

Decoding Rankings:

Consumer Search, Price Competition, and Platform Regulation

Hanlin Zhao*

September 24, 2026

[Please click here for the latest version.](#)

Abstract

Online platforms often use monetized rankings: product positions reflect both predicted consumer value and commission revenue. Monetization can steer costly search toward high-commission offers, motivating two policies: neutrality that removes commission from the ranking criterion, and transparency that discloses the ranking's monetization. These policies may substitute if transparency suffices to let consumers discount steering, but complement if improved recommendations under neutrality remains misread with uncorrected beliefs. The determinant is therefore the gap between actual and perceived monetization. Using Expedia data, we estimate an equilibrium model of consumer search and hotel price competition that separately identifies the two. Consumers substantially underestimate monetization. Transparency corrects reliance on the monetized ranking but attractive options remains costly to discover, reducing consumer welfare. Neutrality instead facilitates search but uncorrected beliefs imply excessive skepticism, limiting the welfare gain. The policies are therefore complements. Together they align what recommendations deliver with what consumers believe and raise on-platform bookings by 26% even after hotel repricing. The same mechanism also reverses the platform's revenue comparison: neutrality earns less commission with uncorrected beliefs but more under transparency.

*Toulouse School of Economics. Email: hanlin.zhao@tse-fr.eu.

1 Introduction

Consumers rely on intermediaries whose recommendations may favor the transactions from which they earn more. Online platforms are a clear case where advices often take the form of a ranked list (e.g., Amazon, Google, Uber, Expedia, etc.). An option can be placed near the top because it suits the consumer or because its price generates more commission for the platform. Regulatory responses distinguish disclosing this commercial influence from restricting the ranking itself. The EU’s Digital Services Act and the UK’s hotel-booking commitments require ranking transparency, while the European Commission’s May 2026 proposal calls for ranking neutrality to exclude commercial influences ([European Union, 2019](#); [Competition and Markets Authority, 2019](#); [European Union, 2022](#); [European Commission, 2026](#)). This distinction raises an economic question: are ranking transparency and neutrality substitutes or complements? They may be substitutes if informed consumers can discount commercial steering themselves, reducing the need to change the ranking. But they may instead be complements if neutrality improves recommendations while uncorrected beliefs cause consumers to misinterpret them.

The answer depends on what consumers infer from rank. Consider a relatively expensive hotel at the top of Expedia’s recommended list. Its position reduces the cost of reaching it and can signal a good match before the traveler inspects its page. A consumer who attributes little commercial influence to the ranking may perceive the high position as evidence that the hotel’s quality justifies its price. A consumer who expects more commission-driven promotion discounts that evidence. We call this inference rank decoding. Because inspection is costly, consumers examine only a few hotels, and these beliefs determine which hotels enter consideration. A better alternative the consumer never inspects cannot be booked. These beliefs therefore distort demand through the discovery process even though consumers learn a hotel’s match before booking. Because hotels price against this demand, the distortion carries over to price competition in equilibrium.

This paper asks to what extent consumers’ beliefs about commercial influence differ from the platform’s ranking rule they face, and how this belief gap shapes the interaction between transparency and neutrality policies through consumer search and price competition. To answer, we develop and estimate a structural model using Expedia search and booking data. We represent the platform’s organic ranking by a learning-to-rank algorithm trained on a score that weighs predicted consumer surplus against commission revenue.¹ This score reflects the two roles of rank. Rank is informative, because the platform predicts part of a hotel’s match that consumers learn only by inspection, and

¹An organic ranking is the ordering of listings not separately placed as advertisements.

transmits that prediction through the hotel's position that guides consumers' inspections. Rank is also commercial, to a degree we call the actual monetization weight: the weight on commission in the score. At zero the platform ranks by predicted surplus alone, and at one by commission alone. Consumers hold a belief about the weight and use this belief to interpret the ranking's information about uninspected matches. We call the mean of that belief the perceived monetization weight. Our estimates suggest an actual weight of 0.78 and a perceived weight of 0.25. Consumers thus substantially under-perceive the commercial influence in the platform's ranking.

The estimated belief gap resolves the policy ambiguity: transparency and ranking neutrality are complements because disclosure moves consumers' reliance in opposite directions under the monetized and neutral rankings. A transparency policy informs consumers about the actual monetization weight without changing the ranking rule. Under the estimated monetized rule, this reduces reliance on top recommendations. Yet attractive alternatives remain costly to inspect, and consumer surplus net of inspection costs falls. A neutral-ranking policy instead removes commission revenue from the ranking criterion while retaining consumers' misperceived beliefs. It promotes offers with higher predicted consumer surplus and expands bookings, but leaves consumers skeptical by attributing commercial influence to a ranking from which it has been removed. Adding transparency to the neutral ranking now removes this skepticism, and encourages use of the improved recommendations: bookings rise further and inspection cost falls relative to neutrality alone. Informed consumer is therefore not a costless substitute for a less commercial ranking. Disclosure changes the interpretation of positions, whereas neutrality changes which offers occupy them. These demand responses also change the platform's return to the two rules: the neutral ranking generates less commission with uncorrected beliefs but more when both rules are transparent.

The empirical analysis uses Expedia Group's research dataset for June–July 2021, covering more than 800,000 users and about 2.5 million searches worldwide. The data link displayed hotel lists to the hotels inspected, subsequent search refinements, and bookings. A session is a displayed list and a trip links a user's sessions for the same destination, dates, traveler, and room requirements. Current estimation is based on the data's most searched destination, containing 8,792 trips and 13,016 sessions by 5,311 users. Two features are especially useful. Repeated hotel appearances within a trip provide variation in position for a given consumer and hotel. Consumers also choose whether to follow the recommended order, filter the list, or sort it by an explicit criterion such as price. 22% of sessions use a filter or sort. These choices reveal how consumers use and replace recommendations, connecting within-list inspection to the decision to seek a different ranking list.

The demand model covers search within a ranking list and the choice to generate another. Consumers know prices and visible hotel attributes on the list but must costly inspect a hotel’s page to learn its match for their trip. When entering each list, consumers beliefs determine expectations about uninspected matches and the value of further search. Consumers inspect sequentially following a modified [Weitzman \(1979\)](#), trading expected gains against a cost that increases with the displayed position, and stop when another inspection is no longer worthwhile. They then book an inspected hotel, leave, or open another ranking list by reissuing the query, filtering, or sorting. Inspected offers remain available for recall. Consumers update their subjective beliefs by Bayes’ rule using inspected matches when continuing to a new ranking list. Hotels choose prices before search under Nash–Bertrand competition. They know their demand schedule including how price directly affects utility and indirectly visibility. Re-solving prices under each counterfactual policy captures the response to changes in demand and substitution induced by search.

The central identification problem is that price is chosen by hotels and monetized rank by the platform, both in response to demand we do not fully observe. On the consumer side, we distinguish price preferences, inspection costs, and beliefs. We identify price sensitivity from within-hotel variation in price quintiles across lists. Conditional on hotel effects, trip characteristics, observed histories, and list controls, we assume this variation is unrelated to omitted demand and ranking inputs. Deep inspection can reflect either low inspection costs or a belief that high positions convey little information. The key distinction is that inspection costs depend on the position reached, while rank decoding further depends on the hotel occupying it. For an expensive hotel, a consumer who expects more monetization attributes more of its high rank to commission and less to match quality. The interaction between price and organic rank therefore informs perceived monetization separately from the cost of reaching that position. Consumer-sorted lists provide a complementary comparison: their order follows the consumer’s chosen rule and carries no platform-generated match signal in the model. Interleaved advertisements move displayed positions without changing organic ranks, providing additional variation in inspection costs under conditional exogeneity. We model the choice to sort or filter to account for selection across list types, and use first-session behavior to anchor the trip-entry belief before within-trip inspection evidence accumulates.

On the platform side, actual monetization is identified from observed orderings conditional on demand and the calibrated commission rate. Price enters the economic ranking score twice: it reduces predicted consumer surplus but increases commission. Pairs of hotels ordered differently by predicted surplus and by commission reveal the platform’s trade-off. A higher price can push a hotel up the list to the extent of the monetization

weight, but also lowers predicted consumer surplus and thus pushes it down. The actual monetization weight is a mixture of these two extremes and is therefore identified by how observed rankings resolve these conflicts.

We use a two-step likelihood estimator: demand is estimated at each candidate weight, and the weight is selected by its fit to observed rankings. The estimator integrates over unobserved inspection order and preference heterogeneity. Hotel pricing optimality conditions then recover marginal costs, using nightly rates constructed from the recorded quintiles and auxiliary Expedia price data from 2013.

We evaluate neutrality and transparency policies, separately and together, against the status quo of the platform's monetized ranking with consumers' estimated beliefs. Each comparison is made first at fixed hotel prices and then after hotels reprice. At fixed prices, neutrality increases on-platform bookings by 27.0% and lowers the average booked nightly rate by 25.4%. This rate decline is entirely a shift in which offers consumers book: the two cheapest price groups' share of bookings rises by 39 percentage points while the most expensive group's share falls by 23 percentage points. Suppressing rank decoding while holding other parameters fixed reduces the booking response to 12.9%. Promoting an offer matters not only because it becomes easier to reach but because its position becomes a more encouraging signal. Half of neutrality's effect thus runs through what consumers infer from a hotel's position rather than through the inspection cost. A model that misses the decoding channel would produce misleading policy predictions. After hotel repricing, the booking gain is 24.1%, and the gain in consumer surplus net of inspection costs declines from 1.1% to 0.7%. We find that cheap hotels hold prices to attract marginal consumers while expensive ones lose demand but raise prices to extract more surplus from the remaining loyal consumers. Seller adjustment thus attenuates the gains from improved discovery, but only modestly. Because a neutral ranking penalizes price as consumers do, a hotel that raises its price loses visibility as well as demand. This limits sellers' market power despite market expansion and substitution, and intensifies competition.

The cost of correcting beliefs depends on which ranking consumers face. At the monetized ranking, transparency produces fewer inspections but greater total inspection costs as the remaining inspections occur in costlier positions. Adding transparency to neutrality instead increases bookings by a further 1.8% and lowers inspection costs by 2.6% relative to neutrality alone at fixed prices. With both interventions and equilibrium repricing, on-platform bookings rise by 26.3% relative to the status quo, and surplus net of inspection costs increases by 0.8%, compared with a 0.9% decline under transparency alone. The complementarity arises because disclosure aligns consumers' reliance with what the ranking delivers. Neutral ranking improves recommendations but uncorrected beliefs imply

excessive skepticism, which limits the gain in bookings and consumer surplus. Adding transparency removes this skepticism and raises reliance on the improved recommendations. The return to correcting beliefs therefore rises with the quality of the ranking, and the return to improving the ranking rises with consumers' willingness to follow it.

The same discovery mechanism changes the platform's revenue comparison. With equilibrium prices and uncorrected beliefs, the neutral rule earns 3.1% less commission than the monetized rule. When transparency applies in both regimes, it earns 4.1% more. Disclosure removes demand sustained by underperceived commercial influence under monetization and restores demand discouraged by excessive skepticism under neutrality. Removing the monetized rule's belief-related revenue advantage accounts for most of this reversal. Consumer understanding thus changes the relative return to facilitating transactions with informative rankings rather than steering consumers toward higher-commission offers, even though the commission contract is unchanged.

Related literature. Our contribution is to quantify beliefs about the commercial content of a ranking separately from that content itself, and to trace the resulting belief gap through consumer search and ranking regulation. This builds on research showing that prominence affects search and that positions can convey information ([Armstrong et al., 2009](#); [Athey and Ellison, 2011](#); [Ursu, 2018](#)). [Fong et al. \(2026\)](#) experimentally distinguish position-specific search costs from beliefs about returns to search, showing that both contribute to position effects. We give those beliefs an economic object: the weight consumers attribute to commission revenue. The comparison between this perceived weight and the estimated actual weight identifies the direction of misperception and permits separate policy interventions on the ranking and its interpretation.

[Compiani et al. \(2024\)](#) show that rankings can improve discovery by promoting products whose value exceeds consumers' presearch expectations. Their counterfactuals hold fixed the process generating rank-based expectations. [Greminger \(2026\)](#) estimates beliefs about ranking quality, allows them to adjust under alternative rankings, and finds limited trade-offs among revenue, transactions, and consumer welfare. We distinguish changing consumers' knowledge of commercial influence from changing the influence itself in an equilibrium framework. This distinction explains why disclosure can reduce reliance on a monetized ranking yet increase reliance on a neutral one. It also connects the demand generated by each ranking to the platform's revenue comparison across belief regimes and provides insights to the design of recommendations ([Nocke and Rey, 2024](#); [Janssen et al., 2026b,a](#)).

The closest equilibrium analysis is [Kaye \(2026\)](#), who models ranking-based expectations and hotel repricing and shows that seller price responses can overturn consumer-surplus

gains from personalization. Our analysis shares the link between recommendations, consumer search, and hotel pricing, but studies a different distortion: consumers may misperceive the platform’s commercial ranking objective. Separating actual from perceived monetization allows us to evaluate a restriction on that objective and a disclosure of it within the same framework. Hotel repricing attenuates the effects in our application, while the interaction between the ranking rule and consumers’ beliefs determines the contrast between the two interventions. This complements work linking search design to price competition (Choi et al., 2018; Dinerstein et al., 2018; Salz, 2022; Calvano et al., 2026).

Finally, the paper relates to biased intermediation and its regulation. Hagiu and Jullien (2011) study incentives to divert search, and Heidhues et al. (2023) show how consumer fallibility changes the consequences of platform steering. Boleslavsky et al. (2026) show that transparency can help or harm consumers by changing the platform’s chosen algorithm and its ability to steer and persuade. Our transparency experiment instead holds the ranking rule fixed and changes consumers’ understanding of its commercial content. The resulting demand responses reveal how disclosure interacts with a separate ranking restriction and changes the revenue return to the two rules.

Section 2 presents the data, Section 3 the model, and Section 4 identification and estimation. Sections 5 and 6 report the estimates and counterfactuals.

2 Institutional Setting, Data, and Motivating Evidence

2.1 Hotel Booking on Expedia

Expedia Group is one of the largest online travel platforms in the world. One of its main services is intermediating hotel bookings. In this market, the platform assembles a ranked hotel list in response to a consumer’s query, and thus controls what the consumer sees and in what order. The platform is paid through a commission on each booking, and thus has a commercial interest in both the booking likelihood and the price of the offer booked.² This commercial interest can differ from the consumer’s interest in finding a suitable hotel at a low price. Therefore, when hotels compete in prices, they consider the trade-off between position on the list and consumers’ utilities. Expedia states publicly that commission and additional compensation for visibility can affect the ordering of properties with similar offers (Expedia, 2026a), while the consumer does not know to what extent that matters and

²Expedia distinguishes merchant from agency contracts. Under the merchant contract, Expedia is the merchant of a listing and earns merchant revenue. Under the agency contract, the supplier is the merchant of listing and Expedia earns a commission (Expedia Group, Inc., 2022). Our model approximates platform revenue under both contracts as a share of transaction value.

make search and booking decisions based on her belief about the platform’s commercial interest.

A search query specifies a destination, check-in and check-out dates, the composition of travelers, and the number of rooms. The platform responds to the query with a list of up to 100 hotel cards in a default “Recommended” order. Each card shows the hotel’s star rating, guest review score and review count, free-cancellation and discount information, and a nightly rate. Clicking a card opens the hotel’s page, where the consumer can examine additional information such as the hotel’s amenities and learn its match quality. We call this clicking action a costly inspection. The consumer can proceed to booking only on the hotel page and thus after inspection. Appendix C.2 illustrates the distinction between the list page and the hotel page in the Expedia interface.

