality dominates), Policy VH behaves like Policy V: all consumers prefer product 1 regardless of their horizontal type. When r is small (horizontal fit dominates), Policy VH behaves like Policy H: consumers sort purely by type.

the fraction $1-\alpha$ who receive random rankings. Thus $d_1(\alpha) = \alpha+(1-\alpha)/2 = (1+\alpha)/2$, and similarly $d_2(\alpha) = (1-\alpha)/2$. Total matches equal $1-\alpha/2$, which *decreases* in α.

Case 2: $r < 1$ **(horizontal fit dominates).** Type-A consumers prefer product 1 and type-B consumers prefer product 2. Under full personalization ($\alpha = 1$), each consumer is directed to her preferred product: demand splits 50-50 and everyone matches. Under random ranking ($\alpha = 0$), demand also splits 50-50 by symmetry. For intermediate α, demand is $d_1(\alpha) = d_2(\alpha) = 1/2$ for all α. Everyone matches regardless of α.

We have the following result (see Appendix A for proof).

Proposition 2 (Utility-Congestion Frontier)**.** *(i) When $r > 1$ (vertical dominates), the match rate decreases from 1 to 1/2 as α increases from 0 to 1. Average utility of matched consumers increases in α, because more consumers are directed to the high-quality product. A trade-off between match utility and congestion emerges.*
(ii) When $r < 1$ (horizontal fit dominates), the match rate equals 1 for all α, and average utility increases in α because more consumers are directed to the product matching their horizontal type. There is no trade-off: personalization raises utility without increasing congestion.

Proposition 2 traces a frontier in (utility, congestion) space that slopes upward when $r > 1$ and is flat when $r < 1$. The α-indexed policy is a natural way to move along it, but it need not be the most efficient one at a given match rate.[7] The model, while stylized, provides intuition for the analysis that follows, in a rich empirical setting with many rooms and many attributes.

4 Empirical Model and Estimation Results

We model user behavior in two stages corresponding to the two decisions users make: clicking and requesting. In the second stage (estimated first), a user who has clicked on a room decides whether to request it, having now observed all room features. In the first stage, a user decides which rooms to click on from the search results page, based on visible features and position.

Request utility. The utility of requesting search result s in search S by user i is

$$U_{is} = X_{1s}\beta_1 + X_{2s}\beta_2 + (X_{2s} \otimes Z_i)\beta_{XZ} + \varepsilon^r_{is}, \tag{2}$$

where X_{1s} contains characteristics visible on the search page (price, number of tenants, location), X_{2s} contains characteristics visible only after clicking (amenities, days since

[7] For instance, when $r > 1$, consider a market-level lottery that draws a single mode for the entire market, implementing Policy V with probability η and Policy H with probability $1-\eta$. Setting $\eta = \alpha$ yields the same match rate as the α-policy but strictly higher average utility for any $\alpha < 1$, because it replaces the α-policy's random-ranking component with Policy H: consumers who are not personalized receive their horizontal match with certainty rather than a randomly selected product.

publication, landlord preference match with user), Z_i is a vector of user characteristics, and ε^r_{is} is an iid extreme-value type 1 shock. The interaction $X_{2s} \otimes Z_i$ captures, for example, whether user i has the gender preferred by the landlord of listing s. Within each search, users request their highest-utility options among those clicked, yielding a rank-ordered logit with ties.[8]

Click propensity. The propensity to click on search result s is

$$I_{is} = f(pos_{is})\beta_{pos} + \mathbb{E}[U_{is} \mid X_{1s}, pos_{is}]\,\beta_U + pos_{is} \times \mathbb{E}[U_{is} \mid X_{1s}, pos_{is}]\,\beta_{pos\times U} + \varepsilon^k_{is}, \quad (3)$$

where $f(\cdot)$ is a polynomial in position with point masses at positions 1, 2, and 3, and the expected utility $\mathbb{E}[U_{is} \mid X_{1s}, pos_{is}]$ is constructed from first-stage estimates by projecting X_{2s} onto (X_{1s}, pos_{is}). The click model also takes the form of a rank-ordered logit with ties.

Estimation. There are two sequential stages in our estimation procedure, proceeding backwards. In the first step, we estimate the parameters of the request utility (2) via maximum likelihood. After estimating $\hat{\beta}_1$, $\hat{\beta}_2$, and $\hat{\beta}_{XZ}$, we project the covariates in X_{2s} onto X_{1s} and pos_{is} (including numerous interactions for flexibility) to construct the estimated expected request utility:

$$\hat{U}_i(X_s) = X_{1s}\hat{\beta}_1 + \hat{X}_{2s}\hat{\beta}_2 + (\hat{X}_{2s} \otimes Z_i)\hat{\beta}_{XZ}. \quad (4)$$

We then estimate the click propensity model (3) using $\hat{U}_i(X_s)$ as a covariate, interacted with position.

4.1 Sample description

Table 1 reports descriptive statistics for the estimation sample. On average, users are 29 years old, 51% are female, 35% are students, and 81% are employed. Among rooms, 27% specify a preferred gender for renters (23% prefer females, 4% prefer males) and 25% specify a preferred occupation (22% prefer not renting to students, 3% prefer students only). In terms of congestion, while all rooms appear at least once as a search result, only 64% receive at least one click and only 15% receive at least one request. On average, each user conducts 144 searches, browses over 1,700 search results, clicks on 75 rooms, and submits requests to 7 rooms.

Table 2 reports summary statistics of room characteristics and mean differences between clicked and requested rooms versus all search results.[9] On average, rooms cost around €400 per month. Clicked rooms are €12 cheaper than average and requested rooms

[8] Beggs, Cardell, and Hausman (1981); Allison and Christakis (1994). We also considered semi-parametric versions of this model (Fox, 2007; Yan and Yoo, 2017), and our findings are qualitatively similar.

[9] We provide characteristics at the level of search results since this is the unit of observation in the analysis.

Table 1: Sample description

Number of users (renters)	1,202
Average age	29.44
Are females (%)	51.4
Are students (%)	34.61
Are employed (%)	81.03
Number of rooms	45,462
Specify a preferred tenant gender (%)	27.2
Prefer female tenants (%)	23.13
Prefer male tenants (%)	4.07
Specify a preferred tenant occupation (%)	25.4
Prefer no student tenants (%)	22.7
Prefer not employed tenants (%)	2.68
Are clicked on at least once (%)	63.56
Are requested at least once (%)	14.69
Total number of searches	173,223
Average per user	144.11
Total number of search results	2,066,147
Average per user	1,718.92
Average per search	11.93
Total number of clicks	89,624
Average per user	74.56
Prob. search result is clicked	4.3%
Total number of requests	8,542
Average per user	7.11
Prob. click is requested	9.5%

Notes: Table reports sample sizes and descriptive statistics of the users (renters), rooms, searches, search results, clicks, and requests in the estimation sample.

are €15 cheaper, confirming price sensitivity. Renters are more likely to click on and request listings where their gender, age, and occupation match the landlord's stated preferences. Amenities vary across rooms, and renters tend to click and request rooms with balconies, exterior views, and heating. In terms of location, renters prefer rooms near the city center.[10]

[10]The city of Barcelona proper has 10 districts, which we group into five by merging together Sarrià-Sant Gervasi and Gràcia, Ciutat Vella and Saint Martí, Sants-Montjuic and Les Corts, and Horta-Guinardó, Nou Barris and Sant Andreu (into North Barcelona). We also pool rooms within 30km of the city center but outside the 10 districts into "Greater Barcelona."