Consumers can also replace the current list. Re-issuing the query generates another default list. Filtering restricts the offers, for example by price, stars, or amenities. Sorting orders them by an explicit criterion, such as price-ascending. We call filtering and sorting refinement. A filtered list keeps the platform’s ordering within the eligible set, whereas a sorted list follows the consumer’s criterion. Filtering therefore changes the set to which the platform’s recommendations apply, while sorting replaces the ordering itself. When consumers acquire new information across different list pages, an inspected offer remains available on the opened hotel page at its revealed attributes throughout the trip. This motivates a perfect-recall assumption that is standard in the literature (e.g., Ursu, 2018; Greminger, 2026).

Hotels paying for advertisements are called sponsored listings, contrary to the organic listings.³ Sponsored listings have the same hotel card structure except with an additional label “Ad”. They are placed in the list by a separate rule that is independent of the platform’s organic ranking we study. We thus treat the sponsored placement as exogenously given. Sponsored listings are interleaved with organic listings. Two position measures therefore describe where a hotel appears on the list: the displayed slot and the organic rank. Slot counts every card including advertisements, while the rank counts only organic listings. An advertisement can therefore move an organic hotel to a deeper slot without changing its organic rank. This double-coordinate feature allows us to separate the cost of reaching a position from the information conveyed by that position, which is helpful for the identification analysis in Section 4.⁴ Advertised hotels remain available for inspection

³Expedia sells a small number of sponsored positions on the list page by auction, is paid per click, and labels those positions as “Ad” (Expedia, 2026b). Three lists in four carry at least one “Ad”, the median list carries four, and about one card in eleven is an advertisement. Sponsored hotels are placed in the list by a separate rule. We focus on the organic (i.e., non-sponsored) listings, which are the majority of the page and the ones that the platform’s ranking algorithm directly orders.

⁴Specifically, consumers’ inspection costs depend on the slot on the page as they scroll down, while their beliefs

and booking, but their paid positions supply no organic-rank signal, and their revealed matches do not enter learning about the ranking rule. Sponsored listings are paid per click at a rate that we do not observe in the data. We thus exclude advertising receipts from the modeled commission revenue.

These interface features separate rank, inspection, and refinement. The ordering can signal match quality while reflecting commission incentives. The list page card reveals much about a hotel before any click and the hotel page reveals the rest, so an inspection is a purchase of additional information whose expected value depends partly on that signal transmitted by the platform ranking. The consumer can abandon the platform’s order and build an alternative list of her own with filtering or sorting at an effort cost. Section 3 models these choices jointly.

2.2 Data

We use Expedia Group’s research dataset that contains worldwide search and booking logs from June 1 to July 31, 2021, covering more than 800,000 users and about 2.5 million searches on desktop and mobile devices.⁵ For each search, the data record the query, filters and sorting criterion, the displayed list and its order, the hotel card attributes including sponsored indicators, and the hotels clicked and booked. A companion property table supplies amenities that consumers learn on the hotel page. These linked records connect consumers’ inspections within a list to their choices of subsequent lists.

A session is one displayed list. A trip groups a user’s sessions for the same destination, dates, traveler composition, and room requirements.⁶ At a session’s boundary, the consumer books, leaves without booking, or opens and continues to another list by re-issuing the default query, filtering, or sorting. Leaving means that no further session or booking is observed for that trip on Expedia, which includes both abandoning the trip and booking elsewhere, which the data do not distinguish. Click and booking labels follow the provider’s search-level attribution rule, detailed in Appendix C.1, separately from the model’s maintained availability of previously inspected offers.

For computational ease, estimation uses the destination with the most searches in the data. The final sample contains 8,792 trips and 13,016 sessions from 5,311 users. This is a large city market visited mainly by couples: nearly three quarters of trips are for two adults without children, with a median stay of three nights. The median booking window is one month ahead of check-in. Lists are long with the median showing the full 100 hotels, and

about match quality and the platform’s commercial interest depend on organic ranks in default and filtered lists.

⁵A user identifier is a browser cookie. Destination and hotel identifiers are anonymized.

⁶A gap between two sessions exceeding 24 hours starts a new trip.

busy with sponsored listings where 76.5% of sessions include at least one advertisement. The displayed median hotel has approximately four stars and a guest review score near four, commonly provides free cancellation and discounted prices. Inspection is selective. 28% of sessions and 39% of trips include at least one inspection, and around 4% of trips end in a booking. Continuation is common: 31% of trips go beyond their first list, and 22% of all sessions are lists the consumer refined rather than the platform’s default. Table 6 reports summary statistics of the estimation sample and Appendix C.1 provides more details on the sample construction.

Three data limitations affect the empirical work. First, the recorded price is a quintile within the displayed list, rather than the nightly rate in dollars.⁷ The same hotel can move between quintiles at an unchanged nightly rate when its displayed competitors change. We construct a dollar-price proxy using stratum–quintile centers from Expedia’s 2013 Kaggle dataset, with hotel-specific adjustments from our data’s price-filtered and sorted sessions.⁸ Appendix C.4 provides details on the construction. Second, the data record the set of hotels inspected in a session but not their inspection order, so the likelihood integrates over compatible orders. Third, hotel ownership is unobserved, so the pricing model treats each property as maximizing its own profit.

2.3 Motivating Evidence

The role of position in inspection. Figure 1 shows the inspection rate at each position on lists with at least 30 hotels. Approximately 5% of hotels in the first slot are inspected, compared with about 2% in the second and third, 1.2–1.7% through the nineteenth, and 0.5% by the thirtieth. This descriptive gradient is consistent with costly access to less prominent offers and thus motivates a position-specific inspection cost, as in the consumer search literature (Ursu, 2018; Greminger, 2026). Besides, conditional on inspection, consumers often reach relatively deep in the list: the median deepest inspected slot is 16, and one quarter of sessions reach slot 30 or deeper. This suggests that consumers do not strictly follow the recommendation order, which motivates a sequential search model à la Weitzman.

Refinement and the current list. Consumers do not passively accept the platform’s recommendation order. They often deviate with refinement actions. Filtering accounts for 8.1% of sessions and sorting for 14.3%, together approximately 22% refined sessions. Among

⁷For lists with at least five hotels, available hotels are distributed evenly across the five price buckets. Lower buckets correspond to lower prices. The data code an unavailable hotel’s price as missing.

⁸The auxiliary 2013 Kaggle dataset is widely used in the literature such as Ursu (2018); Compiani et al. (2024); Abaluck et al. (2026); Greminger (2026); Kaye (2026).




Figure 1. Inspection by position on the page

Notes: Share of displayed hotels inspected at each position, on the 11,222 sessions whose list has at least thirty entries (336,640 cards, 4,290 inspections). The band is a 95% confidence interval with standard errors clustered by trip. Position counts advertisements.

trips with more than one session, 69% of second lists are filtered or sorted rather than a re-issued default.

We further explore the reason of refinement. Consumers may refine a list because of the perceived commercial content in the ordering. For instance, a list that visibly favors expensive hotels could make consumers more willing to replace it. Table 1 thus examines what feature from the preceding default list triggers refinement. We investigate four features that reflect the list's monetization degree. Two features use prices. The price-rank slope measures how systematically expensive hotels occupy high ranks, and the top-five price share measures the concentration of the two most expensive quintiles near the top. The other two features describe advertising. *Ceteris paribus*, a list with one standard deviation more tilt toward expensive hotels is associated with a 4.5 percentage point higher probability of refining the next list. A higher share of expensive hotels in the top-5 is associated with a rise of 2.9 percentage points. The advertising coefficients are not statistically different from zero in this regression. Thus, consumers read the monetization of the ranking off the prices and thus commissions, not off the advertisements, and act on it by building the next list of their own rule. This pattern motivates endogenous refinement choices and list-dependent refinement values.

Session endings before and after inspection. Figure 2 conditions on sessions without a booking. The figure shows the share of sessions after which the consumer opens no further session in the trip, by the number of hotels inspected in the session. It investigates

Table 1. Refinement and the preceding default list

Feature of the default list just seen	Next list is refined
Expensive hotels placed near the top (price–rank slope)	4.50 (1.05)
Share of the top-5 in the two most expensive quintiles	2.85 (0.97)
Advertisements’ share of the list	−1.31 (1.43)
Advertisements in the top-10	−0.46 (1.37)

Notes: Linear probability models of whether the session that follows a default list is a filtered or sorted list. The dependent variable is whether that next session is a refined list rather than a re-issued default. Coefficients are in percentage points per one standard deviation of the feature, with standard errors clustered by trip in parentheses. The price–rank slope is the within-list slope of the price quintile on the negative logarithm of rank (so that a higher value means expensive hotels sit nearer the top). All four measures are standardized, and the regression controls for session order within the trip, the list’s mean and dispersion of price quintiles, mean star rating, mean review score, share on discount, length, and share of unavailable hotels, and whether the consumer inspected a hotel on the list and how the inspected hotels compared with the list. The estimation sample is 3,180 default-list sessions followed by another session in the same trip. The base rate that the next list is refined rather than a re-issued default is 68.1%.

whether consumers more often stop searching after inspecting a hotel. Among sessions with no inspection, 60% end the trip. Among those with one inspection, 83% end the trip, and the termination share remains that high as inspections further accumulate. Sessions with any inspection are thus 23 percentage points more likely to end the trip than sessions with no inspection, which is consistent with consumers inspecting, finding an unattractive match, and leaving.⁹ This pattern motivates modeling both stopping without inspection and choosing the outside option after inspection.

3 Model

The model connects the platform’s recommendations, consumers’ interpretation of them, and hotel price competition. The platform’s monetization weight, α^P , describes the trade-off between predicted consumer surplus and commission revenue in the organic ranking. Consumers may attribute a different commercial influence to that ranking. Their trip-entry belief about monetization has mean α^C , the perceived monetization weight. Their beliefs affect which hotels they inspect and hence which offers become available for booking. Hotels set prices before search, knowing the commission rate, marginal cost and the predicted demand at different prices, and recognizing both the utility and visibility effects of their own price. The equilibrium is conditional on the platform’s rule and we do not model the platform’s optimal choice of monetization weight.

⁹The higher termination rate for sessions with an inspection also appears at each of the first four session orders, although its magnitude varies (Appendix C.3).




Figure 2. Leaving the platform after a session, by hotels inspected

Notes: Share of sessions after which the consumer opens no further session in the trip, among the 12,637 sessions that end without a booking, by the number of hotels inspected in the session. Bars are 95% confidence intervals with standard errors clustered by trip. The difference between sessions with any inspection and sessions with none is 23.7 percentage points ($t = 28.6$).

Figure 3 shows the timing of moves. The platform sets the ranking rule and hotels set prices before the trip. The consumer search problem is a dynamic decision process with sequential inspection and multiple sessions. A trip contains at most $\bar{S}$ sessions. At the entry of each session, the consumer reads the list, decodes its order to predict uninspected matches, and chooses whether to costly inspect another hotel or stop. Boundary shocks arrive after stopping, which opens the choice among booking the best inspected offer, leaving, and continuing to another session. If the consumer continues, she decides to generate the new list by re-issuing the default query, filtering, or sorting. The consumer uses her session-entry belief and continuation values through this choice. After choosing continuation, she updates the ranking belief using inspected matches and replans at the next entry. The consumer has perfect recall of inspected offers across sessions. Sections 3.1–3.3 specify utility, ranking, and beliefs. Section 3.4 characterizes the search process. Section 3.5 defines hotel pricing.

3.1 Consumer Utility

Consumer i searches for a hotel for trip t . Session s displays the hotel set $\mathcal{J}_{its}$. Hotel j occupies displayed slot ℓ_{jts} and, when it is organic on a default or filtered list, organic rank r_{jts} . As in Section 2.1, slots count advertisements, whereas organic ranks describe the platform’s non-sponsored ordering. The hotel card displays attributes x_{jts} , including star




Figure 3. Timing of a trip

Notes: Shaded boxes show pre-trip primitives and list formation. The dashed box records consumer's session-entry information. Plain boxes are consumers' actions and ellipses are their terminal payoffs. The dashed containers separate inspection from the boundary choice. Consumer-trip indices i, t are suppressed. W_s is visible list information, F_{js}^C the entry match distribution, ζ_{js} the reservation value, and Ψ_s the expected boundary payoff. Default and filtered lists use the platform's organic order while sorted lists use the consumer's chosen order. Sponsored listings are inserted exogenously after organic ordering and before display, determining slots ℓ_{js} while preserving organic ranks r_{js} . The session-entry belief b_s and continuation values of the next session $V_s^a(\cdot; \omega_s)$ remain fixed through inspection and the boundary choice. Stopping is available without inspection. Only after stopping are mean-zero shocks ε_{as}^B observed for all feasible boundary actions. After a continuation choice, the likelihood L_s uses organic inspections on default or filtered lists to update b_{s+1} . Sorted, sponsored, and repeated inspections supply no evidence for such updates. Inspected offers remain available for recall. Booking requires an inspected offer. At $s = \bar{S}$ there is no continuation, only booking and leaving are available.

rating, review score and count, labels of free cancellation, discount and advertisement, and the nightly rate p_{jts} . Utility from booking the hotel is

$$U_{ijt} = v_{ijts} + m_{ijt}, \quad v_{ijts} = x'_{jts}\beta + \beta_{p,i} p_{jts} + e_{ijt}. \quad (3.1)$$

The consumer knows v_{ijts} before inspection: it combines the card information, her preferences $(\beta, \beta_{p,i})$, and a pre-search taste e_{ijt} that the researcher does not observe. The consumer learns match m_{ijt} only after inspecting the hotel page. Booking requires inspection, so beliefs about an uninspected match affect whether the hotel is discovered. This (m, v) information division follows the empirical search literature (Kim et al., 2010; Ursu, 2018).

The match combines tastes for amenity and an innovation:

$$m_{ijt} = a_{ijt} + \nu_{ijt}, \quad a_{ijt} = x_j^A \beta_{it}^A, \quad \nu_{ijt} \sim N(0, 1). \quad (3.2)$$

Amenities x_j^A are valued through trip-dependent tastes β_{it}^A . For instance, a family values a pool, a long stay values laundry. The innovation is drawn once for each consumer–hotel–trip combination. Both the match and the pre-search taste shock persist across sessions.

3.2 Platform Ranking

Expedia-affiliated research describes rankings that balance relevance and compensation. Nguyen et al. (2017) study re-ranking a relevance ordering to improve marketplace compensation. Abdollahpouri et al. (2020); Cunha et al. (2021) describe relevance, compensation adjustments, and other business components in Expedia-brand recommendation scores. These designs motivate our economic representation of the relevance–compensation conflict, rather than recovering the platform’s black-box algorithm. In our model, relevance is represented by consumer surplus and compensation by commission revenue.

The platform ranks the organic listings by the ideal economic score

$$s_{ijts} = (1 - \alpha^P) (v_{ijts} + m_{ijt}) + \alpha^P \gamma p_{jts}, \quad (3.3)$$

where γ is the commission rate and $\alpha^P \in [0, 1]$ is the platform’s monetization weight.¹⁰ At $\alpha^P = 0$, the ideal score orders hotels by consumer surplus, while at $\alpha^P = 1$ by commission revenue. The monetization weight locates the ranking between these endpoints.

The ideal score is not feasible because the platform ranks at the query-time before the consumer’s match is revealed. Let $\mathcal{I}_{ijts}^P$ be the platform’s query-time information, including

¹⁰In estimation, we replace γ by $\bar{\gamma} = \gamma|\beta^p|$ with the population mean price coefficient $\beta^p < 0$, so a dollar of commission is converted into utility units using $|\beta_p|$.

visible attributes and historical data. The platform's expected score is

$$\hat{s}_{ijts} \equiv \mathbb{E}[s_{ijts} \mid \mathcal{I}_{its}^P] = (1 - \alpha^P) \left(\underbrace{\mathbb{E}[v_{ijts} \mid \mathcal{I}_{its}^P]}_{\bar{v}_{ijts}} + \underbrace{\mathbb{E}[m_{ijts} \mid \mathcal{I}_{its}^P]}_{m_{ijts}^P} \right) + \alpha^P \gamma p_{jts}. \quad (3.4)$$

We assume the platform knows the distribution of price preferences and pre-search tastes when forming the expected visible utility $\bar{v}_{ijts}$. The platform's prediction of the match uses information that the consumer has not yet acquired by inspection. In particular,

$$m_{ijts} = m_{ijts}^P + \eta_{ijts}^P, \quad m_{ijts}^P = \mathbb{E}[a_{ijts} \mid \mathcal{I}_{ijts}^P] + \mathbb{E}[\nu_{ijts} \mid \mathcal{I}_{ijts}^P], \quad (3.5)$$

where the platform's prediction error η_{ijts}^P satisfies $\mathbb{E}[\eta_{ijts}^P \mid \mathcal{I}_{ijts}^P] = 0$ by construction.¹¹ The predictable component m_{ijts}^P is the information about match that the ranking can transmit. The observed ordering therefore serves as a noisy signal of the match, which motivates the consumer's decoding behavior in Section 3.3. The perceived commercial influence affects how much of this information a position conveys to the consumer.

Let $\mathbf{r}_{its} = (j_1, \dots, j_{K_{its}})$ denote the organic order, j_k the hotel at rank k , and $\mathcal{R}_k$ for the set of remaining hotels when rank k is filled. With independent unit-scale type-I extreme-value ranking disturbances ε_{ijts} conditional on the score inputs, the order generated by $\hat{s}_{ijts} + \varepsilon_{ijts}$ has probability

$$\Pr(\mathbf{r}_{its} \mid \hat{\mathbf{s}}) = \prod_{k=1}^{K_{its}} \frac{\exp(\hat{s}_{ij_kts})}{\sum_{l \in \mathcal{R}_k} \exp(\hat{s}_{illts})}. \quad (3.6)$$

This is the Plackett–Luce model widely used in the computer science literature (Luce, 1959; Plackett, 1975; Cao et al., 2007; Burges, 2010). The Plackett–Luce probability fills the first rank by a logit choice among all organic hotels, the second among those remaining, and so forth. Ranking disturbances ε_{ijts} captures the randomness in the ranking process and the platform's information that the researcher does not observe.¹²

The platform-order model applies to organic listings on default and filtered lists. Sorting replaces that order, while sponsored listings are interleaved through a separate placement mechanism that we take as exogenously given. Sponsored hotels remain inspectable and bookable, with inspection effort determined by their displayed slots. Their paid positions

¹¹If the platform knows a_{ijts} and $\mathbb{E}[\nu_{ijts} \mid \mathcal{I}_{ijts}^P] = 0$, then $m_{ijts}^P = a_{ijts}$ and $\eta_{ijts}^P = \nu_{ijts}$. Additional hotel or history predictors the platforms use in $\mathcal{I}_{ijts}^P$ would break this.

¹²For instance, the platform may generate randomness to blunt a hotel's attempt to gain rank by undercutting a competitor by a dollar.

do not convey the organic-rank signal.

3.3 Rank Decoding and Learning

Consumers use the observed ranking to predict uninspected hotels' match, which we call rank decoding. They also use the revealed match values and their associated rank positions to reassess the commercial influence behind the observed ranking, which we call monetization learning. The consumer's trip-entry belief b_{it1} of the platform's monetization weight is a probability distribution over $[0, 1]$ with mean α^C , the perceived monetization weight. The belief carried into session s is b_{its} . Modeling this belief separately from α^P distinguishes the commercial influence in the ranking from the influence consumers attribute to it.

3.3.1 Consumer's subjective model of the ranking

The consumer knows the score's functional form, her visible utilities v_{ijts} , the commission revenue γp_{jts} , and the distribution of the ranking disturbances. She does not know the platform's weight α^P , its predictions m_{ijts}^P , or the ranking disturbances' realized values ε_{ijts} . The consumer's subjective representation of an uninspected match is

$$m_{ijts} = m_{ijts}^{P,C} + \eta_{ijts}^C, \quad m_{ijts}^{P,C} \sim N(0, \sigma_{m^{P,C}}^2), \quad \eta_{ijts}^C \sim N(0, 1),$$

where $m_{ijts}^{P,C}$ is the latent predictable component in her model of the ranking.¹³ Its prior $N(0, \sigma_{m^{P,C}}^2)$ is independent across hotels, and the innovation η_{ijts}^C is independent of these components $m_{ijts}^{P,C}$ and the ranking disturbances ε_{ijts}^C in the subjective model.

At candidate weight α , the deterministic subjective score is

$$s_{ikts}^C(\alpha, m_{ikts}^{P,C}) = (1 - \alpha)(v_{ikts} + m_{ikts}^{P,C}) + \alpha \gamma p_{kts}. \quad (3.7)$$

Adding independent unit-scale type-I extreme-value disturbances ε_{ijts}^C give the subjective order likelihood $\Pr^C(\mathbf{r}_{its} \mid m^{P,C}, W_{its}, \alpha)$ through (3.6):

$$\Pr^C(\mathbf{r}_{its} \mid m^{P,C}, W_{its}, \alpha) = \prod_{k=1}^{K_{its}} \frac{\exp(s_{ij_kts}^C)}{\sum_{l \in \mathcal{R}_k} \exp(s_{ilt_s}^C)},$$

where the non-order information W_{its} contains the displayed hotel set and its visible features.¹⁴

¹³The mean is normalized to zero, with a common level absorbed in the utility intercept.

¹⁴All subjective probabilities also condition on the consumer's known visible utilities, including her private tastes,

3.3.2 Rank decoding: what does the order reveal about a hotel?

Assumption 1 (Role of the order). *The consumer uses the organic order on a default or filtered list to infer hotel matches at each candidate monetization weight, but the order itself does not update the ranking weight belief b_{its} . The displayed set is taken as given. Beliefs about the weight are revised only using newly inspected matches.*

This assumption establishes that the consumer reads ranks as recommendations for the unobserved matches, and the only experience that can challenge her belief about the ranking itself is what an inspection reveals: a poor match found at a high rank. The assumption ensures that first sessions reflect the entry belief alone, which Section 4 uses for identification.¹⁵

For a fixed candidate *alpha*, Bayes' rule tilts the unconditional prior of predictable matches $m_{ijt}^{P,C} \sim N(0, \sigma_{m^{P,C}}^2)$ with the subjective likelihood of the observed order $\Pr^C(\mathbf{r}_{its} \mid m^{P,C}, W_{its}, \alpha)$:

$$\pi_{its}(m^{P,C} \mid D_{its}, \alpha) \propto \Pr^C(\mathbf{r}_{its} \mid m^{P,C}, W_{its}, \alpha) \prod_{k=1}^{K_{its}} \frac{1}{\sigma_{m^{P,C}}} \varphi\left(\frac{m_{ijk}^{P,C}}{\sigma_{m^{P,C}}}\right), \quad (3.8)$$

where φ is the standard normal density. The values of predictable matches $m_{ijt}^{P,C}$ receive more weight when they make the observed ranking $\mathbf{r}_{its}$ more likely. Adding the independent innovation η_{ijt}^C , we obtain the resulting α -specific CDF of the actual match $m_{ijt} = m_{ijt}^{P,C} + \eta_{ijt}^C$:

$$F_{ijts}^C(m \mid D_{its}, \alpha) = \int_{-\infty}^{\infty} \Pr(m_{ijt}^{P,C} + \eta_{ijt}^C \leq m \mid D_{its}, \alpha) \pi_{its}(m^{P,C} \mid D_{its}, \alpha) dm_{ijt}^{P,C},$$

where $D_{its} = (\mathbf{r}_{its}, W_{its})$ denotes the displayed list. We then integrate over the consumer's belief b_{its} about the monetization weight to obtain the rank-conditioned CDF of the actual match m_{ijt} :

$$F_{ijts}^C(m) = \int_0^1 F_{ijts}^C(m \mid D_{its}, \alpha) b_{its}(d\alpha), \quad (3.9)$$

which is a mixture of the CDFs at each perceived weight value, with weights given by the consumer's belief about the monetization weight.¹⁶ Appendix A.1 gives the normalizing integral, marginalization, and convolution for a general list.

this conditioning is suppressed in the notation.

¹⁵The unrestricted Bayesian model conditioning on all observed information in which the order itself updates the belief of the ranking is left for future work.

¹⁶By construction, the mean of the perceived match distribution F_{ijts}^C is the mean of $m_{ijt}^{P,C}$ under the weight belief mixture, and its variance is that mixture's variance plus one that comes from the innovation η_{ijt}^C .




Figure 4. What a rank says about a hotel, under two beliefs

Notes: Expected match (mean of the distribution (3.9)) given rank, evaluated at a fixed weight and averaged over the organic listings of the default lists shown in the first session of each trip in the estimation sample (8,787 lists). The two curves use the estimates of perceived weight $\alpha^C = 0.25$ and platform's actual weight $\alpha^P = 0.78$ reported in Section 5. The higher weight produces a lower and flatter profile over the ranks shown. At the perceived weight the top of the list is read as a strong signal of a good match and the signal fades with rank. At the platform's weight the profile is flat and low because a consumer who knew how monetized the ranking is would read little into any rank.

The distribution in (3.9), shifted by visible utility, supplies the returns to inspection and determines the reservation values that govern the optimal inspection strategy in Section 3.4.5. The distribution has two features. First, a high perceived weight raises the expected match. Figure 4 illustrates this by evaluating the decoding on the sample's first-session default lists. At the higher weight, the plotted expected-match profile is lower and flatter because these top positions receive less credit for match quality. Second, for an expensive hotel, part of its high rank is attributed to the commission, so the credit about the match is smaller. This implies that the same rank informs less about an expensive hotel's match than about a cheap one. This price-by-rank interaction helps identify perceived monetization in Section 4. Appendix A derives the continuous match distribution in a two-hotel example and shows how the same ranking can induce different first inspections when only the perceived weight changes.

A sorted position or a sponsored slot supplies no platform-rank signal. An uninspected hotel in either case keeps the predictable-component unconditional prior $N(0, \sigma_{m^{P,C}}^2)$ and hence the full-match forecast $N(0, \sigma_{m^{P,C}}^2 + 1)$. An already inspected hotel instead keeps its known match and utility, with no reinspection nor decoding. At a new list, decoding for uninspected hotels are rebuilt from that list and the updated weight belief.

3.3.3 Monetization learning: what do inspections reveal about commercial influence?

Inspection reveals whether a recommended hotel matches the trip. The consumer can use that experience to revise the weight she attributes to commission. Let $\mathcal{C}_{its}^{\text{new}}$ collect hotels inspected in an organic session s where the list is not sorted and thus is informative about the platform's monetization weight.¹⁷ Denote consumer's rank-conditioned subjective density of the actual match m_{ijt} by $f_{ijts}^C(m \mid D_{its}, \alpha)$. At each candidate weight, the consumer perceives the revealed matches as independent draws from her subjective law $f_{ijts}^C(m \mid D_{its}, \alpha)$. The subjective likelihood $L_{its}(\alpha)$ of the observed matches is therefore the product of these densities over the inspected hotels:

$$L_{its}(\alpha) = \prod_{j \in \mathcal{C}_{its}^{\text{new}}} f_{ijts}^C(m_{ijt} \mid D_{its}, \alpha). \quad (3.10)$$

If the consumer chooses to continue in a new session $s + 1$, her session-entry belief about the monetization weight b_{its} is updated by Bayes' rule and becomes

$$b_{it,s+1}(d\alpha) = \frac{L_{its}(\alpha)b_{its}(d\alpha)}{\int_0^1 L_{its}(a)b_{its}(da)}, \quad (3.11)$$

that she carries into the next session. A hotel that was ranked high and turned out to match poorly, or one that was ranked low and matched well, is more likely under greater commercial influence, therefore moves the belief of the monetization weight up. By contrast, a good match at a high rank confirms her trust and moves the perceived weight down. The direction is jointly determined by candidate-specific likelihoods $L_{its}(\alpha)$, taking account of the hotel's visible utility, price, and the rest of the displayed list. An empty inspection set gives $L_{its}(\alpha) = 1$ and leaves the weight belief unchanged.

It is important to highlight that the product in (3.10) is a subjective conditional-independence likelihood for learning. It is not the joint match density implied by the rank-conditioned law in (3.8) for inspections. (3.10) multiplies densities at each candidate before updating the weight belief, whereas inspection uses the already mixed entry marginals in (3.9). Bayes' formula is applied relative to the learning likelihood $L_{its}(\alpha)$ without tempering: the learning speed is fixed at one.¹⁸

When are beliefs revised? We treat each displayed list as a planning unit. The consumer immediately evaluates the match revealed by an inspection as it changes the set of known offers and may raise the best-found utility $\bar{u}$. However, the consumer defers reassessment of

¹⁷The session needs to display an organic listing on a default or filtered list. Sorted and sponsored inspections are excluded.

¹⁸A learning rate λ would use L_{its}^λ in (3.10).

the platform’s monetization weight after she decides to open another list. In fact, assessing whether a particular hotel suits the trip involves a different inference from assessing how the platform trades match quality against revenue. The long lists and selective inspections described in Section 2.2, together with the filtering and sorting options in Section 2.1, motivate this division of attention. The following assumption formalizes this.

Assumption 2 (Entry-fixed forecasts and post-boundary updating). *Inspection, stopping, and boundary decisions use the session-entry weight belief and predictive marginals of match. At entry, each continuation action is assigned a complete expected-payoff function $V_s^a(\cdot; \omega_s)$ of the best-found utility. This function remains fixed through the boundary choice. Current inspections affect its value only through $\bar{u}$. After choosing to continue, the consumer updates the weight belief by (3.10) and carries it with all inspected offers into the next session.*

Future search thus enters the current stopping decision, but the consumer does not anticipate how current inspections would revise the rule belief and continuation functions until the boundary stage. This separation of learning from anticipation is related to anticipated utility in Cogley and Sargent (2008). Gabaix et al. (2006) also provide experimental support for partially myopic information acquisition. In our model, updating first affects inspections on the next list, not the choice to open that list. The current list and the offers already discovered can nevertheless change that choice at an unchanged belief, as in the refinement margin motivated by Table 1.

Besides, the post-boundary updating is a behavioral restriction that makes the within-session problem tractable, and it is a fair description of how a list is used: while the consumer is inspecting, she focuses on what the pages reveal about the hotels, and what they reveal about the platform is read when she sits down to the next list. The consequence is that the consumer’s subsequent replanning need not coincide with an earlier fully contingent plan.

3.4 The Trip as a Dynamic Problem

At each list, the consumer decides how much to inspect before choosing whether to book, leave, or seek another list. Stopping is available at session-entry and after every inspection. It ends inspection of the current list, not necessarily the trip: the stopping value includes the option of further sessions. The boundary shocks are observed only after stopping, until when the session-entry beliefs, predictive marginals of match and expected continuation values remain fixed and govern the ensuing choice. We suppress consumer–trip indices i, t where unambiguous.

3.4.1 States and inspection

The session-entry state is $z_s = (\bar{u}_s^0, \omega_s)$. The best-found utility, $\bar{u}_s^0$, equals $-\infty$ when no offer has been inspected. ω_s contains the session index, displayed list and order, weight belief b_s , predictive marginals of match, list covariates, feasible continuation actions, and expected continuation values. Let $\mathcal{K}_{s-1}$ be the identities inspected earlier in the trip. The new inspection opportunities are $\text{NI}_s^0 = \mathcal{J}_s \setminus \mathcal{K}_{s-1}$, including uninspected sponsored hotels. The full history also contains each inspected offer's match, price, and attributes, together with the evidence already used for monetization learning.