Table 2: Summary statistics and mean differences

	mean	sd	median	min	max	clicked − all	requested − all	requested − clicked
Price	400.591	102.643	390	0	750	-12.377***	-15.288***	-3.740***
Days since first published	100.669	166.695	14	0	744	5.648***	-5.017***	-11.495***
User has gender preferred by landlord	0.812	0.391	1	0	1	0.008***	0.034***	0.029***
User has age preferred by landlord	0.654	0.476	1	0	1	0.006***	0.026***	0.022***
User has occupation preferred by landlord	0.902	0.297	1	0	1	0.006***	0.013***	0.008**
Has AC	0.241	0.428	0	0	1	-0.004***	-0.017***	-0.014***
Has exterior view	0.492	0.500	0	0	1	0.014***	0.022***	0.009
Has balcony	0.417	0.493	0	0	1	0.015***	0.040***	0.027***
Has heating	0.470	0.499	0	0	1	0.004**	0.011**	0.008
In Ciutat Vella or Saint Martí	0.198	0.399	0	0	1	0.003**	0.043***	0.044***
In L'Eixample	0.263	0.440	0	0	1	0.016***	0.036***	0.022***
In North of Barcelona	0.124	0.329	0	0	1	-0.022***	-0.043***	-0.025***

Notes: the table reports selected summary statistics of the variables used in the model. The last three columns report mean differences between (a) clicked and all search results; (b) requested and all search results; (c) requested and clicked search results. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Variables *Price* and *Days since first published* are winsorized on the right at the 99th percentile.

4.2 Estimation results

Table 3 reports the parameter estimates of the request utility. We estimate a negative and statistically significant price sensitivity across all specifications. The coefficients are only identified up to scale. To interpret magnitudes, we divide estimates by the absolute value of the price coefficient, yielding willingness-to-pay in euros. The largest coefficient is on "user has gender preferred by landlord" (around €520 in the preferred specification). This is the type of horizontally-differentiated attribute that the theoretical framework identifies as most valuable for reducing congestion. Having a balcony is valued at €113 and heating at €102. Beyond gender match and price, only a few characteristics have statistically significant effects on request probabilities, among them balcony, heating, and location in North Barcelona.

Table 3: Parameter estimates of request utility

	(1)	(2)	(3)	(4)
Price	-0.0007**	-0.0007**	-0.0008***	-0.0009***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Days since first published		-0.0003**	-0.0003**	-0.0002**
		(0.0001)	(0.0001)	(0.0001)
User has gender preferred by landlord		0.4865***	0.4906***	0.4818***
		(0.0571)	(0.0572)	(0.0574)
User has age preferred by landlord		-0.0068	0.0027	0.0095
		(0.0536)	(0.0538)	(0.0539)
User has occupation preferred by landlord		0.0771	0.0769	0.0799
		(0.0851)	(0.0853)	(0.0855)
Has balcony			0.1149***	0.1049***
			(0.0381)	(0.0382)
Has heating			0.0892**	0.0945**
			(0.0395)	(0.0397)
In North of Barcelona				-0.5561***
				(0.1947)
Num Searches	52,737	52,737	52,737	52,737
Num Results	89,624	89,624	89,624	89,624
Pseudo-R2	0.0005	0.0087	0.0111	0.0150
Log-likelihood	-5112.57	-5070.45	-5058.32	-5038.15

Notes: Table reports rank-ordered logit parameter estimates (selected coefficients shown). Estimation via maximum likelihood. Standard errors in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Full results available upon request.

Table 4 reports the parameter estimates of the click propensity. Given the stylized facts presented earlier, position is the dominant driver of clicking: being shown one position lower reduces click propensity by approximately €250 in equivalent terms, and being shown at position 1 has an additional large positive effect. Expected utility has a positive but statistically insignificant effect (with bootstrapped standard errors that account for the generated regressor). The ranking determines clicks through attention, but conditional on clicking, request decisions are driven by room features rather than position.

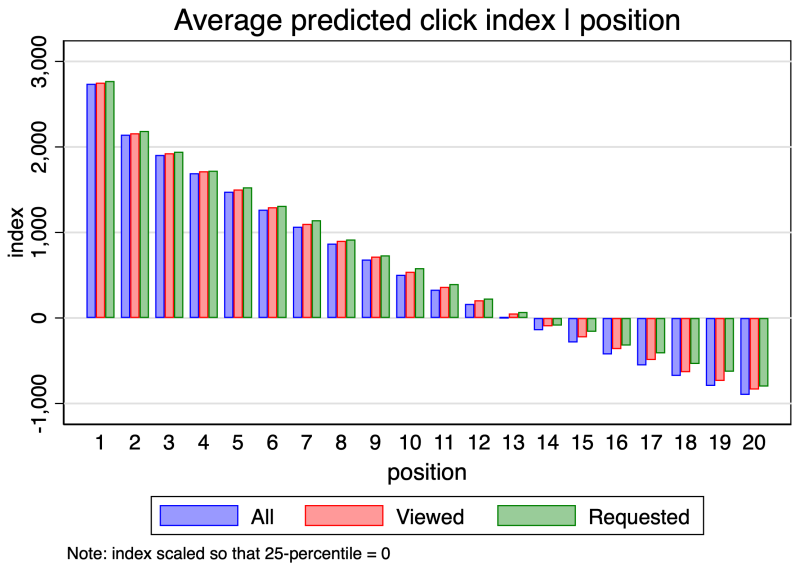
Table 4: Parameter estimates of click propensity

	(1)	(2)	(3)	(4)
Position	-0.0538***	-0.0834***	-0.0698***	-0.0698***
	(0.0010)	(0.0029)	(0.0052)	(0.0119)
Position^2		0.0016***	0.0010***	0.0010***
		(0.0002)	(0.0002)	(0.0002)
1(Position = 1)			0.0991***	0.0987***
			(0.0216)	(0.0217)
$\mathbb{E}[u\|I]$				0.3027
				(0.2806)
$\mathbb{E}[u\|I]$ × Position				0.0320
				(0.0217)
Num Searches	173,223	173,223	173,223	173,223
Num Results	2,066,147	2,066,147	2,066,147	2,066,147
Pseudo-R2	0.0143	0.0147	0.0149	0.0161
Log-likelihood	-1.644e+05	-1.644e+05	-1.643e+05	-1.641e+05

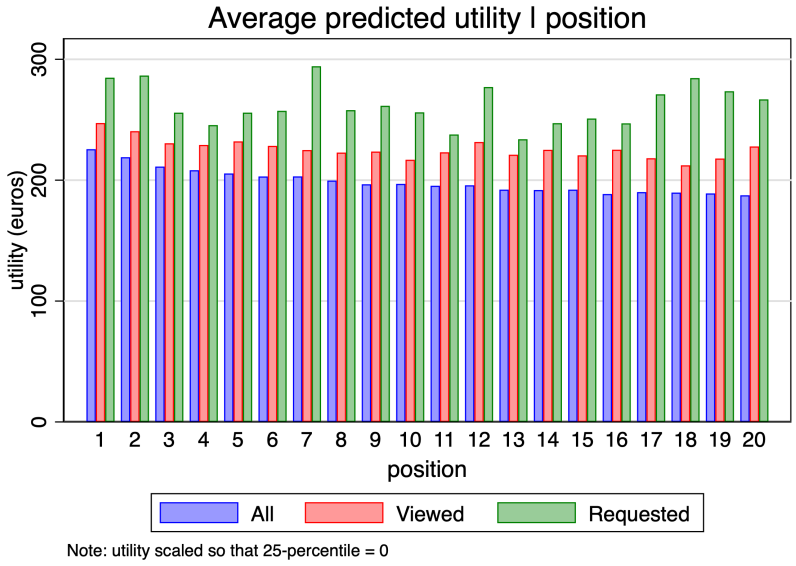
Notes: Table reports rank-ordered logit parameter estimates. Bootstrapped standard errors in parentheses (1,000 draws); *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Figure 4 plots the estimated click propensity and request utility against position. Click propensity decreases monotonically with position for all rooms, while predicted request utility is flat in position, especially for rooms that are actually requested. This is consistent with the pattern in the raw data: position operates through the attention channel, not through preferences.