Conditional on ω_s , an inspection decision uses the current best utility $\bar{u}$ and remaining non-inspected set NI_s . Inspecting $j \in \text{NI}_s$ reveals m_j and hence U_j at cost

$$c(\ell) = c_I (1 + \ell)^\psi, \quad (3.12)$$

where $c_I > 0$ and $\psi \geq 0$. The cost is nondecreasing in the displayed slot and does not depend on hotel identity or organic rank. The inspection adds the offer to recall, removes j from NI_s , and changes the best utility to $\max\{\bar{u}, U_j\}$. The consumer then makes the inspect-or-stop comparison again, using the same entry belief and predictive marginals of match.

3.4.2 Boundary actions

After stopping, the consumer can book an inspected offer, leave, or, when $s < \bar{S}$, generate another list by re-issuing the default query, filtering, or sorting. Booking has deterministic value $\bar{u}$ and is unavailable without an inspected offer. Leaving has value

$$u_{0,it} = z'_{it} \vartheta, \quad (3.13)$$

a function of trip characteristics. Booking urgency can therefore shift the value of leaving. Leaving represents no booking on Expedia for this trip, including both abandoning the trip and booking elsewhere. Every feasible boundary action a receives its own independent type-I extreme-value shock ε_{as}^B , with mean zero and standard deviation $\sigma_B > 0$. These shocks are independent of inspection rewards and previous sessions' boundary shocks, and are observed only once inspection has stopped.¹⁹

¹⁹Denote $\tau_B = \sqrt{6}\sigma_B/\pi$ as the Gumbel scale and γ_E for Euler's constant. Equivalently, $\varepsilon_{as}^B = \tau_B(G_{as} - \gamma_E)$ for independent standard maximum-Gumbel draws G_{as} . Thus the shock location is $-\gamma_E\tau_B$.

3.4.3 The continuation value of another list

For a continuation action $a \in \{\text{def}, \text{filter}, \text{sort}\}$, the deterministic value is

$$V_s^a(\bar{u}; \omega_s) = -c_{\text{sess}} + \kappa_a + \theta_a \mathbf{1}\{a = k_s\} + b'_a x_s^{\text{list}} + \underbrace{\mathbb{E}[V_{s+1}^E(\bar{u}, \omega_{s+1}) \mid a, \omega_s]}_{\widehat{\mathcal{W}}_s^a(\bar{u}; \omega_s)}, \quad (3.14)$$

where entering a new session costs c_{sess} . The action-specific intercept $\kappa_{a \in \{\text{filter}, \text{sort}\}}$ shifts the continuation value relative to the default action. The indicator $\mathbf{1}\{a = k_s\}$ equals one when the consumer reuses the tool k_s that generated the current list s : a positive θ_a raises the value of reusing that tool, while a negative coefficient lowers it. The covariates x_s^{list} describe the current list's price level, price-rank pattern (e.g., how systematically the current list places expensive hotels near the top), and advertising share. Their coefficients allow refinement to respond to the current recommendations at an unchanged weight belief, as motivated by Section 2.3. The final term is the session-entry forecast of future payoff from continuing to another session, including future inspection costs and boundary payoffs. It is defined as

$$\widehat{\mathcal{W}}_s^a(u; \omega_s) = \int \widehat{V}_{s+1|s}^E(u, \omega') \widehat{\Gamma}_a(d\omega' \mid \omega_s),$$

where $\widehat{\Gamma}_a$ is the session-entry forecast transition law for the next list and $\widehat{V}_{s+1|s}^E$ its forecast value. Default and filtered lists use the anticipated platform order, the latter within a restricted set, and sorting uses the consumer's chosen order. The forecast law and value function are held fixed through the boundary. A newly inspected offer can nevertheless change $\widehat{\mathcal{W}}_s^a$ through the utility u carried into further search. Holding a function fixed thus does not hold its value constant. Appendix B.2 gives sufficient monotonicity and integrability conditions. Given the computational burden, Appendix D.3 describes one additional numerical approximation of one-list-ahead forecast.

If the consumer chooses continuation, she then updates belief of the platform's monetization weight b_s by (3.10), using only organic inspections on default or filtered lists. Sorted, sponsored, and repeated inspections supply no new weight evidence for updates. All inspected offers are carried forward. At the entry of a new list, the consumer forms new forecasts for hotels not previously inspected. The update therefore first affects inspections on the next list.

3.4.4 Conditional Bellman equation

Let V_s^E , V_s^I , and V_s^B denote entry, inspection, and boundary values conditional on the session-entry forecasts. They satisfy

$$\begin{aligned} V_s^E(\bar{u}_s^0, \omega_s) &= V_s^I(\bar{u}_s^0, \text{NI}_s^0; \omega_s), \\ V_s^I(\bar{u}, \text{NI}; \omega_s) &= \max \left\{ \mathbb{E}[V_s^B(\bar{u}; \omega_s, \varepsilon_s^B)], \right. \\ &\quad \left. \max_{j \in \text{NI}} \left[-c(\ell_j) + \mathbb{E}[V_s^I(\max\{\bar{u}, U_j\}, \text{NI} \setminus \{j\}; \omega_s)] \right] \right\}, \\ V_s^B(\bar{u}; \omega_s, \varepsilon_s^B) &= \max \left\{ \bar{u} + \varepsilon_{\text{books}}^B, u_0 + \varepsilon_{\text{quits}}^B, \right. \\ &\quad \left. \max_{a \in \{\text{def}, \text{filter}, \text{sort}\}} [V_s^a(\bar{u}; \omega_s) + \varepsilon_{as}^B] \right\}. \end{aligned} \quad (3.15)$$

The stopping expectation integrates boundary shocks not yet observed, and the inspection expectation integrates U_j under its session-entry forecast. The booking term is omitted when no offer is inspected, and the continuation terms are omitted at $s = \bar{S}$. Maxima over empty sets equal $-\infty$. Thus zero inspection can lead to booking an earlier inspected offer, leaving, or continuing. At the final session only booking and leaving remain, and the consumer must leave without any inspected offer. Booking and leaving end the trip.

3.4.5 Optimal inspection within a session

The inspection policy compares the cost of inspecting an offer with the improvement it can bring to the value of arriving at the boundary. We adapt the optimal stopping policy of [Weitzman \(1979\)](#) after transforming each utility by the expected payoff of stopping. Appendix B provides the complete set of assumptions and proofs. Two restrictions are needed to obtain an independent-box representation of the problem.

Assumption 3 (Independent inspection and scalar stopping value).

(a) Inspection independence. At fixed entry state ω_s , the consumer evaluates uninspected hotels as independent draws from their entry utility marginals F_{js}^U , obtained by shifting (3.9) with visible utility via $U_{js} = m_j + v_{js}$. Revealing one match does not change the other hotels' current forecasts.

(b) Scalar stopping-value sufficiency. For every feasible inspection history h at that entry, the expected boundary payoff is $B_s(h; \omega_s) = \Psi_s(\bar{u}(h); \omega_s)$ for a finite, nondecreasing map Ψ_s on $\{-\infty\} \cup \mathbb{R}$.²⁰

²⁰ $\bar{u} = -\infty$ when no hotel has been inspected and $\bar{u} \in \mathbb{R}$ when at least one hotel has been inspected. The regular integrability conditions in Assumption 5 apply.

These assumptions are behavioral restrictions on the consumer’s beliefs and expectations, which allow the inspection problem to be solved by a simple reservation-value comparison and therefore substantially reduce the computational burden.

Assumption 3(a) is a conditional independence assumption that requires the consumer to evaluate each uninspected hotel independently using its session-entry marginal.²¹ It keeps rank-specific inspection returns while permitting the conditional index policy in Section 3.4.5. The limitation of imposing Assumption 3(a) is that it suppresses cross-hotel inferences that could arise from rank conditioning or shared ranking rule uncertainty. Correlated search is a natural extension, requiring a larger information state or a justified solution method that we leave for future work.²²

Assumption 3(b) states that the expected payoff of stopping depends only on the best-found utility $\bar{u}$ and not on the full inspection history. At each new list Assumption 3(b) permits updating and replanning, so it keeps monetization learning across lists and only restricts assigning current inspections a strategic value for learning the commercial rule, which is plausible in our context, given the limited attention and selective inspections described in Section 2.2. Assumption 3(b) is a restriction on the stopping value. We do not have to impose it because Assumption 2 already imposes a stronger action-level restriction that each continuation function $V_s^a(\bar{u}; \omega_s)$ is fixed at session-entry as a function of $\bar{u}$. Appendix B.3 distinguishes scalar sufficiency from fixed action-level functions (Assumption 3(b) versus Assumption 2).²³

With nondecreasing action forecasts and the mean-zero shocks, Assumption 3(b) allows us to express the value of arriving at the boundary with best utility $\bar{u}$ and state ω_s as

$$\Psi_s(\bar{u}; \omega_s) = \tau_B \log \left[\exp \left(\frac{\bar{u}}{\tau_B} \right) + \exp \left(\frac{u_0}{\tau_B} \right) + \sum_{a \in \{\text{def}, \text{filter}, \text{sort}\}} \exp \left(\frac{V_s^a(\bar{u}; \omega_s)}{\tau_B} \right) \right], \quad (3.16)$$

with τ_B denotes the scale parameter of the boundary shock.²⁴ The third continuation terms are dropped at $s = \bar{S}$. The booking term is zero at $\bar{u} = -\infty$. The continuation sum is zero at the final session. Appendix B.3 fully derives this inclusive value.

If the consumer inspects hotel j and then stops, the gross expected gain over stopping

²¹The joint posterior in (3.8) generally creates dependence, and uncertainty about a common weight can create further dependence. Assumption 3(a) behaviorally restricts the theory for tractability.

²²Hodgson and Lewis (2025) estimate cross-product learning in a different consumer-search environment. Gergatsouli and Tzamos (2023) develop conditional reservation algorithms with approximation guarantees for a cost-minimization problem against a partially adaptive benchmark. Those work however does not guarantee any exact fixed-index theorem for the ordinal decoder we study here.

²³With EVT1 shocks and fixed $\bar{u}$, individual history-dependent continuation action values could satisfy the scalar condition if their exponentiated sum were invariant across those histories, from which we derive Assumption 3(b). Assumption 2 instead holds each continuation action forecast fixed.

²⁴There is no additional Euler-constant term under the mean-zero normalization.

now is

$$\mathbb{E}[\Psi_s(\max\{\bar{u}, U_j\}; \omega_s) - \Psi_s(\bar{u}; \omega_s)] = \mathbb{E}[(\Psi_s(U_j; \omega_s) - \Psi_s(\bar{u}; \omega_s))^+],$$

where the equality follows from monotonicity of Ψ_s .²⁵ A better offer raises the value of booking and may also improve the option carried into later search. Accordingly, the relevant reward is $\Psi_s(U_j; \omega_s)$ and the transformed rewards are independent boxes with outside value $\Psi_s(\bar{u}; \omega_s)$ and positive opening costs $c(\ell_j)$. Applying the independent-box argument of [Weitzman \(1979\)](#) to these transformed rewards gives hotel j 's reservation value ζ_{js} :

$$c(\ell_j) = \int [\Psi_s(u; \omega_s) - \zeta_{js}]^+ dF_{js}^U(u), \quad (3.17)$$

where F_{js}^U is the distribution of consumers' booking utility U_{ij} . The consumer's rank decoding enters the reservation value ζ_{js} and therefore affects consumer's consideration set and final choices through this distribution. Under the regularity conditions in [Appendix B.2](#), every finite positive inspection cost $c(\ell_j)$ has a unique finite solution ζ_{js} .

The reservation value is a hotel-specific index measured in boundary-value units. It is the stopping value at which that inspection's expected gain equals its cost. Conditional on the session-entry forecasts, a lower inspection cost, a first-order stochastically larger transformed reward, or a mean-preserving spread of the transformed reward weakly raises the index. Under [Assumption 3](#), consumers' inspections change the stopping value against which the index is compared but not the index itself. The reservation value remains fixed during the session and is therefore a sufficient statistic to characterize the optimal inspection policy as follows.

Proposition 1 (Optimal inspection conditional on entry forecasts). *Under [Assumption 3](#) and the regularity conditions in [Appendix B](#), there exists an optimal policy conditional on the session-entry forecasts:*

- (i) *inspects an uninspected hotel with the largest reservation value ζ_{js} whenever that value exceeds $\Psi_s(\bar{u}; \omega_s)$;*
- (ii) *stops, including before any inspection, when no hotel remains uninspected or the best remaining reservation value is below the transformed reward of the best found utility, i.e., when*

$$\Psi_s(\bar{u}; \omega_s) \geq \max_{j \in \text{NI}_s} \zeta_{js}; \quad (3.18)$$

- (iii) *after observing the boundary shocks, selects the feasible action maximizing [\(3.15\)](#), using the session-entry continuation functions;*
- (iv) *if booking, selects the highest-utility among inspected offers.*

²⁵Because Ψ_s is nondecreasing, $\Psi_s(\max\{\bar{u}, U_{ij}\}) = \max\{\Psi_s(\bar{u}), \Psi_s(U_{ij})\}$.

At $s = \bar{S}$, if the consumer has inspected at least one offer, booking if $\bar{u} + \varepsilon_{\text{book},s}^B \geq u_0 + \varepsilon_{\text{quit},s}^B$ and leaving otherwise. If no offer has been inspected, leaving is optimal.

Appendix B.5 proves the result by capping each transformed reward at its reservation value. Applying the policy anew after each continuation and belief update defines behavior across sessions in the trip. Its optimality is conditional on session-entry forecasts without considering the full anticipation of subsequent belief updating and replanning.

3.5 Hotel Pricing

Hotels choose nightly rates simultaneously before consumers search in a Nash–Bertrand game. Each property knows its own marginal cost, commission rate, and induced demand schedule over relevant prices, and best responds to rivals’ prices.²⁶ We assume single-property price setting because ownership is unobserved.

Hotel j earns a commission-net margin on each booking, with marginal cost c_j per occupied room-night. Its profit is

$$\pi_j(p) = [(1 - \gamma) p_j - c_j] D_j(p), \quad (3.19)$$

where $D_j(p)$ is the expected number of bookings of hotel j at the price vector p , integrating over consumers and trips, the order law (3.6), rank decoding, inspection, continuation, and booking.

Price affects demand through both utility and ranking. It enters each consumer’s visible utility with coefficient $\beta_{p,i}$ in (3.1), changing both the reservation value of inspection and the booking value of an inspected offer. It also changes the order, and hence visibility, inspection costs, and decoded returns. Holding other economic-score components in (3.3) fixed, the net price loading is $\phi = (1 - \alpha^P)\beta_p + \alpha^P\gamma$. The score therefore penalizes price when the platform’s monetization weight is below $\alpha^P = \beta_p/(\beta_p - \gamma)$ and rewards it above that threshold. Consequently, within a certain range of α^P , decreasing monetization weight may not lower every hotel’s markup or posted price and stimulate competition, because the score can remain net increasing in price, so a higher price can improve visibility and inspection returns. The net effect on demand is therefore ambiguous, and the platform’s monetization weight can create a non-monotone relationship between price and demand. This analysis illustrates how changes in recommendations alter consideration and substitution, connecting ranking design to price competition as in Dinerstein et al. (2018).

²⁶We do not require each hotel to have separate knowledge of the actual or perceived monetization weight or the ranking algorithm.

At an interior differentiable optimum, the hotel's first-order condition is

$$(1 - \gamma) D_j(p) + [(1 - \gamma) p_j - c_j] \frac{\partial D_j(p)}{\partial p_j} = 0, \quad (3.20)$$

where the demand derivative includes both utility and rank-mediated effects. A conditional equilibrium is a price vector at which every property maximizes this profit given rivals' prices and expected demand.

4 Estimation

The empirical analysis separates consumers' perceived monetization from the actual one in the platform's ranking. Inspection, list choice, and booking inform consumer parameters conditional on the displayed lists. Observed organic orders inform actual monetization through fitted ranking functions. This section specifies the empirical inputs, discusses the identifying variations, and presents the estimator. Appendix D provides more details of the numerical implementation.

4.1 Empirical Specifications

4.1.1 Preferences and private information

The visible utility in (3.1) is specified as

$$v_{ijts} = x'_{jts} \beta + \beta_{p,i} p_{jts} + \xi_j + \tilde{e}_{ijt}, \quad \beta_{p,i} = \beta_p + \sigma_p \eta_{p,i}, \quad (4.1)$$

where x_{jts} contains a constant, star rating, review score and count, dummies for free cancellation, discount, and advertisement. The price draw $\eta_{p,i} \sim N(0, 1)$ gives an individual coefficient with population mean β_p and standard deviation σ_p . The hotel (group) effect ξ_j is separated from the pre-search taste with $e_{ijt} = \xi_j + \tilde{e}_{ijt}$. The remaining taste shock $\tilde{e}_{ijt}$ has a unit-scale type-I extreme-value distribution (Kim et al., 2010; Ursu, 2018). It is known to the consumer before search and persists across sessions for the consumer–hotel–trip combination. The researcher integrates over this taste and price heterogeneity jointly with the persistent matches.

4.1.2 Constructed nightly rates

The economic price p_{jts} is continuous, but the data record only the within-list quintile q_{jts} . We construct the dollar-price input as

$$p_{jts} = c(g_{ts}, q_{jts}) + \delta_j, \quad (4.2)$$

where the stratum g_{ts} combines stay length, Saturday-night inclusion, and booking window. For each stratum, Expedia's auxiliary 2013 data supply the median across searches of the within-search 10th, 30th, 50th, 70th, and 90th percentile prices, which gives one center value $c(g_{ts}, q_{jts})$ for each quintile q_{jts} . The constant hotel offset δ_j combines price-sort comparisons, price-filter bands, and average quintiles, with greater weight on the direct sort and filter comparisons and zero offset when this information is unavailable. Appendix C.4 provides more details.

4.1.3 Matches revealed by inspection

Hotel-page amenities enter (3.2) through:

$$m_{ijt} = \sum_{g=1}^4 \beta_g^A(z_{it}) n_{gj} + \nu_{ijt}, \quad \nu_{ijt} \sim N(0, 1), \quad (4.3)$$

where n_{gj} counts the number of amenities in group g that hotel j has. We distinguish amenities into four groups: services (bar, gym, spa, laundry), free extras (breakfast, parking, wifi, airport transfer, internet), leisure amenities (pools, hot tub), or home-style features (air conditioning, washer and dryer, pets allowed). Each group has its own taste coefficient $\beta_g^A(z_{it})$ that includes a constant and the following trip interactions: number of children with leisure, the length of stay with home-style features, and number of travelers with services. The persistent innovation ν_{ijt} is drawn once per consumer–hotel–trip combination.

4.1.4 Platform's predicted matches and learning-to-rank algorithm

The actual platform prediction is $m_{ijts}^P = \mathbb{E}[m_{ijt} \mid \mathcal{I}_{ijts}^P]$. The researcher does not observe matches m_{ijt} nor the platform's full information $\mathcal{I}_{ijts}^P$. We first construct a model-based match $\tilde{m}_n(\theta)$ from (4.3) at demand parameters θ for observation n . We then fit an approximation

$$\hat{m}_{ijts}^P$$

to the platform’s prediction m_{ijts}^P using regressors Z_{ijts}^P , including hotel (group) indicators χ_j and the overlap h_{ijts} between the hotel’s amenities and those of hotels inspected earlier in the trip:

$$\begin{aligned}\hat{b}^P(\theta) &= \arg \min_b \sum_n [\tilde{m}_n(\theta) - Z_n^{P'} b^P]^2, \\ \hat{m}_{ijts}^P(\theta) &= Z_{ijts}^{P'} \hat{b}^P(\theta).\end{aligned}\tag{4.4}$$

In practice, Expedia generates the observed ranking with complicated proprietary algorithms that are not known to the researcher. The platform’s ranking per se is a black box, but it is informed by the features and historical data the researcher observes. We therefore approximate the platform’s actual ranking technology with a learning-to-rank function trained on the fitted match and other observed query-time features. This approach allows us to capture the essential elements of the platform’s ranking behavior while maintaining flexibility in the mapping from our model’s economic objectives (3.4) to realized rankings. Following [Kaye \(2026\)](#), we use LambdaMART ([Burges, 2010](#)) to approximate Expedia’s ranking technology. The difference is that we discipline the ranking objective with an interpretable economic score (3.4) and the monetization weight α^P provides a policy-relevant design parameter.

For candidate weight α , we plug the fitted prediction $\hat{m}_{ijts}^P(\theta)$ into the economic score (3.4) to produce the training target for LambdaMART:

$$T_{ijts}(\alpha, \theta) = (1 - \alpha)(\bar{v}_{ijts} + \hat{m}_{ijts}^P(\theta)) + \alpha\gamma p_{jts},\tag{4.5}$$

where $\bar{v}_{ijts}$ averages over private tastes and price preferences. We calibrate the commission rate to $\gamma = 0.15$.²⁷ LambdaMART learns a function $f_{\alpha, \theta}$ from observed query-time features and these targets. Its output supplies the index g^R used in the empirical ranking likelihood (4.10) to estimate the actual monetization weight. Appendix D.1 describes more details of this process.

4.1.5 Beliefs and inspection

The trip-entry belief b_{it1} is a Beta distribution with estimated mean α^C and fixed concentration, represented on a grid over $[0, 1]$. We fix the subjective predictable-component standard deviation at $\sigma_{m^{P,C}} = 1$ and the learning speed in (3.10) at $\lambda = 1$. The reason is that with few inspections per trip in our context, prior dispersion and signal precision are difficult to distinguish from learning speed ([Erdem and Keane, 1996](#); [Crawford and Shum,](#)

²⁷Hotel commissions in this industry range from 10 to 25% ([Ursu, 2018](#)). Expedia Group reported 2021 gross bookings of \$72.4 billion, revenue of \$8.6 billion, which leads to a company-wide revenue margin of 11.9% ([Expedia Group, Inc., 2022](#)).

2005; Ursu et al., 2024). The estimated α^C and implied subsequent learning are conditional on these calibrations. Appendix D.4 provides details on the implementation.

Inspection costs $c(\ell) = c_I(1 + \ell)^\psi$ has two parameters, c_I and ψ , with displayed slots ℓ counting advertisements. Reservation values solve (3.17) on the transformed reward, hotel by hotel. Under the timing and independence restrictions in Section 3.4.5, inspection updates the best-found offer while leaving entry forecasts fixed. Recalled identities are excluded from new inspections. Only first organic inspections on default or filtered lists supply monetization weight learning evidence, and updating follows a continuation choice.

4.1.6 Boundary values

The outside option $u_{0,it} = z'_{it}\vartheta$ depends on the trip's characteristics z_{it} including booking window, stay length, numbers of adults, children and rooms, and the check-in month. We estimate the standard deviation σ_B of the mean-zero boundary shocks. Continuation values in (3.14) include the list-opening cost c_{sess} , filter and sort intercepts ($\kappa_{\text{filter}}, \kappa_{\text{sort}}$) relative to default, and the same-tool shifter $\theta_a \mathbf{1}\{a = k_s\}$.

For filter and sort, x_s^{list} includes three standardized summary statistics that describe how monetized a current list looks: mean price quintile, the within-list slope of price quintile on log rank, and the advertisement share. The last term $\widehat{\mathcal{W}}_s^a$ is the expectation of the next session's value over the list the continuation action returns. Following De los Santos and Koulayev (2017), we take that expectation over lists the consumer expects to see: for a re-issued default query, fresh draws of the market's hotels in the order she anticipates the platform's rule produces; for a filter, the same draws with the share of hotels a filter keeps on average in the data; for a sort, the draws in price order. Appendix D.3 provides details on the implementation.

4.1.7 Hotel costs

Marginal costs recovered from the pricing first-order condition are related to hotel and trip characteristics by

$$c_{jt} = x'_j\delta + z'_t\delta_z + \epsilon_{jt}, \quad (4.6)$$

where x_j contains an intercept, star rating, review score and count, free cancellation and discount information, and the number of amenities. z_t contains booking window, stay length, rooms, and traveler size.

4.2 Identification Discussion

The identification problem is the search-market version of a classic simultaneity problem. Hotels price with knowledge of quality, and the platform ranks with expected demand. Observed price and rank therefore need not be independent of unobserved quality and taste shocks. Identification arguments combine conditional variation in prices, displayed slots, and organic orders with the model’s restrictions on information and choice as follows.

Assumption 4 (Identifying restrictions).

*(a) **Ranking representation.** The expected economic score uses the population price coefficient β_p and calibrated commission γp . Platform prediction, training target construction, ranker training, and the learned-score and ranking-noise conventions link candidate economic scores to observed-order probabilities exogenously.*

*(b) **Price and list variation.** Conditional on included hotel (group) effects, observed attributes, trip context, relevant prior actions, and list controls, the variation used to identify price responses is unrelated to omitted demand-relevant utility and ranking inputs. The constructed dollar price does not bring endogenous variations. Continuation choice and personalization are accounted for through the specified boundary model and observed histories.*

*(c) **Slots and signals.** Inspection cost follows the same displayed-slot function on every list. Sponsored placement is unrelated to omitted determinants of inspection. Sorted positions and sponsored slots provide no platform-generated organic-rank signal in the consumer’s model.*

Consumer parameters: preferences and slot costs. The regular price endogeneity concern is that price could be correlated with unobserved quality. In our context, repeated hotel appearances in different lists provide within-hotel variation in price. Hotel (group) effects further control for unobserved quality differences. Therefore, under Assumption 4(b), inspection and booking responses to this price input inform β_p and its heterogeneity, conditional on observables and hotel (group) effects.

Another identification concern is that consumers may self-select into deep inspection and booking either because they face low slot costs or because those hotels are more attractive. To distinguish inspection costs from unobserved utility, we exploit the gap between organic rank and displayed slot. Sponsored placement separates organic rank from the effort of reaching a hotel. At a fixed organic order, inserting advertisements can shift an organic hotel to different displayed slots across different lists without changing the rank signal. Holding its price, visible attributes, decoding inputs, and session-entry belief fixed, this displacement changes its inspection cost but not its direct utility or match forecast. Conditional exogeneity in Assumption 4(c), together with controls for the list

characteristics, makes the inspection response informative about c_I and ψ .²⁸

Perceived monetization versus slot costs. A deep inspection can reflect low inspection costs or weak reliance on the recommended ranking due to a high perceived monetization. The distinction is that inspection cost $c(\ell)$ only depends on the slot, whereas the expected return of rank decoding also depends on the occupying hotel’s attributes such as price. At a given slot, a price change leaves the inspection cost fixed but changes both visible utility and the match inferred from organic rank. For an expensive hotel near the top, a consumer attributing more weight to commission assigns more of that position to its price and less to match quality. The resulting price-by-rank variation helps distinguish perceived monetization from a common slot-cost gradient.

Table 2 shows the reduced-form evidence of this price-by-rank interaction that informs the perceived monetization. We regress a linear probability models of whether a displayed hotel is inspected on its price quintile, interacted with its rank, conditional on hotel (group) effects, rank-bin controls, and session effects. We focus on first-session behavior that anchors the perceived monetization weight, where no within-trip inspection evidence has yet accumulated and Assumption 1 prevents the order alone from changing the weight belief. We find that the inspection probability falls by 0.16 percentage points per quintile step for the top five ranks, and by 0.09 percentage points per quintile step in general (with all ranks entered as $-\log$).

Table 2. The price-by-rank interaction in first sessions

	(1)	(2)	(3)	(4)
Price quintile	−1.11 (0.13)	−1.60 (0.13)	−4.10 (0.68)	−5.41 (0.69)
Price quintile × top five ranks	−1.62 (0.61)	−2.25 (0.62)		
− log rank			3.54 (0.43)	5.96 (0.46)
Price quintile × − log rank			−0.86 (0.18)	−1.11 (0.18)
Hotel effects	yes	yes	yes	yes
Rank bins	yes	yes	yes	yes
Session effects		yes		yes

Notes: Linear probability models of whether a displayed hotel is inspected, restricted to the first session of each trip. Columns (1)–(2) interact price quintile with a top-five-rank indicator; columns (3)–(4) interact it with continuous negative log rank measure. All columns include hotel effects and rank-bin controls. Columns (2) and (4) additionally include session effects. Standard errors in parentheses are clustered by trip.

²⁸ Around 14% of organic listings have no advertisement above them, and around 67% have three or more.

Consumer-sorted lists further empowers this identification strategy. The idea is that sorting removes the platform’s ordering and the perceived monetization does not affect the inspection probability. The price-by-rank interaction should be absent on sorted lists, which serves as a benchmark. The additional inspection response to price-by-rank interaction on default and filtered lists, relative to sorted lists, informs the perceived monetization. Higher perceived monetization would suggest a stronger (i.e., more negative) inspection response to the price-by-rank interaction. Within a trip, the displayed hotels in default and sorted lists are the same, which controls for the list’s impact on inspections.

Figure 5 compares the price-by-rank interaction across list types. As expected, we find that the interaction has a significant impact on inspection on default and filtered lists, but is not significant on sorted lists. Besides, the small magnitude of the interaction’s impact suggests that the perceived monetization may not be large, which is consistent with the structural estimate of α^C in Section 5.1. Finally, Figure 5 also shows that the main effect of rank is positive and significant across all list types, which is consistent with the identification intuition that inspection costs are present on all lists and can be separated from the perceived monetization effect.




Figure 5. Price–position associations across list types

Notes: Coefficients from linear probability models of inspection on centered price quintile, negative log rank, and their interaction, estimated separately on default, filtered, and sorted lists with session fixed effects. Filled circles show the interaction; open squares show the negative-log-rank coefficient at the price centering value. Bars are 95% intervals from standard errors clustered by trip. On a sorted list, position follows the chosen sort key and is not a organic rank under the platform’s rule.

A new concern arises that the observed differences across list types may be due to different consumer composition in these lists rather than the presence of monetization. This motivates us to structurally model the choice to sort or filter as a decision to account for self-selection into these list types (De los Santos and Koulayev, 2017).

Unobserved inspection order. The data records the inspected set, not the order. The estimation needs to account for the latent order because the stopping decision is compared with what has been found. An identification concern is that the order and the belief cannot be separated. For instance, a consumer who thinks rank is informative about surplus may inspect from the top down, while a consumer who thinks rank is heavily monetized may inspect from the bottom up. The concern is valid if both scenarios fit the data equally well.

The pattern of skips separates them. Suppose a consumer inspected the first and the tenth hotel and skipped the eight between. Every compatible order must select the tenth over all eight skipped ones and then stopped with them uninspected. The index and stopping inequalities must therefore hold for those comparisons under the same persistent tastes and realized matches. Changing perceived monetization changes their decoded returns, in ways that also depend on prices and visible attributes. For instance, under a low perceived weight the decoded expected match falls steeply with rank, so the tenth hotel's reservation value is far below the skipped ones' and that choice is improbable in any order. By contrast, under a high perceived weight the decoded profile is flat and the choice is cheap to rationalize. The frequency of deep inspections against skipped shallow ones, and the prices of the hotels skipped, therefore inform the belief without the order being observed.