We summarize this section by considering how the estimation results map to the earlier model (Section 3.2). The estimated preferences are primarily vertical: price and position dominate clicking behavior, while the horizontal gender-match variable dominates requesting behavior only after clicking. To quantify the balance between vertical and horizontal preferences, recall that we defined $r = \lambda\Delta v/[(1-\lambda)b]$ as the ratio of vertical to horizontal prominence in consumer preferences. Our estimates suggest $r > 1$: the gender-match coefficient (the strongest horizontal component) has a willingness-to-pay of approximately €520, which is substantial but smaller than the overall range of variation in predicted utility across rooms driven by price, location, and amenities. This places the market in the vertically-dominated case of the framework, where increasing personalization (α) raises both match utility and congestion because all users agree on the best rooms (Proposition 2). We next turn to the counterfactual scenarios.



(a) Click index



(b) Request utility

Figure 4: Estimated click index and request utility by position

5 Counterfactual scenarios

Based on the estimation results, we simulate counterfactual ranking algorithms and examine their effects on congestion and match utility, mapping to the discussion in Section 3.2. We first establish a baseline frontier by varying how much of each user's *total* estimated utility is incorporated into the ranking (Section 5.1). We then implement the maximal contrast principle directly by ranking on the horizontal component of utility only (Section 5.2). Finally, we perform a garbling exercise that destroys all cross-user preference correlation, showing that the trade-off largely disappears when preferences are purely horizontal (Section 5.3). Appendix C reports additional results for an alternative specification that allows for latent consumer clusters estimated from search filter behavior. Because users in different clusters rank the same rooms differently, full personalization no longer pushes all users toward the same set of rooms, and the resulting frontier becomes more concave than the baseline, in the direction Proposition 2 would suggest.

In computing the counterfactuals, we maintain several simplifying assumptions. We hold fixed 1) the composition of search results within each search and vary only the order in which rooms are shown; 2) the number of clicks and requests each user makes. Predictions use only the estimated deterministic indices, without drawing utility shocks. This approach is conservative in that it relies solely on observable heterogeneity.[11]

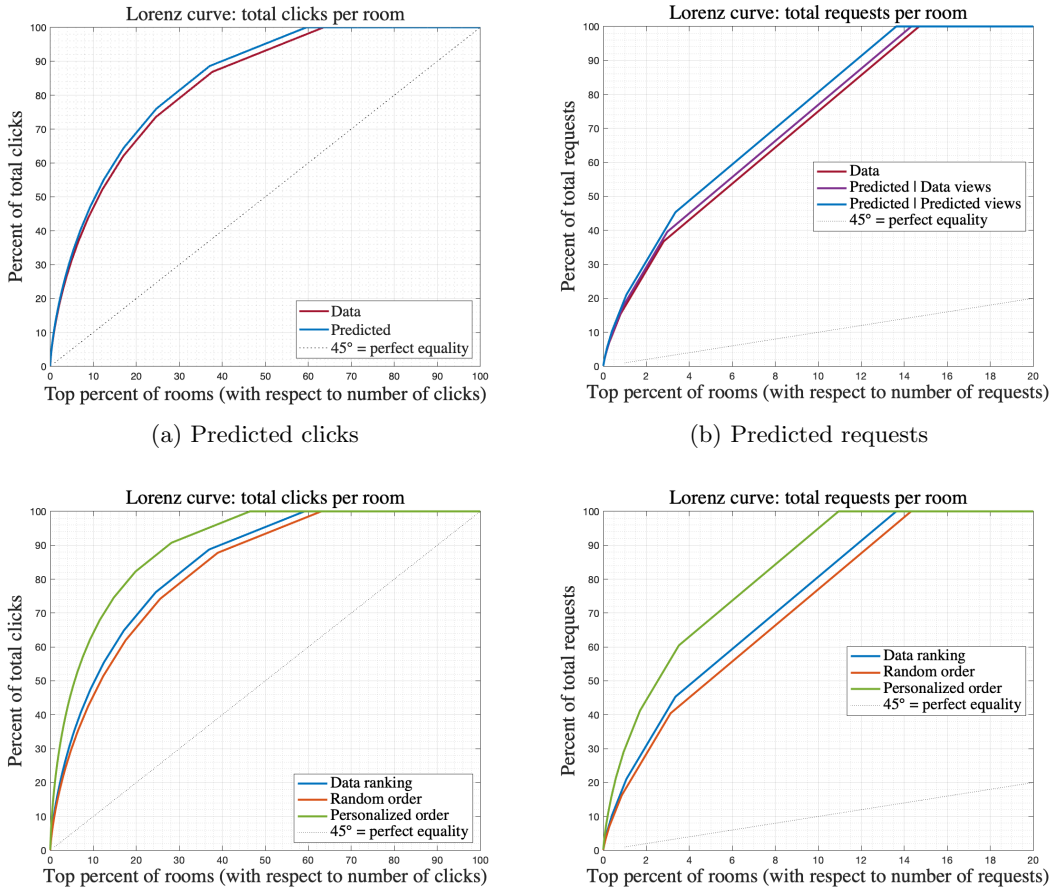
5.1 Baseline: ranking algorithms and the utility-congestion frontier

We consider a family of ranking algorithms indexed by $\alpha \in [0, 1]$: at $\alpha = 1$, rooms are ranked for each user according to her estimated utility (full personalization); at $\alpha = 0$, rooms are ranked in a random order that differs across users. This α is the personalization parameter from Section 3.2.1: higher α means the ranking incorporates more preference information (which in this market is vertically dominated).

We measure congestion using Lorenz curves and Gini coefficients of the distribution of clicks and requests across rooms, since realized matches are not observed. We measure utility by the average estimated utility of rooms that users are predicted to click or request.

Figure 5 presents Lorenz curves for several scenarios. In the top row, the predicted and observed Lorenz curves are close to one another. The bottom row shows that full personalization ($\alpha = 1$, green) substantially increases congestion relative to the data, while random order ($\alpha = 0$, orange) slightly decreases it. Full personalization steers all users toward the same top-utility rooms because preferences are primarily vertical in

[11] Our findings are also qualitatively robust to semiparametric methods that relax the logit assumption (Fox, 2007; Yan and Yoo, 2017).