Actual ranking monetization. Conditional on demand parameters and the calibrated commission rate, the ranking's response to price reveals how the platform trades-off predicted consumer surplus against commission revenue, which informs the actual monetization weight. A higher price lowers the surplus component, holding other inputs fixed, but raises commission. Comparisons between hotels at different positions with similar predicted consumer surplus and different prices inform the weight on commission under Assumption 4(a). Besides, pairs of hotels where one hotel has higher surplus but lower price and the other has lower surplus but higher price are also informative.²⁹

The net price loading ϕ makes the trade-off explicit. Collecting the direct price terms in (3.3), with $\beta_p < 0$, gives

$$\phi = (1 - \alpha^P)\beta_p + \alpha^P\gamma \implies \alpha^P = \frac{\phi - \beta_p}{\gamma - \beta_p}. \quad (4.7)$$

Holding the other score components fixed, ϕ moves from the consumer price penalty β_p at $\alpha^P = 0$ to the commission reward γ at $\alpha^P = 1$, crossing zero at $\beta_p/(\beta_p - \gamma)$, and the actual weight is where on that segment the observed ranking sits. Monetization can weaken the penalty while price still lowers the score when the net loading remains

²⁹This conflicting pair of hotels is common in the data. Among pairs in the top ten positions, 87% conflict.

negative. Equation (4.7) shows that three terms pin down α^P . ϕ is from conditional ranking responses, β_p from demand, and γ from calibration.³⁰

4.3 Estimator

We use a two-step simulated likelihood procedure. Let

$$\theta = (\beta, \beta^p, \sigma_p, \xi, \beta^A, \alpha^C, c_I, \psi, \sigma_B, \vartheta, c_{\text{sess}}, \kappa, \theta, b)$$

collect the demand, belief, inspection cost, outside option, and boundary parameters. We first estimate demand conditional on the displayed rankings,

$$\hat{\theta} \in \arg \max_{\theta} \ell_{Y|r(\alpha^P)}(\theta), \quad (4.8)$$

where $\ell_{Y|r(\alpha^P)}$ is the log likelihood of the observed inspection sets, boundary actions, and any booked hotel, given the observed rankings. At this conditional estimate, we construct model-based matches, fit (4.4), form the targets $T_{ijts}(\alpha^P, \hat{\theta})$, and train and evaluate the LambdaMART ranker. We select the weight by its ranking likelihood:

$$\hat{\alpha}^P \in \arg \max_{\alpha \in \mathcal{A}} \ell_r(\alpha, \hat{\theta}), \quad (4.9)$$

where ℓ_r is the log likelihood of the observed orders.

4.3.1 Ranking likelihood

Suppress consumer-trip indices. The ranker learns the function $f_{\alpha, \theta}$ and supplies $g_{js}^R(\alpha, \theta)$, the index entering a latent score $g_{js}^R + \varepsilon_{js}^R$. Independent type-I extreme-value ranking disturbances ε_{js}^R with scale $\sigma_R = 1$ write the likelihood of the observed organic order $\mathbf{r}_s = (j_1, \dots, j_{J_s^R})$ as

$$R_s(\mathbf{r}_s \mid w_s, H_{s-1}^R) = \prod_{k=1}^{J_s^R} \frac{\exp(g_{j_k s}^R / \sigma_R)}{\sum_{l \in \mathcal{R}_k} \exp(g_{l s}^R / \sigma_R)}, \quad (4.10)$$

where J_s^R counts the organic ranked hotels, w_s contains recorded non-order list inputs, H_{s-1}^R the recorded pre-list history, and $\mathcal{R}_k$ the remaining organic hotels. The criterion ℓ_r sums these log probabilities on the evaluated default and filtered lists. Sorted orders and

³⁰We can show that the ratio of the two weights, α^C / α^P , does not depend on γ that is calibrated, nor on the scale of utility.

sponsored placements do not contribute to the ranking likelihood. Appendix D.1 gives the training and evaluation details.

4.3.2 Inspection likelihood

Let Λ_{it} contain the persistent price heterogeneity, pre-search tastes $\tilde{e}_{ijt}$, and hotel–trip match innovations ν_{ijt} . For a draw Λ_{it} and a completed session-entry history h_s , the consumer’s visible utilities and inspected offers are determined. Her forecasts for uninspected hotels use the subjective law.

Let $\mathcal{C}_s$ be the newly inspected set and $\pi_s = (j_1, \dots, j_n)$ a compatible order. After k inspections, $\bar{u}_k$ is the best utility among all retained offers and NI_k is the remaining inspectable set. Proposition 1 gives the conditional inspection likelihood

$$I_s(\pi_s \mid \Lambda, h_s) = \prod_{k=1}^n \mathbf{1} \left\{ \zeta_{j_k s} > \Psi_s(\bar{u}_{k-1}), j_k = \arg \max_{l \in \text{NI}_{k-1}} \zeta_{ls} \right\} \times \mathbf{1} \left\{ \Psi_s(\bar{u}_n) \geq \max_{l \in \text{NI}_n} \zeta_{ls} \right\}. \quad (4.11)$$

4.3.3 Boundary action and booking likelihoods

Boundary shocks are observed only after stopping. Integrating these shocks, conditional on the same draw Λ_{it} , history h_s , and compatible order π_s , gives the boundary action likelihood

$$A_s(a_s \mid \Lambda, h_s, \pi_s) = \frac{\exp(V_s^{a_s}(\bar{u}_n; \omega_s)/\tau_B)}{\sum_{a \in \mathcal{A}_s} \exp(V_s^a(\bar{u}_n; \omega_s)/\tau_B)}, \quad (4.12)$$

for a_s in the feasible action set $\mathcal{A}_s$, and zero otherwise. We have $V_s^{\text{book}} = \bar{u}_n$, $V_s^{\text{quit}} = u_0$, and continuation values are (3.14).

When booking occurs at the j^* hotel, this contributes to the booking likelihood

$$B_s(j^* \mid \Lambda, h_s, \pi_s) = \mathbf{1} \left\{ j^* \in \mathcal{K}_s, j^* = \arg \max_{j \in \mathcal{K}_s} U_{ijt} \right\}, \quad (4.13)$$

where $\mathcal{K}_s$ contains all inspected hotels. On a nonbooking boundary, $B_s = 1$.

4.3.4 Demand likelihood

Let $\mu_s(d\Lambda, dh_s)$ be the researcher’s conditional law of the persistent draw and unobserved earlier inspection history, given H_{s-1}^R , w_s , and $\mathbf{r}_s$. This law preserves selection from earlier

outcomes and the current observed list. For the session outcome Y_s , define

$$G_s = I_s A_s B_s,$$

with every factor evaluated at the same draw and path. We have

$$\begin{aligned} d_s(Y_s \mid \mathbf{r}_s, w_s, H_{s-1}^R) &= \int \sum_{\pi_s \in \Pi(\mathcal{C}_s)} G_s(Y_s, \pi_s \mid \Lambda, h_s) \mu_s(d\Lambda, dh_s), \\ \ell_{Y|r} &= \sum_{i,t,s} \log d_{its}. \end{aligned} \tag{4.14}$$

The likelihood factors are multiplied before integration. After a continuation choice, each compatible history carries its offers forward and updates the consumer's weight belief using organic inspections. Appendix D.2 details the computation.

4.3.5 Inference

The reported MLE standard errors are clustered by trip. However, the uncertainty from the demand estimates may propagate through the ranking stage. We use bootstrap resampling of trips to construct confidence intervals. This calculation is in progress.

5 Estimation Results

5.1 Parameter Estimates

Table 3 reports the estimates. Panel A contains preferences, match, entry-belief, inspection-cost, and the monetization weight. Panel B contains the boundary stage parameters including standard deviation of the boundary shocks, the outside option and continuation parameters.

5.1.1 Actual and perceived monetization

The estimated actual monetization weight is $\hat{\alpha}^P = 0.78$. The ranking places the direct price loading in (4.7) 78% of the way from the population price penalty β_p toward the commission reward γ . At the calibrated commission rate $\gamma = 0.15$, the net price loading remains negative in the score, but monetization substantially weakens the price penalty.³¹ By contract, the mean of consumers' trip-entry belief is only $\hat{\alpha}^C = 0.25$. Their ratio is

³¹The net price loading is $(1 - \alpha^P)\beta_p + \alpha^P\gamma = 0.22 \times (-1.27) + 0.78 \times 0.15 = -0.16$.

$\hat{\alpha}^C/\hat{\alpha}^P = 0.32$. Consumers behave as if the monetization weight were about one-third of its actual value. Therefore, every rank is credited with more match than it carries, and an expensive hotel near the top is discounted less than it should be.

The belief gap matters for the value consumers attach to inspecting a recommendation. An expensive hotel’s high position is more readily attributed to a favorable match when consumers expect little commercial influence. Figure 4 illustrates the implication on the sample lists. Evaluating the decoder at the estimated actual weight yields a lower and flatter expected-match profile than evaluating it at the perceived weight, rounded to 0.25 in the figure. It shows how the same prominent offers can appear more promising under the estimated entry belief.

5.1.2 Visible preferences and price heterogeneity

The visible-utility estimates assign positive value to star rating, review score, and review count. The star-rating coefficient is 0.49, compared with 0.07 for review score and 0.28 for log review count. The population price coefficient is $\hat{\beta}_p = -1.27$, with substantial dispersion $\hat{\sigma}_p = 1.46$. These attributes are known before inspection and therefore help determine both which hotels are worth inspecting and which retained offer is preferred.

5.1.3 Matches revealed by inspection

Amenities contribute to the hotel–trip match in (4.3). The baseline services coefficient is positive, at 1.50. The estimates do not show significant effects for free extras, leisure, or home-style amenities, or for the interactions of children with leisure, nights with home-style features, and party size with services.

5.1.4 Inspection costs

The estimated slot-cost function in (3.12) has $\hat{c}_I = 0.03$ and $\hat{\psi} = 0.71$. The reported costs are 0.04 utility units at the first displayed slot, 0.14 at the tenth, 0.30 at the thirtieth, and 0.69 at the hundredth. Costs rise with depth, with smaller increments at deeper slots. The increment per slot is about \$1.20 at the top of the list and \$0.70 around the tenth slot, the same order as the \$1.92 per position that Ursu (2018) recovers from randomized rankings on this platform. Because slots count advertisements, prominence reduces inspection expenditure separately from the information conveyed by organic rank. A deeper hotel must offer a sufficiently attractive distribution of potential returns to compensate for its higher inspection cost.

5.1.5 The outside option and boundary shocks

The outside option is the value of ending the trip without an Expedia booking, including both booking elsewhere and abandoning the trip. Longer booking windows, longer stays, and more rooms raise this value conditional on the other regressors. The estimated standard deviation $\hat{\sigma}_B = 1.72$ captures variation in boundary payoffs beyond their deterministic values.

5.1.6 Continuation and reuse

Opening another list costs $\hat{c}_{\text{sess}} = 6.15$ utility units. Its attractiveness also depends on the expected return to further search, the best-found offer, and the outside option. Relative to re-issuing the default query, the filter and sort intercepts shift their deterministic values by -2.85 and -4.44 . The reuse term $\theta_a \mathbf{1}\{a = k_s\}$ then lowers the value of another default list by 0.43 when the current list is default, but raises the value of reusing a filter by 3.03 and of reusing a sort by 1.29 . The positive filter and sort terms favor persistence in those tools.

Current-list composition further affects refinement. A one-standard-deviation increase in mean price quintile raises sorting value by 0.62 and filtering value by 0.78 . The negative price–rank coefficients, -0.26 for sorting and -0.48 for filtering, likewise raise these values when expensive hotels are concentrated near the top.³² A larger advertisement share lowers sorting value, with coefficient -0.80 , whereas its filtering coefficient is not statistically significant. These list responses operate at the session-entry belief.

5.2 Model Fit

The model reproduces search intensity: 5,544 inspections predicted against 5,822 observed. Figure 6 compares the predicted and observed profiles by rank. The continuation block reproduces how often consumers ask for another list. Of first sessions, 27.3% are predicted to be followed by another against 30.7% observed. After a first session with no inspection, 38.3% of consumers open another list against 39.6% observed; after one with an inspection, 11.8% against 11.7%.

³²The structural covariate is the within-list slope of price quintile on positive log rank, standardized before entering continuation values. A more negative slope places expensive offers nearer the top.

Table 3. Parameter estimates

	Estimate	(s.e.)
<i>Panel A. Preferences, inspection, and monetization</i>		
Ranking and belief		
Monetization weight, α^P	0.78	(0.007)
Perceived monetization weight, α^C (mean of the entry belief)	0.246	(0.032)
Ratio, α^C / α^P	0.316	(0.041)
Visible utility		
Constant	1.374	(0.239)
Star rating	0.493	(0.017)
Review score	0.068	(0.015)
Number of reviews, $\log^\dagger$	0.276	(0.029)
Free cancellation or discount badge	0.859	(0.063)
Nightly dollar price, β_p	-1.270	(0.058)
Dispersion of the price coefficient, σ_p	1.464	(0.058)
Hotel effects, dispersion †	0.791	
Match: amenity groups and taste interactions		
Services (bar, gym, spa, laundry) †	1.497	(0.441)
Free extras (breakfast, parking, wifi, transfer) †	-0.219	(0.166)
Leisure (pools, hot tub) †	-0.082	(0.168)
Home-style (air conditioning, washer and dryer, pets) †	0.297	(0.449)
Children $\times$ leisure †	-0.278	(0.316)
Nights $\times$ home-style †	0.067	(0.142)
Party size $\times$ services †	-0.235	(0.192)
Slot cost		
Intercept, c_I	0.026	(0.014)
Slope, ψ	0.711	(0.061)
<i>Panel B. Boundary values</i>		
Standard deviation of boundary shocks, σ_B	1.718	(0.067)
Outside option		
Booking window, per 100 days	1.258	(0.055)
Length of stay, per 10 nights	1.462	(0.149)
Adults, per 10	0.762	(0.340)
Children, per 5	0.155	(0.320)
Rooms, per 2	3.869	(0.285)
Check-in month	-1.035	(0.083)
Continuation		
Continuation cost, c_{sess}	6.151	(0.194)
Filter intercept, κ_{filter}	-2.845	(0.305)
Sort intercept, κ_{sort}	-4.440	(0.323)
Reuse: default, θ_{def}	-0.430	(0.130)
Reuse: filter, θ_{filter}	3.032	(0.194)
Reuse: sort, θ_{sort}	1.290	(0.262)
Sort: mean price quintile of the list	0.624	(0.081)
Sort: price-rank slope of the list	-0.264	(0.059)
Sort: share of advertisements	-0.800	(0.081)
Filter: mean price quintile of the list	0.779	(0.104)
Filter: price-rank slope of the list	-0.482	(0.074)
Filter: share of advertisements	0.081	(0.057)

Notes: Standard errors in parentheses are clustered by trip.




Figure 6. Inspection by rank, observed and predicted

5.3 The Belief Channel

5.3.1 Decoding in the fitted behavior

Suppressing rank decoding at the estimated parameter vector worsens the reported fit of inspections by 16% and bookings by 25%. Booking fit changes because the interpretation of rank changes which offers consumers discover, even though their matches are known when booking occurs. Thus the informational role of position matters for transactions as well as for inspection.

5.3.2 Learning within the trip

The fitted model implies slow updating. The reported mean is 2 inspections per session conditional on at least one inspection. We find an average update of around 0.5% of the distance to the platform's weight, with a tail figure of around 7.5%. Consider that each consumer conducts in average 2–3 sessions per trip, in average a trip would close only 4–6% of the gap to the platform's weight.

A possible explanation for small updates comes from the platform's design. Indeed, a large amount of the information the platform uses to rank hotels is displayed on the list page to retain the consumer's attention, which the consumer sees it before any click. Thus the consumer's click reveals relatively little about the ranking.

5.4 Marginal Costs

Marginal costs are recovered from the hotel pricing first-order condition (3.20) and related to hotel and trip characteristics through (4.6). They are pricing-implied opportunity costs, rather than accounting expenditures. Their monetary scale is the 2013-dollar scale of the constructed nightly prices. Table 4 reports the estimates of the marginal cost shifters.