(a) Predicted clicks

(b) Predicted requests

(c) Counterfactual clicks

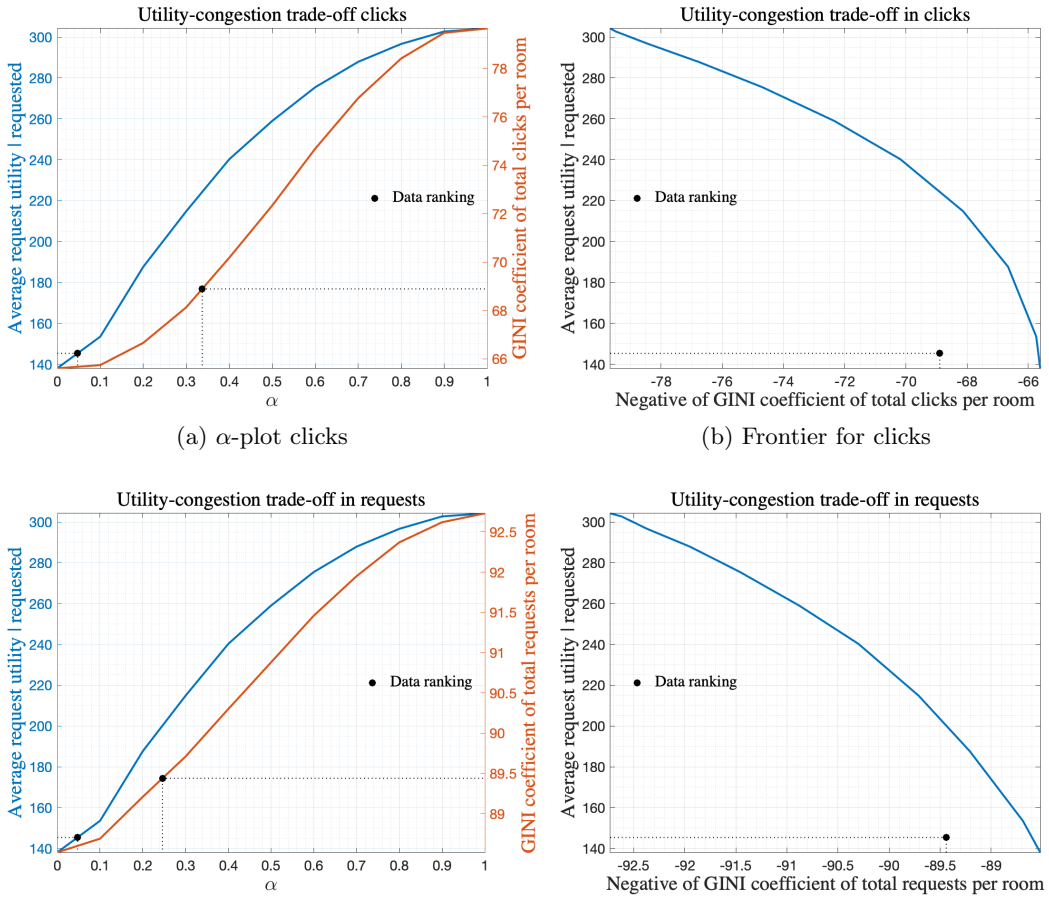
(d) Counterfactual requests

Figure 5: Lorenz curves of clicks and requests across rooms. Top row: data (red) vs. model (blue). Bottom row: full personalization (green) and random order (orange) vs. data.

this market. From a congestion perspective, the current algorithm is preferable to full personalization.

Figure 6 traces the frontier as α varies continuously from 0 to 1. In the left panels, the blue curve shows average utility increasing with α while the red curve shows congestion also increasing. The right panels plot the achievable (utility, congestion) pairs.

The black dot marks the platform's actual outcome. It lies below the frontier for both clicks and requests, implying that the platform is inefficient: we could achieve the same utility with less congestion, or higher utility at the same congestion, by choosing a different α. The platform sits below the frontier because its algorithm generates congestion through the common ranking (everyone sees the same order) without delivering the utility benefit of personalization (the ranking is not based on individual preferences). The α-family dominates because each member either randomizes (reducing congestion) or personalizes (raising utility), while the platform's algorithm does neither effectively.



(a) α-plot clicks

(b) Frontier for clicks

(c) α-plot requests

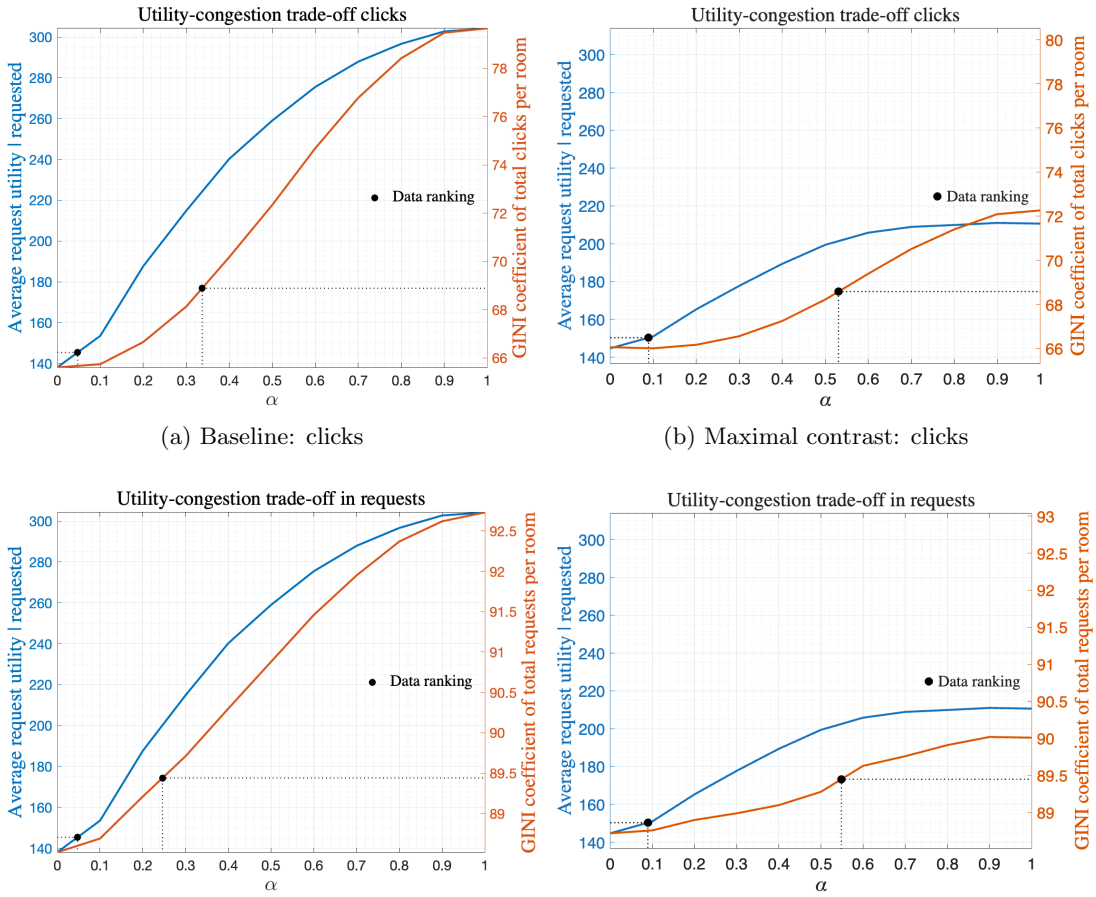
(d) Frontier for requests

Figure 6: Utility and congestion trade-off as α varies from 0 to 1

5.2 The maximal contrast counterfactual

To study the maximal contrast principle directly, we construct a ranking algorithm based exclusively on the *horizontal* component of estimated utility. Specifically, we rank rooms within each search according to the interaction terms $(X_{2s} \otimes Z_i)\hat{\beta}_{XZ}$ in the request utility, which capture whether user i matches the landlord's preferences for tenant gender, age, and occupation. Everything common across users (price, location, amenities) is dropped from the ranking key. Under this algorithm, a female student sees rooms whose landlords prefer female students at the top, while a male worker sees a different ordering of the same results. The algorithm personalizes, but exclusively along the dimension where users disagree.