Conditional on the remaining regressors, one additional star is associated with a \$16.10 higher cost and one additional review-score point with \$11.51. Amenity count also has a positive loading. The fitted cost averages \$70.9 per occupied room-night, with a 5th percentile of \$41.2 and a 95th percentile of \$106.9. McClure (2025) recovers marginal costs for U.S. hotels from a structural model with capacity constraints and reports, for 2015 to 2018, averages of \$44.5 for economy hotels, \$55.0 for midscale, \$80.8 for upper midscale, \$93.7 for upscale, and \$122.8 for upper upscale, which are of a similar order of magnitude.

Table 4. Marginal-cost

	Coefficient	(s.e.)
Star rating	16.10	(0.49)
Review score	11.51	(0.63)
Number of reviews, log	−3.64	(0.23)
Free cancellation or discount badge	−14.52	(2.03)
Number of amenities	0.58	(0.18)
Booking window, per day	0.02	(0.02)
Length of stay, per night	1.67	(2.99)
Rooms	−17.81	(11.43)
Guests	−0.84	(0.35)
Constant	10.64	(1.20)

Notes: Standard errors in parentheses are clustered by trip.

6 Counterfactual Analysis

The belief gap affects discovery before it affects transactions because consumers must decide which offers to inspect before they can choose among their realized matches. This misperception favors platform monetization and softens hotels’ price competition. We consider two interventions that could reduce the belief gap: a neutral-ranking policy and a transparency policy. The former removes the monetization weight from the ranking, while the latter reveals it to consumers. Whether those policies are substitutes or complements is not obvious. They may be substitutes if informed consumers can discount commercial steering themselves, reducing the need to change the ranking. But they may instead be complements if neutrality improves recommendations while uncorrected beliefs cause

consumers to misinterpret them. To answer this question, we evaluate the two interventions separately and together, first at fixed prices and then after hotels re-optimize their prices.

6.1 Policy Experiments

The neutral-ranking policy sets $\alpha^P = 0$ in the economic ranking criterion (3.4), removing commission revenue so that the platform orders by predicted net surplus alone. Ideal transparency instead leaves the ranking rule unchanged and reveals the actual monetization weight. The belief b_{its} is therefore a point mass at α^P in every session. This is knowledge of commercial influence, rather than of uninspected matches or the platform’s private predictions.

These experiments resonate with two instruments in policy discussions. The CMA’s 2019 hotel-booking principles call for clear, prominent disclosure when remuneration can affect search rankings ([Competition and Markets Authority, 2019](#)). EU consumer law also requires information on the main ranking parameters and paid prominence ([Competition and Markets Authority, 2019](#); [European Union, 2019](#); [Netherlands Authority for Consumers and Markets, 2020](#)). Direct restrictions on ranking have a related precedent in the European Commission’s May 2026 proposal on multimodal transport booking, whose Article 7 would require neutral display of offers and exclude paid prominence from default ranking criteria ([European Commission, 2026](#)). These examples motivate the distinction between informing consumers and changing the ordering. The transparency policy is a benchmark for effective information about the monetization weight, related to the general requirement to describe ranking parameters. The neutral-ranking policy removes the commission revenue component of organic hotel rankings, related to the proposal’s restriction on paid prominence in transport.

We compare four regimes: the monetized rule with inherited entry beliefs, (M, E) , which is the status quo; neutrality with those beliefs, (N, E) ; the monetized rule with transparency, (M, T) ; and neutrality with transparency, (N, T) . The monetized weight is 0.78. Inherited beliefs use the estimated trip-entry distribution with mean 0.25 and the subsequent updating in (3.10). Each regime is evaluated at fixed prices and at re-solved equilibrium prices. The fixed-price comparison holds the baseline constructed nightly rates unchanged. In the second comparison, hotels simultaneously best respond to rivals’ prices under the imposed rule and belief regime, using the profit definition and pricing equilibrium in Section 3.5. Table 5 reports percentage changes from (M, E) , the status quo, within each price regime. Appendix E gives the outcome formulas and the references for the additional comparisons.

Table 5. Neutral ranking vs. transparency: counterfactual outcomes

	Percentage change from the status quo		
	Neutral ranking	Transparency	Both
<i>Panel A. Prices fixed</i>			
On-platform bookings	+27.0	−2.2	+29.3
Inspections	+6.9	−7.8	+8.6
Cost per inspection (utility units)	+0.0	+15.2	−4.2
Total inspection cost (utility units)	+6.9	+6.1	+4.1
Outcome surplus (utility units)	+1.5	−0.3	+1.4
Surplus net of inspection costs (utility units)	+1.1	−0.8	+1.2
Average booked nightly rate (\$)	−25.4	−4.7	−26.4
Commission revenue (\$)	−5.1	−6.8	−4.9
<i>Panel B. Prices re-solved</i>			
On-platform bookings	+24.1	−2.8	+26.3
Inspections	+6.9	−7.8	+8.7
Cost per inspection (utility units)	+0.0	+15.2	−4.2
Total inspection cost (utility units)	+6.9	+6.1	+4.1
Outcome surplus (utility units)	+1.1	−0.4	+1.0
Surplus net of inspection costs (utility units)	+0.7	−0.9	+0.8
Average booked nightly rate (\$)	−21.4	−3.8	−22.3
Commission revenue (\$)	−3.1	−6.5	−2.7

Notes: Each number is a percentage change from monetized ranking with inherited entry beliefs, (M, E) , within the panel’s price regime. The columns are neutrality with those beliefs, (N, E) ; the monetized rule with transparency, (M, T) ; and neutrality with transparency, (N, T) . All use the same trip population, including noninspectors and nonbookers. Cost per inspection pools inspection expenditure and counts. Outcome surplus includes booked-hotel or outside-option utility while the inspection-net measure subtracts inspection costs. Nightly rates use the constructed 2013-dollar scale.

6.2 A Neutral Ranking

6.2.1 At fixed prices

Removing commission from the ranking criterion promotes offers with higher predicted consumer surplus. Their greater prominence can make them less costly to reach and, at the inherited belief, more encouraging to inspect. At fixed prices, neutrality raises on-platform bookings by 27.0% and lowers the average booked nightly rate by 25.4%. Figure 7 shows a reallocation of booked hotels: the two cheapest quintiles’ share of bookings rises from 24% to 63%, while the most expensive quintile’s share falls from 32% to 9%. Because hotel rates are unchanged, the lower booked-rate mean comes entirely from choosing different offers. Improved recommendation thus redirect consumers to cheaper offers already available but were hidden before.

Inspections and total inspection expenditure both rise by around 6.9%. Outcome surplus rises by 1.5%, and surplus net of inspection costs by 1.1%. This shows that better recommendations do not necessarily reduce information acquisition. They can make additional inspections worthwhile and support more transactions.

Platform’s commission nevertheless falls by 5.1% despite the market expansion. The




Figure 7. Booking reallocation at fixed prices

Notes: Shares of on-platform bookings by the booked hotel’s price quintile under monetized ranking and neutrality, both with inherited entry beliefs and fixed hotel prices. Solid bars show the status quo, (M, E) ; hatched bars show (N, E) .

commission loss per additional booking is about \$11. A commission-funded intermediary benefits from matching consumers to products, but also from selecting which successful matches it intermediates. Commercial steering under a monetized ranking can trade away transaction creation for transaction value. Neutral ranking changes that trade-off. Our results show that the current booking expansion does not offset the lower average commission per booking. This is the underlying limitation of the neutral-ranking policy that the transparency policy will later address.

6.2.2 The role of decoding

Suppressing rank decoding at the estimated parameter vector reduces neutrality’s fixed-price booking gain from 27.0% to 12.9% and its commission decline from 5.1% to 2.5%. Suppressing decoding thus roughly halves the booking response. This shows the importance of the interpretation of rank. With rank decoding, promoting an attractive hotel up changes the inspection cost but also what consumers expect to discover through the perceived distribution of inspection returns. The platform thus does not merely sell access to attention but also supplies an interpretation of where that attention should go. The same inference mechanism that steers demand under the monetized rule also amplifies the response to a more consumer-oriented ordering. Eliminating reliance on rankings is therefore not equivalent to eliminating only the harmful part of steering. This is further illustrated by the transparency-only exercise below.

6.2.3 Equilibrium repricing

Hotels' price responses attenuate the effects. Bookings rise by 24.1% rather than 27%, the average booked nightly rate falls by 21.4% rather than 25.4%, and surplus net of inspection costs rises by 0.7% rather than 1.1%. The commission decline also narrows to 3.1% compared to 5.1%.

Changes in recommendations lead to market expansion and shift consumer inspections, which reshape hotel substitution and demand, against which hotels reprice. Their responses feed back into both the utility and the visibility of offers. A monetized ranking can thus make the market expensive by controlling which existing offers enter effective consideration, and soften the price competition among them.

We find that most cheap hotels gain demand in general and hold their prices, because the customers they gain are the marginal ones whose demand could be more elastic. Expensive hotels lose demand in general and raise prices, because the inframarginal customers they keep are the ones who sought them out deep in the list and could be less price-sensitive. The cheap-hotel side reflects the discipline that a price-sensitive design imposes on sellers in [Dinerstein et al. \(2018\)](#); [Eisfeld \(2023\)](#). The expensive-hotel side reflects the pricing power that [Kaye \(2026\)](#) finds when a recommendation moves demand towards a seller. However, the overall effect is that the average booked nightly rate falls, because neutrality shifts bookings to cheaper hotels more and promotes cheaper hotels, which disciplines their prices and stimulates competition.

6.2.4 Who bears monetization more?

Figure 8 compares the outcome change from the neutral-ranking to the status quo monetized ranking at fixed prices. Booking reductions are 18–24% across groups. Reductions in surplus are more dispersed, with larger losses for short, last-minute, solo, and cheap trips. In other words, monetization “taxes” most heavily the trips with the least surplus. The reason is that for a short, solo, or last-minute trip, the cheapest acceptable hotel is often the best option, but the monetized ranking often hides it due to a low commission revenue.

6.3 Transparency of the Monetization Weight

Transparency leaves the monetized rule in place but changes the inference drawn from it. Greater perceived commercial influence attributes an expensive hotel's high position more to commission and less to match quality. On the sample lists, Figure 4 illustrates the decoded-match profile under the revealed $\alpha^C = \alpha^P = 0.78$ as lower and flatter near the top. This weakens the attraction of prominent recommendations.




Figure 8. Bookings and surplus: monetization relative to neutral-ranking

At fixed prices, transparency reduces bookings by 2.2% and inspections by 7.8%. The remaining inspections cost more on average: cost per inspection rises by 15.2%, and total inspection expenditure rises by 6.1%. Given the weakened attraction, consumers find fewer hotels worth inspecting and exit. The remaining inspections are concentrated in deeper, costlier slots. Fewer inspections therefore coexist with greater expenditure. Outcome surplus falls by 0.3%, and surplus net of inspection costs falls by 0.8%. Correcting the perceived weight changes how consumers search, but does not supply the easier access to attractive alternatives that neutrality provides. Consumer sophistication therefore is not a frictionless substitute for a better recommendation environment.

Commission falls by 6.8% in this fixed-price comparison. We call the commission difference between inherited beliefs and transparency platform's informational rent from the belief gap. The transparency results are qualitatively invariant to hotel repricing. Bookings fall by 2.8%, the average booked nightly rate by 3.8%, surplus net of inspection costs by 0.9%, and commission by 6.5%.

6.4 The Two Levers Together

Neutral-ranking improves the ordering but only expands bookings to a limited extent because consumers misinterpret the ranking rule. In fact, neutrality reverses the direction of the inherited belief error. Consumers who initially overtrust the monetized ranking now

become overly skeptical of the neutral ranking. They then give the improved recommendations less credit for match quality than they should. Adding transparency reduces this excess skepticism and strengthens reliance on the promoted offers. Transparency therefore discourages reliance under monetization but encourages it under neutrality. Its effect depends on what the ranking actually recommends.

Relative to neutrality alone, adding transparency increases bookings by approximately 1.8% in either price regime. At fixed prices it also lowers total inspection expenditure by approximately 2.6% relative to neutrality, even though inspections increase. The cost per inspection falls as consumers use the more promising offers in less costly slots. This is the contrast with transparency under monetization, where inspections decline but expenditure rises. Information about a better rule can economize on where discovery takes place rather than simply induce skepticism about all recommendations.

The total joint-policy booking gain is 29.3% at fixed prices and 26.3% after repricing. The joint policy raises surplus net of inspection costs by 1.2% and 0.8%, respectively, compared with 1.1% and 0.7% under neutrality alone. Its average booked nightly rate is 26.4% lower at fixed prices and 22.3% lower after repricing. This shows that the two levers are complements: a neutral ranking performs better when consumers know it to be neutral, and transparency is less costly when the ranking is neutral.

The complementarity concerns the interaction of recommendations and their use. Changing beliefs alone leaves the monetized ordering intact; changing the ordering alone leaves the inherited biased interpretation in place. Combining the interventions improves discovery through both margins. Appendix [E.4](#) provides additional evidence of complementarity.

6.5 The Revenue Comparison Between Ranking Rules

The same ranking-belief alignment mechanism also changes the platform's commission return from monetized versus neutral rankings. With inherited biased beliefs, neutrality earns the platform 5.1% less commission than monetization at fixed prices and 3.1% less after repricing. However, when transparency applies in both ranking regimes, neutrality instead earns 2.0% more than monetization at fixed prices and 4.1% more after repricing. Figure 9 makes the distinction visible.

Transparency removes demand supported by underperceived commercial influence under monetization, while recovering demand discouraged by excess skepticism under neutrality. Once consumers know which rule they face, the platform's commission from additional neutral-ranking bookings is sufficient to outweigh the lower booked rates.




Figure 9. Bookings and commission under the four regimes

Notes: Each point reports the changes in on-platform bookings and commission revenue from Table 5, relative to (M, E) status quo, within the same price regime. Filled markers hold hotel prices fixed; open markers use equilibrium prices. “Correct beliefs” denotes transparency under the monetized rule, (M, T) ; “Both” denotes the joint-policy intervention, (N, T) . The origin is the corresponding status quo. The revenue comparison reported in the paper compares the the yellow and green points.

Consumers’ interpretation affects not only which transactions occur, but also the platform’s relative return to generating transactions with useful recommendations rather than favoring higher-commission offers. This implies that a platform may have an incentive to implement a less monetized ranking if consumers are informed about the monetization weight. Effective transparency can therefore make commercial steering less attractive relative to facilitating transactions.

6.6 Economic Implications

The distortion operates through discovery. Consumers learn realized match quality before booking. The belief gap therefore changes which matches become available for choice, rather than leaving consumers mistaken about the hotel they buy. A ranking shapes this discovery process through both access costs and the perceived value of inspecting an offer. More inspections can accompany better recommendations when improved search opportunities make further discovery worthwhile. Hotels, in turn, compete for demand after it has been filtered by search, so changes in discovery alter the demand levels and elasticities against which they price. Neutrality’s lower booked rates arise substantially from demand reallocation across hotels. A lower average rate need not therefore reflect lower prices at every hotel.

This pricing channel also creates scope for a search externality across consumers. When consumers face common hotel prices, changes in one group's interpretation of the ranking can alter the demand elasticities hotels face and, through equilibrium repricing, affect consumers whose own beliefs remain unchanged. This mechanism is related to [Salz \(2022\)](#), who studies how intermediation affects the prices paid by buyers who do not use intermediaries. In our setting, the potential spillover arises from consumers' interpretation of recommendations. Its sign depends on how belief correction reallocates search and changes hotel-level demand elasticities. Our counterfactuals incorporate equilibrium price adjustment, but do not separately quantify this cross-consumer effect. This can be achieved by modelling consumers with heterogeneous beliefs and simulating the resulting hotel pricing equilibrium, which is being calculated.