As in the baseline, we mix this ranking with a uniformly random permutation in proportion $\alpha \in [0, 1]$. At $\alpha = 0$ both exercises produce the same random order, providing a common starting point and a consistency check. Sweeping α from 0 to 1 traces out a second utility-congestion frontier to compare against the baseline.



(a) Baseline: clicks

(b) Maximal contrast: clicks

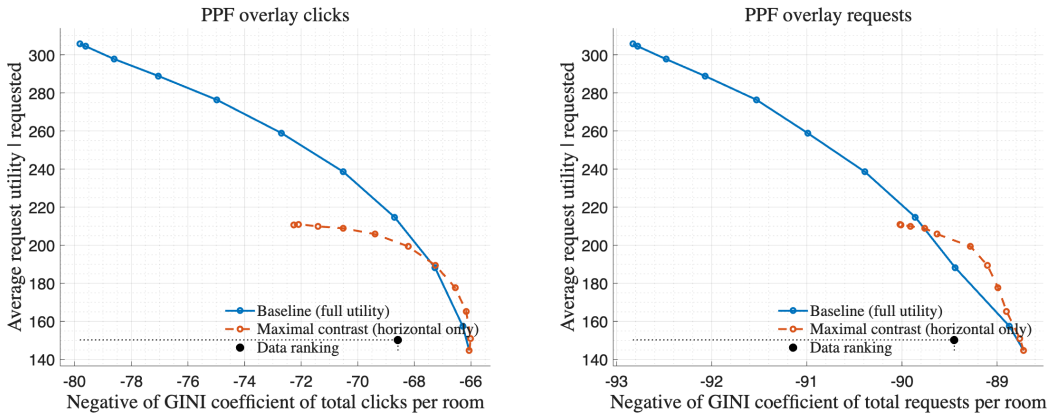
(c) Baseline: requests

(d) Maximal contrast: requests

Figure 7: Utility and congestion as α varies from 0 to 1. Left column: baseline ranking (full estimated utility). Right column: maximal-contrast ranking (horizontal component only). Top row: clicks; bottom row: requests. Left axis: average utility in euros; right axis: Gini coefficient. Common axes within each row.

Figure 7 plots the two exercises side by side against α, with utility and Gini on common axes so the columns can be compared directly. The baseline ranking (panels a and c) pulls utility steeply upward as α increases because it sorts on total utility, which is vertically dominated: higher α steers more users toward the same high-quality rooms, and the Gini coefficient rises in tandem. The maximal-contrast ranking (panels b and d) raises utility more slowly and flattens out past $\alpha \approx 0.6$, having exhausted the available horizontal differentiation. In return the Gini rises far more slowly, and in requests barely moves at all.

Figure 8 overlays the two frontiers in utility-congestion space. The maximal-contrast frontier is the flatter of the two: it delivers less utility at any given level of personalization but also generates less congestion. Neither frontier dominates everywhere, as they cross. At the right-end of the graphs, corresponding to high randomization (low-α), the maximal-contrast ranking is more efficient: at $\alpha = 0.2$ it raises the utility of



(a) Clicks (b) Requests

Figure 8: Utility-congestion frontiers for the baseline and maximal-contrast rankings, overlaid. Each point corresponds to one value of $\alpha \in [0, 1]$; the two curves share a common origin at $\alpha = 0$. The black dot marks the platform's actual outcome.

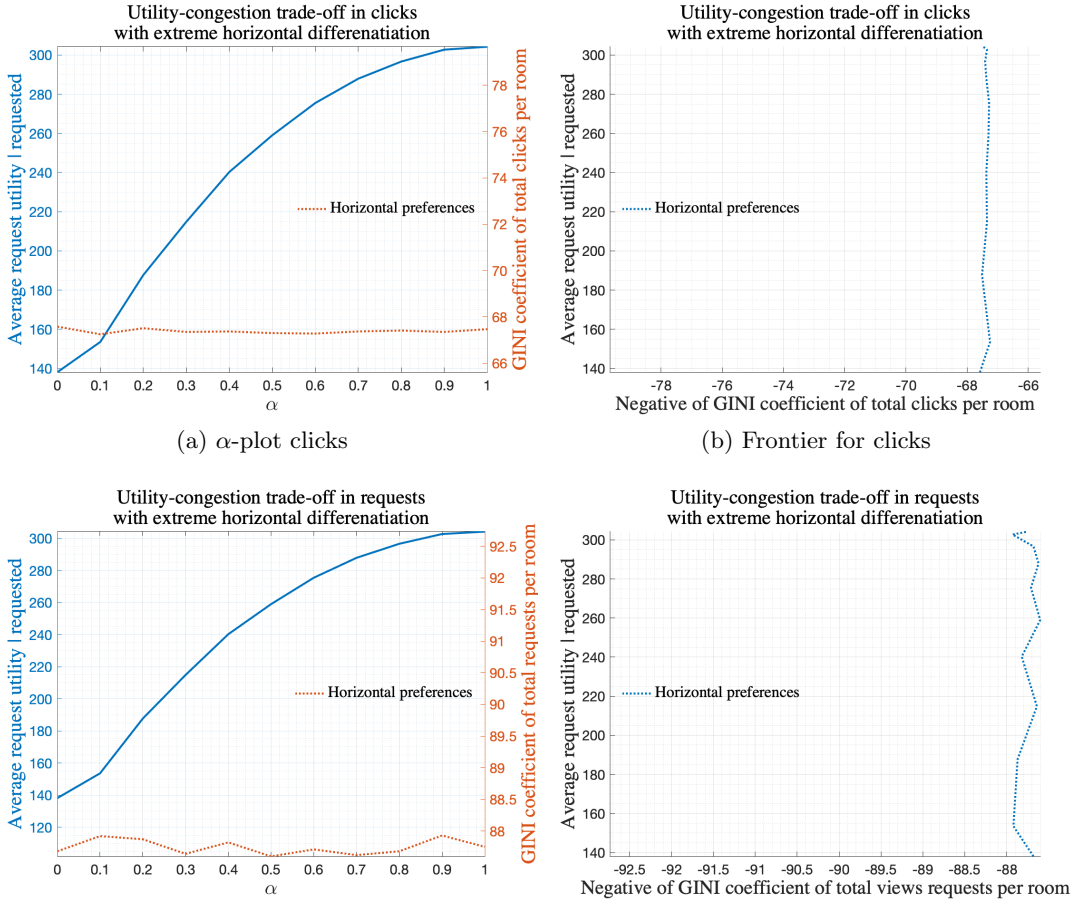
a requested room from €144.68 to €165.30 while barely moving click congestion (Gini 66.06 $\rightarrow$ 66.17), whereas the baseline held to the same congestion would deliver only €150.79, a gain of €14.51. The two click frontiers cross at $\alpha \approx 0.42$; past that point the baseline dominates. For requests the crossing comes later, around $\alpha \approx 0.72$, and the maximal-contrast ranking is the more efficient choice over most of the range.

The lower congestion of the maximal-contrast ranking echoes Proposition 1: personalizing along the horizontal dimension disperses demand rather than concentrating it. But horizontal differentiation in this market is limited. Most of what makes a room worth requesting is common to all users, so a ranking that leaves out the vertical component eventually forgoes too much match quality, which is why the two frontiers cross.

The crossing implies that the best algorithm depends on how much congestion the platform is willing to tolerate. Near the random end, where the platform's current outcome lies, the maximal-contrast ranking is the more efficient way to add personalization. Starting from the current inefficient position (the black dot), an improvement path involves moving toward moderate personalization, where the maximal-contrast ranking dominates. When the platform is willing to accept higher congestion in exchange for higher average match utility, the baseline eventually dominates. The platform's actual outcome lies below both frontiers, so the current algorithm is inefficient regardless of which objective applies.

5.3 Garbling: the purely horizontal limit

To isolate the role of preference structure, we now perform an artificial exercise that eliminates the vertical component entirely. In the model, the trade-off between utility and congestion depends on whether preferences are primarily vertical or horizontal. To



(a) α-plot clicks

(b) Frontier for clicks

(c) α-plot requests

(d) Frontier for requests

Figure 9: Utility and congestion trade-off under garbled preferences (extreme horizontal differentiation)

examine the opposite extreme, we perform a "garbling" exercise: we randomly relabel room identities across searches, destroying the correlation of preferences across users while preserving the utility distribution within each search. Under garbling, preferences become purely horizontally differentiated in the sense that knowing one user's preferred room tells you nothing about another user's. This corresponds to the horizontally-dominated case of the framework.

Figure 9 shows the result. Under garbled preferences, personalization no longer increases congestion. The frontier becomes approximately vertical: as α increases, utility rises but congestion stays flat. This is consistent with Proposition 2: when preferences are horizontal (the $r < 1$ case), personalization does not create congestion because different users' favorites are different. The trade-off observed in Figure 6 is a consequence of vertical preference structure, not an inherent feature of personalization.

6 Conclusion

This paper shows that capacity-constrained matching platforms face a trade-off between match quality and congestion when personalizing search rankings to emphasize vertical attributes (on which consumer preferences agree), but this trade-off weakens or disappears when emphasizing horizontal attributes (on which consumers disagree) instead. This trade-off between match utility and congestion depends on the relative importance of vertical and horizontal differentiation in preferences: when quality dominates, there is a real cost to dispersing demand (consumers are steered away from the best product); when horizontal fit dominates, dispersion is costless.

The empirical analysis demonstrates that this mechanism operates quantitatively on a real platform. The current algorithm ranks rooms identically for all users based on attributes (recency, verification) that function as common quality signals. This generates congestion without delivering utility gains from personalization. We show that a family of personalized algorithms traces out a utility-congestion frontier above the platform's current outcome, but that this frontier itself exhibits a trade-off because estimated preferences are vertically dominated.

We then show how this trade-off weakens as the ranking algorithm emphasizes more horizontal attributes. The maximal contrast counterfactual, which personalizes rankings along the demographic-match component of utility only, traces out a flatter frontier that is more efficient than the baseline at low levels of personalization; the two frontiers cross at intermediate values of α, after which the baseline dominates. The garbling exercise, which destroys cross-user preference correlation entirely, shows that the trade-off largely disappears when preferences are purely horizontal.

The maximal contrast principle applies broadly to any capacity-constrained matching platform. Dating platforms, labor market platforms, and school choice systems all face the same tension between showing users their "best" options and spreading demand across available supply. A general prescription emerging from this paper is that the platform's ranking design should emphasize features that users value heterogeneously rather than the features that users value equally.

References

Allison, P. D. and N. Christakis (1994): "Logit models for sets of ranked items," in *Sociological Methodology*, ed. by P. V. Marsden, Blackwell, vol. 24, 123–126. 12

Armstrong, M. and J. Zhou (2022): "Consumer Information and the Limits to Competition," *American Economic Review*, 112, 534–577. 3

Arnosti, N., R. Johari, and Y. Kanoria (2021): "Managing congestion in matching markets," *Manufacturing & Service Operations Management*, 23, 620–636. 3

Ashlagi, I., A. K. Krishnaswamy, R. Makhijani, D. Saban, and K. Shiragur (2022): "Assortment planning for two-sided sequential matching markets," *Operations Research*, 70, 2784–2803. 3

Bai, J., M. X. Chen, J. Liu, X. Mu, and D. Y. Xu (2022): "Stand Out from the Millions: Market Congestion and Information Friction on Global E-Commerce Platforms," *mimeo*. 3

Beggs, S., S. Cardell, and J. Hausman (1981): "Assessing the potential demand for electric cars," *Journal of Econometrics*, 17, 1–19. 12

Bonhomme, S., T. Lamadon, and E. Manresa (2022): "Discretizing Unobserved Heterogeneity," *Econometrica*, 90, 625–643. 29

Chen, K.-M., Y.-W. Hsieh, and M.-J. Lin (2023): "Reducing Recommendation Inequality via Two-Sided Matching: A Field Experiment of Online Dating," *International Economic Review*, 64, 1383–1421. 3

Chen, Y. and S. Yao (2017): "Sequential search with refinement: Model and application with click-stream data," *Management Science*, 63, 4345–4365. 4

Choi, M. and M. Shum (2024): "The Maximal Contrast Principle for Platform Design," Working paper, University of California Irvine and California Institute of Technology. 3

Dang, C. I., R. Ursu, and P. K. Chintagunta (2022): "Product Revisits and Consumer Consideration Sets," Available at SSRN: https://ssrn.com/abstract=3626451. 4

De los Santos, B. and S. Koulayev (2017): "Optimizing click-through in online rankings with endogenous search refinement," *Marketing Science*, 36, 542–564. 4

Farronato, C., A. Fradkin, and A. MacKay (2023): "Self-Preferencing at Amazon: Evidence from Search Rankings," *AEA Papers and Proceedings*, 113, 239–243. 4

Fox, J. T. (2007): "Semiparametric estimation of multinomial discrete-choice models using a subset of choices," *The RAND Journal of Economics*, 38, 1002–1019. 12, 18

Fradkin, A. (2017): "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb," Available at SSRN: https://ssrn.com/abstract=2939084. 3

Ghose, A., P. G. Ipeirotis, and B. Li (2014): "Examining the impact of ranking on consumer behavior and search engine revenue," *Management Science*, 60, 1632–1654. 4

Hodgson, C. and G. Lewis (2020): "You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search," Cowles Foundation Discussion Papers 2246, Cowles Foundation for Research in Economics, Yale University. 4

Horton, J. J. (2019): "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Management Science*, 65, 3518–3540. 3

Immorlica, N., B. Lucier, V. Manshadi, and A. Wei (2021): "Designing approximately optimal search on matching platforms," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, 632–633. 3

Kaufman, L. and P. J. Rousseeuw (1990): *Finding Groups in Data: An Introduction to Cluster Analysis*, Wiley. 29

Koulayev, S. (2014): "Search for differentiated products: identification and estimation," *The RAND Journal of Economics*, 45, 553–575. 4

Lucier, B., N. Immorlica, M. Mobius, A. Slivkins, D. G. Goldstein, J. M. Hofman, S. Jaffe, and D. M. Rothschild (2026): "Agentic Markets: Equilibrium Effects of Improving Consumer Search," ArXiv:2603.25893. 4

Mekonnen, T. (2023): "Random vs. Directed Search for Scarce Resources," *mimeo*. 3

Roth, A. E. (2008): "What Have We Learned from Market Design?" *The Economic Journal*, 118, 285–310. 3

Roth, A. E. and X. Xing (1997): "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, 105, 284–329. 3

Ursu, R. M. (2018): "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions," *Marketing Science*, 37, 530–552. 4

Yan, J. and H. I. Yoo (2017): "Semiparametric Estimation of the Random Utility Model with Rank-Ordered Choice Data," Working papers, Durham University Business School. 12, 18

A Proofs

Proof of Proposition 2. Case (i): At personalization level α, demand for product 1 is $(1+\alpha)/2$. Each of the $(1+\alpha)/2$ consumers requesting product 1 is matched with probability $(1/2)/[(1+\alpha)/2] = 1/(1+\alpha)$. Their expected utility conditional on matching is $\lambda v_1 + (1-\lambda)b/2$ (they chose the high-quality product; their horizontal match is random since both types request it). The $(1-\alpha)/2$ consumers randomly assigned to product 2 are all matched (demand does not exceed capacity 1/2). Their expected utility is $\lambda v_2 + (1-\lambda)b/2$. As α increases, more consumers target the better product, raising conditional utility of those matched, but the match rate $1-\alpha/2$ falls.

Case (ii): Since type-A and type-B consumers prefer different products, both personalization and randomization split demand evenly. The match rate is 1 at every α. Utility increases in α because personalized consumers are matched to the product of their own horizontal type (utility gain $(1-\lambda)b$), while randomly assigned consumers receive the horizontal bonus only with probability 1/2. □

B Additional Tables

Table 5 reports the normalized coefficients of the request utility (divided by the absolute value of the price coefficient), allowing interpretation in euros.

Table 5: Normalized parameter estimates of request utility (in euros)

	(1)	(2)	(3)	(4)
Price	-1.00	-1.00	-1.00	-1.00
Missing number of tenants		-130.93	-96.72	-85.94
Number of tenants		0.53	1.13	-4.02
Days since first published		-0.40	-0.32	-0.26
User has gender preferred by landlord		708.38	603.54	520.57
User has age preferred by landlord		-9.85	3.38	10.28
User has occupation preferred by landlord		112.34	94.65	86.36
Has doorman			-113.70	-120.95
Has dishwasher			89.54	78.72
Has terrace			-23.48	-14.73
Has AC			10.03	16.12
Is smoker friendly			2.10	-2.08
Has elevator			65.10	44.97
Has exterior view			-31.93	-17.42
Has TV			-14.12	-13.54
Has balcony			141.31	113.34
Has heating			109.70	102.14
In Saint-Gervasi or Gracia				-76.72
In Ciutat Vella or Saint Marti				-109.53
In Eixample				-61.97
In North of Barcelona				-600.84
In Sants-Montjuic of Les Corts				-183.20

Notes: Coefficients normalized to be in euros (divided by absolute value of price coefficient). The category "In Greater Barcelona" is omitted.

Table 6: Normalized parameter estimates of click propensity (in euros)

	(1)
Position	-249.14
Position^2	3.74
1(Position = 1)	352.42
1(Position = 2)	-3.70
1(Position = 3)	-9.68
$\mathbb{E}[u\|I]$	1.00
$\mathbb{E}[u\|I]$ × Position	0.11

Notes: Coefficients normalized to be in euros (divided by absolute value of the normalized utility coefficient).

C Cluster-specific preferences

The baseline frontier in Figure 6 treats all users as having the same preference parameters (up to the observable interactions $X_{2s} \otimes Z_i$). To capture richer heterogeneity, we estimate cluster-specific preferences using observed search behavior.

The intuition behind our clustering approach is that users who have different preferences search differently, for example by specifying different search criteria when using filters. In our platform, a number of search filters are available for price range, location, amenities, and gender preferences. While we have no direct information on the filters users employ, we can infer filtering behavior: if all results within a search share a characteristic, we assume the user filtered for it.

We partition users into clusters using the k-medoids clustering method over the percentage of searches a user makes that are filtered by each room characteristic. The k-medoids method is similar to k-means but uses actual data points (medoids) as cluster centers rather than averages, making it more robust to outliers and more interpretable for categorical data (Kaufman and Rousseeuw, 1990; Bonhomme, Lamadon, and Manresa, 2022). We use $k = 3$ clusters.

The three clusters have natural interpretations (Table 7 and Table 8). Cluster 1 (23% of users) consists primarily of students, split evenly by gender. Cluster 2 (35%) consists mainly of female workers. Cluster 3 (41%) consists mainly of male workers. These demographic characteristics are not used to construct the clusters; they emerge from the search behavior data alone.

We estimate the request utility and click propensity models separately for each cluster, allowing preferences to vary across groups. With cluster-specific preferences, each user's counterfactual ranking is determined by her cluster's estimated parameters. This introduces a form of horizontal differentiation that the pooled model cannot capture: users in different clusters rank the same rooms differently, so full personalization no longer pushes all users toward the same set of rooms.

Figure 10 plots utility and congestion against α for both the baseline ranking (full cluster-specific utility) and the maximal-contrast ranking (horizontal component only), re-estimated on the cluster-specific preference parameters. The two exercises are drawn on common axes so the columns can be compared directly. As in the pooled analysis, the baseline raises utility steeply as α increases while the maximal-contrast ranking rises more slowly and flattens after $\alpha \approx 0.6$. The key difference from the pooled results is in the congestion panels: under cluster-specific estimates, the baseline's Gini coefficient of clicks barely moves as α increases (from 67.93 at $\alpha = 0$ to 75.65 at $\alpha = 1$), while the maximal-contrast ranking's Gini rises even less (67.93 to 72.76). Because the clusters already introduce substantial between-group horizontal differentiation, full per-

Table 7: Medoids of each cluster (k-medoids)

	Cluster 1	Cluster 2	Cluster 3
User has gender preferred by landlord	12.98	20.83	68.54
User has age preferred by landlord	4.21	4.17	15.05
User has occupation preferred by landlord	79.30	4.17	76.33
Has doorman	0.35	0.00	0.27
Has dishwasher	0.35	4.17	0.55
Has terrace	0.00	0.00	0.55
Has AC	0.35	4.17	0.55
Is smoker friendly	2.46	4.17	1.09
Has elevator	0.70	4.17	8.21
Has exterior view	3.51	8.33	3.15
Has TV	4.91	4.17	3.69
Has balcony	0.00	4.17	2.74
Has heating	3.16	4.17	2.05
In Saint-Gervasi or Gracia	0.00	0.00	2.74
In Ciutat Vella or Saint Marti	3.86	0.00	7.66
In Eixample	1.05	0.00	13.27
In North of Barcelona	2.81	0.00	0.14
In Sants-Montjuic or Les Corts	1.05	0.00	6.43

Notes: Clusters are constructed over the percent of searches a user makes filtering by each covariate. A search is said to be filtered if all search results have the covariate = 1.

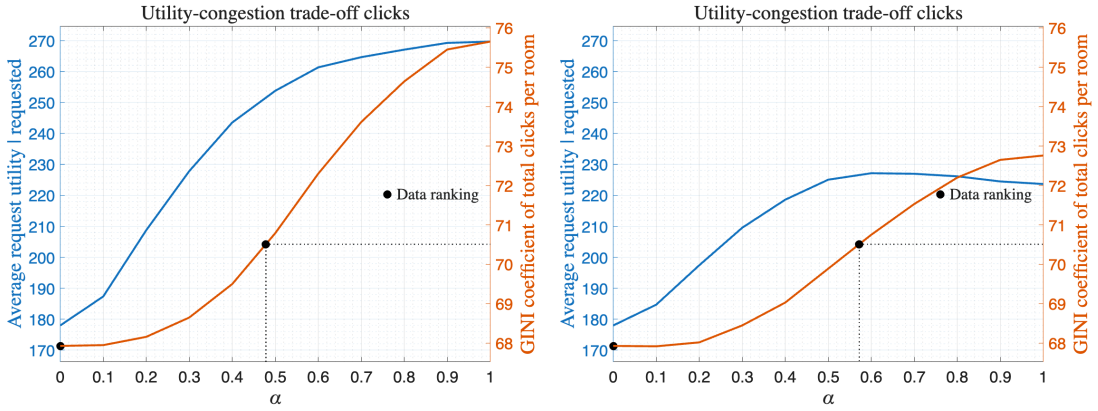
Table 8: Average user characteristics within each cluster

	Cluster 1	Cluster 2	Cluster 3
Percent of users (%)	23.33	35.40	41.27
Age	27.53	29.19	30.64
Is female	0.50	0.93	0.17
Is student	0.59	0.31	0.24
Is worker	0.31	0.93	0.99
Number of searches per user	137.17	174.79	143.61
Number of clicks per user	141.46	169.94	137.77
Number of requests per user	12.58	13.83	13.77

Notes: Clusters are constructed over the percent of searches a user makes filtering by each covariate. User characteristics are not used to construct the clusters.

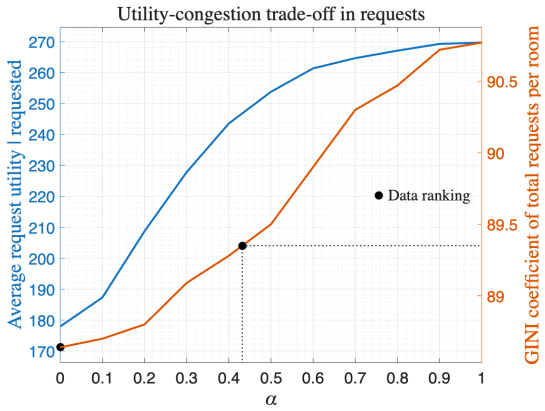
sonalization within the baseline spreads demand more evenly than it did under pooled estimates.

Figure 11 overlays the two frontiers in utility-congestion space. Both frontiers start from the same random-order point (€177.95, Gini 67.93 for clicks) and fan apart as α increases. The maximal-contrast frontier is again the flatter of the two: at full personalization ($\alpha = 1$) the baseline reaches €269.69 at a click Gini of 75.65, while the maximal-contrast ranking delivers only €223.67 at a Gini of 72.76. The two frontiers cross at $\alpha \approx 0.22$ for clicks, much earlier than the $\alpha \approx 0.42$ crossing under pooled estimates. Near the random end the maximal-contrast ranking is more efficient: at $\alpha = 0.2$ it raises utility to €197.40 while barely moving the click Gini (67.93 $\rightarrow$ 68.02), versus €194.45 the baseline would deliver at the same congestion, a gain of €2.95. The gain

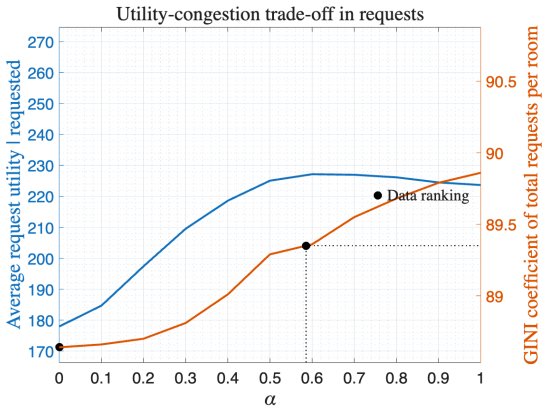


(a) Baseline: clicks

(b) Maximal contrast: clicks



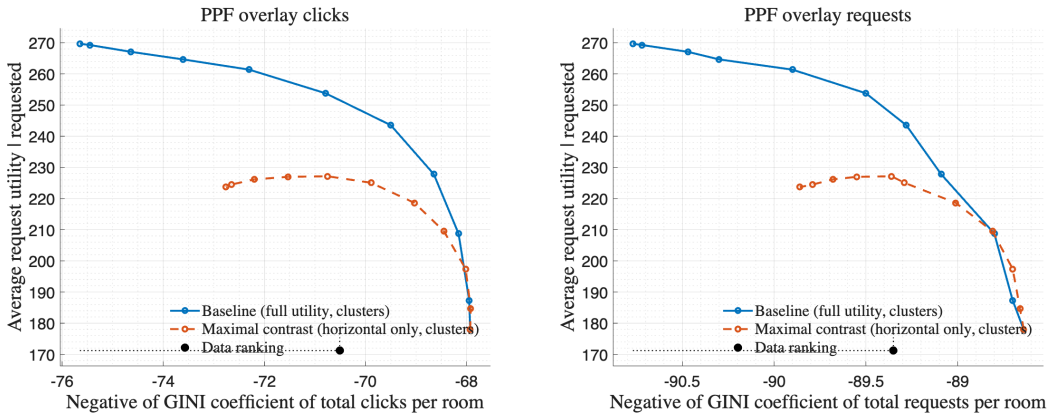
(c) Baseline: requests



(d) Maximal contrast: requests

Figure 10: Utility and congestion as α varies from 0 to 1, cluster-specific estimates. Left column: baseline ranking (full estimated utility). Right column: maximal-contrast ranking (horizontal component only). Top row: clicks; bottom row: requests. Left axis: average utility in euros; right axis: Gini coefficient. Common axes within each row.

is smaller than in the pooled case (€14.51) because cluster-specific preferences already capture much of the horizontal differentiation that the maximal-contrast key is trying to exploit. This is in line with the mechanism described in Section 3.2: as the baseline preference model incorporates more horizontal heterogeneity, the incremental advantage of sorting on the horizontal component alone shrinks.



(a) Clicks (b) Requests

Figure 11: Utility-congestion frontiers for the baseline and maximal-contrast rankings, cluster-specific estimates, overlaid. Each point corresponds to one value of $\alpha \in [0, 1]$; the two curves share a common origin at $\alpha = 0$. The black dot marks the platform's actual outcome.