Policy must align recommendations and reliance. The same uncorrected belief has opposite implications under the two ranking rules. Under monetization, it understates commercial influence and supports excessive reliance on the ranking. Under neutrality, it overstates commercial influence and limits reliance on a more consumer-oriented ordering. The two policies address different discrepancies: neutrality changes the recommendations supplied, whereas transparency changes how consumers use them. Their complementarity follows from the reversal in disclosure's effect across the two rules. At the monetized rule, correcting overtrust does not move attractive alternatives into easier-to-reach positions. At the neutral rule, correcting excessive skepticism makes improved recommendations more effective: relative to neutrality alone, adding transparency raises bookings while lowering aggregate inspection expenditure. Reducing the platform's informational rent is therefore not equivalent to improving consumers' outcomes, and correcting beliefs is not a substitute for improving the recommendation. In these counterfactuals, the policies reinforce one another because consumers respond more appropriately to the change in the ranking rule.

Beliefs change the return to recommendation quality. The commission comparison connects these consumer responses to platform incentives. Transparency removes demand support from underperceived monetization and restores part of the demand lost to excessive skepticism under neutrality. The additional bookings generated by the neutral rule then outweigh their lower commission per transaction. This reverses the revenue ordering of the two rules. The broader lesson is that misperception weakens the link between useful recommendations and the revenue they generate. Effective transparency can strengthen that link through demand, making useful intermediation more rewarding relative to commercial steering without changing the commission contract.

7 Conclusion

Platform recommendations help consumers navigate large product assortments, but the intermediary supplying this information may earn more from some transactions than from others. The consequences of this conflict depend on how consumers interpret the advice they receive. Using Expedia search and booking logs, we develop and estimate a structural model that separates the platform’s actual ranking monetization from consumers’ perceived monetization. The model connects an economically interpretable ranking rule to consumers’ inference from rank, their inspection and refinement decisions, and hotels’ equilibrium pricing.

We find that consumers substantially underperceive the platform’s ranking monetization. To cure this misperception, we evaluate two interventions that are in debate: a neutral-ranking policy that removes commission from the ranking criterion, and a transparency policy that reveals the actual monetization weight. The two policies are complements rather than substitutes. Neutrality improves the recommendations supplied, while transparency changes how consumers use them. The complementarity arises because the two instruments address different parts of discovery: the promoted option and the interpretation of that promotion, so the effect of changing beliefs depends on the rule consumers face. Finally, we show that consumer understanding of the ranking also changes the platform’s relative return to different ranking rules. A platform may have an incentive to implement a less monetized ranking if consumers are informed about the monetization weight.

These findings motivate two directions for further research. The first is to endogenize the platform’s ranking and disclosure choices, allowing competition for consumers and reputation across trips to shape those choices. The revenue reversal between the two rules provides a starting point for studying whether transparency induces platforms to choose a less monetized design. The second is to measure how beliefs respond to actual disclosures and accumulated search experience. Experiments that vary both ranking quality and disclosure could distinguish learning to use better recommendations from becoming uniformly more skeptical of rankings. Together, these directions would connect the transparency benchmark studied here to the design of information policies and to the recommendations platforms choose to supply, with consumers’ ability to recognize and use them.

References

- Abaluck, Jason, Giovanni Compiani, and Fan Zhang (2026) “A Method to Estimate Discrete Choice Models That Is Robust to Consumer Search,” *Journal of Political Economy*, 134 (7), 1967–2022, [10.1086/740223](https://doi.org/10.1086/740223).
- Abdollahpouri, Himan, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato (2020) “Multistakeholder Recommendation: Survey and Research Directions,” *User Modeling and User-Adapted Interaction*, 30 (1), 127–158, [10.1007/s11257-019-09256-1](https://doi.org/10.1007/s11257-019-09256-1).
- Armstrong, Mark, John Vickers, and Jidong Zhou (2009) “Prominence and Consumer Search,” *RAND Journal of Economics*, 40 (2), 209–233, [10.1111/j.1756-2171.2009.00062.x](https://doi.org/10.1111/j.1756-2171.2009.00062.x).
- Athey, Susan and Glenn Ellison (2011) “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, 126 (3), 1213–1270, [10.1093/qje/qjr028](https://doi.org/10.1093/qje/qjr028).
- Boleslavsky, Raphael, Thomas Jungbauer, and Mehdi Shadmehr (2026) “Algorithm Transparency and Search Manipulation: Steering vs. Persuasion,” [10.48550/arXiv.2608.12558](https://arxiv.org/abs/10.48550/arXiv.2608.12558), Working paper, August.
- Burges, Christopher J. C. (2010) “From RankNet to LambdaRank to LambdaMART: An Overview,” Technical Report MSR-TR-2010-82, Microsoft Research, <https://www.microsoft.com/en-us/research/publication/from-ranknet-to-lambdarank-to-lambdamart-an-overview/>.
- Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello (2026) “Artificial Intelligence, Algorithmic Recommendations and Competition,” Working paper.
- Cao, Zhe, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li (2007) “Learning to Rank: From Pairwise Approach to Listwise Approach,” in *Proceedings of the 24th International Conference on Machine Learning*, 129–136, [10.1145/1273496.1273513](https://doi.org/10.1145/1273496.1273513).
- Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim (2018) “Consumer Search and Price Competition,” *Econometrica*, 86 (4), 1257–1281.
- Cogley, Timothy and Thomas J. Sargent (2008) “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49 (1), 185–221, [10.1111/j.1468-2354.2008.00477.x](https://doi.org/10.1111/j.1468-2354.2008.00477.x).
- Competition and Markets Authority (2019) “Online Hotel Booking: Principles for Businesses,” Technical report, Competition and Markets Authority, London, <https://www.gov.uk/government/publications/online-hotel-booking-principles-for-businesses>.
- Compiani, Giovanni, Gregory Lewis, Sida Peng, and Peichun Wang (2024) “Online Search and Optimal Product Rankings: An Empirical Framework,” *Marketing Science*, 43 (3), 615–636, [10.1287/mksc.2022.0071](https://doi.org/10.1287/mksc.2022.0071).
- Crawford, Gregory S. and Matthew Shum (2005) “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73 (4), 1137–1173.
- Cunha, Tiago, Ioannis Partalas, and Phong Nguyen (2021) “Juggler: Multi-Stakeholder Ranking with Meta-Learning,” in *Proceedings of the 1st Workshop on Multi-Objective Recommender Systems (MORS 2021)*, 2959 of CEUR Workshop Proceedings, <https://ceur-ws.org/Vol-2959/paper3.pdf>.

- De los Santos, Babur and Sergei Koulayev (2017) “Optimizing Click-Through in Online Rankings with Endogenous Search Refinement,” *Marketing Science*, 36 (4), 542–564, [10.1287/mksc.2017.1036](https://doi.org/10.1287/mksc.2017.1036).
- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan (2018) “Consumer Price Search and Platform Design in Internet Commerce,” *American Economic Review*, 108 (7), 1820–1859, [10.1257/aer.20171218](https://doi.org/10.1257/aer.20171218).
- Eisfeld, Luise (2023) “How Do Online Product Rankings Influence Sellers’ Pricing Behavior?”, https://luiseeisfeld.github.io/assets/docs/HotelRankings_Eisfeld_TSE.pdf, Working paper, December 11.
- Erdem, Tülin and Michael P. Keane (1996) “Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 15 (1), 1–20.
- European Commission (2026) “Proposal for a Regulation of the European Parliament and of the Council on Multimodal Booking and Repealing Regulation (EC) No 80/2009,” Technical Report COM(2026) 231 final, European Commission, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026PC0231>, Procedure 2026/0113(COD), Brussels, 13 May 2026.
- European Union (2019) “Directive (EU) 2019/2161 of the European Parliament and of the Council of 27 November 2019 as regards the Better Enforcement and Modernisation of Union Consumer Protection Rules,” <https://eur-lex.europa.eu/eli/dir/2019/2161/oj>, Official Journal of the European Union, L 328, 18 December 2019, pp. 7–28.
- European Union (2022) “Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services (Digital Services Act),” Official Journal of the European Union, L 277, 27 October 2022, Article 27.
- Expedia (2026a) “How Our Sort Order Works,” <https://www.expedia.com/lp/b/sort-order-info?langid=1033>, Accessed September 8, 2026.
- Expedia (2026b) “Terms of Service,” <https://www.expedia.com/lp/b/terms-of-service?langid=1033>, Accessed September 8, 2026.
- Expedia Group (2013) “Personalize Expedia Hotel Searches—ICDM 2013,” Kaggle competition and dataset, <https://www.kaggle.com/c/expedia-personalized-sort>.
- Expedia Group, Inc. (2022) “2021 Annual Report,” Form 10-K filed with the U.S. Securities and Exchange Commission, <https://www.sec.gov/Archives/edgar/data/1324424/000132442422000009/expe-20211231.htm>.
- Fong, Jessica, Olivia R. Natan, and Ranmit Pantle (2026) “Consumer Inferences from Product Rankings: The Role of Beliefs in Search Behavior,” *Management Science*, [10.1287/mnsc.2025.00761](https://doi.org/10.1287/mnsc.2025.00761), Articles in Advance.
- Gabaix, Xavier, David Laibson, Guillermo Moloché, and Stephen Weinberg (2006) “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96 (4), 1043–1068, [10.1257/aer.96.4.1043](https://doi.org/10.1257/aer.96.4.1043).
- Gergatsouli, Evangelia and Christos Tzamos (2023) “Weitzman’s Rule for Pandora’s Box with Correlations,” in *Advances in Neural Information Processing Systems*, 36, https://proceedings.neurips.cc/paper_files/paper/2023/hash/29d319f7c1513c9ecd81d3a6e9632a6e-Abstract-Conference.html.
- Greminger, Rafael P. (2026) “Trade-Offs Between Ranking Objectives: Descriptive Evidence and

- Structural Estimation,” *Management Science*, [10.1287/mnsc.2025.00782](https://doi.org/10.1287/mnsc.2025.00782), Articles in Advance.
- Hagiu, Andrei and Bruno Jullien (2011) “Why Do Intermediaries Divert Search?” *RAND Journal of Economics*, 42 (2), 337–362, [10.1111/j.1756-2171.2011.00136.x](https://doi.org/10.1111/j.1756-2171.2011.00136.x).
- Heidhues, Paul, Mats Köster, and Botond Köszegi (2023) “Steering Fallible Consumers,” *Economic Journal*, 133 (652), 1430–1465, [10.1093/ej/ueac093](https://doi.org/10.1093/ej/ueac093).
- Hodgson, Charles and Gregory Lewis (2025) “You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search,” *Econometrica*, 93 (4), 1299–1332, [10.3982/ECTA19576](https://doi.org/10.3982/ECTA19576).
- Janssen, Maarten C. W., Thomas Jungbauer, Marcel Preuss, and Cole Williams (2026a) “Search Platforms: Big Data and Sponsored Positions,” *Economic Journal*, [10.1093/ej/ueag002](https://doi.org/10.1093/ej/ueag002), Advance Access.
- Janssen, Maarten C. W., Thomas Jungbauer, Marcel Preuss, and Cole Williams (2026b) “Optimally Informative Rankings and Consumer Search,” *International Journal of Industrial Organization*, 108, 103282, [10.1016/j.ijindorg.2026.103282](https://doi.org/10.1016/j.ijindorg.2026.103282).
- Kaye, Aaron P. (2026) “The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets,” https://apkaye.github.io/aaronpkaye.website/Kaye_Aaron_JMP.pdf, Working paper, July 18.
- Kim, Jun B., Paulo Albuquerque, and Bart J. Bronnenberg (2010) “Online Demand Under Limited Consumer Search,” *Marketing Science*, 29 (6), 1001–1023.
- Luce, R. Duncan (1959) *Individual Choice Behavior: A Theoretical Analysis*, New York: Wiley.
- McClure, Jonathon (2025) “Markups and Costs under Capacity Constraints: The Welfare Effects of Hotel Mergers,” https://users.ssc.wisc.edu/~jcmccclure2/jmp_jcmccclure_nov23.pdf, Working paper, Purdue University, December.
- Netherlands Authority for Consumers and Markets (2020) “Booking and Expedia Inform Consumers More Clearly about Their Offers Following Action of European Consumer Authorities,” [ACM news release](#), Published 18 December 2020.
- Nguyen, Phong, John Dines, and Jan Krasnodebski (2017) “A Multi-Objective Learning to Re-Rank Approach to Optimize Online Marketplaces for Multiple Stakeholders,” arXiv:1708.00651v2, [10.48550/arXiv.1708.00651](https://arxiv.org/abs/10.48550/arXiv.1708.00651).
- Nocke, Volker and Patrick Rey (2024) “Consumer Search, Steering, and Choice Overload,” *Journal of Political Economy*, 132 (5), 1684–1739, [10.1086/728108](https://doi.org/10.1086/728108).
- Plackett, R. L. (1975) “The Analysis of Permutations,” *Journal of the Royal Statistical Society, Series C*, 24 (2), 193–202, [10.2307/2346567](https://doi.org/10.2307/2346567).
- Salz, Tobias (2022) “Intermediation and Competition in Search Markets: An Empirical Case Study,” *Journal of Political Economy*, 130 (2), 310–345, [10.1086/717349](https://doi.org/10.1086/717349).
- Ursu, Raluca M. (2018) “The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions,” *Marketing Science*, 37 (4), 530–552, [10.1287/mksc.2017.1072](https://doi.org/10.1287/mksc.2017.1072).
- Ursu, Raluca Mihaela, Tülin Erdem, Qingliang Wang, and Qianyun Zhang (2024) “Prior Information and Consumer Search: Evidence from Eye Tracking,” *Management Science*, 70 (12), 8685–8708, [10.1287/mnsc.2021.00611](https://doi.org/10.1287/mnsc.2021.00611).
- U.S. Bureau of Labor Statistics (2026) “Measuring Price Change in the CPI: Hotels and Motels,” <https://www.bls.gov/cpi/factsheets/hotels-and-motels.htm>, Consumer Price Index factsheet, accessed September 8, 2026.

Weitzman, Martin L. (1979) "Optimal Search for the Best Alternative," *Econometrica*, 47 (3), 641–654, [10.2307/1910412](#).

Appendices

Appendix A Rank Decoding: Derivation and a Two-Hotel Illustration

This appendix derives the match forecasts in (3.8)–(3.9) for a general organic list, then specializes them to the two-hotel illustration. The derivation conditions matches on the order at each candidate weight and keeps the session-entry weight belief fixed, as in Assumption 1.

A.1 Decoding a General List

Information and order likelihood. Fix a consumer, trip, and session and suppress their indices. Consider J organic hotels on a default or filtered list, with observed order $\mathbf{r} = (j_1, \dots, j_J)$. The displayed set and visible information are W , and $D = (\mathbf{r}, W)$ represents the displayed list. Throughout, also condition on the consumer's visible utilities v_j , which include price. Her subjective match representation is

$$U_j = v_j + m_j, \quad m_j = m_j^{P,C} + \eta_j, \quad m_j^{P,C} \sim N(0, \sigma_{m^{P,C}}^2), \quad \eta_j \sim N(0, 1). \quad (\text{A.1})$$

The predictable components, innovations, and ranking disturbances are independent before conditioning on the order. Let $g(t) = \varphi(t/\sigma_{m^{P,C}})/\sigma_{m^{P,C}}$, where φ and Φ are the standard normal density and CDF. The derivation supplies forecasts for uninspected hotels. An already inspected hotel keeps its revealed match rather than receiving a new forecast.

For candidate weight α and predictable matches $\mathbf{t} = (t_1, \dots, t_J)$, set $s_j^C(\alpha, t_j) = (1 - \alpha)(v_j + t_j) + \alpha\gamma p_j$. Independent unit-scale type-I extreme-value disturbances give

$$P_\alpha(\mathbf{t}) \equiv \Pr(\mathbf{r} \mid \mathbf{t}, W, \alpha) = \prod_{k=1}^J \frac{\exp(s_{j_k}^C(\alpha, t_{j_k}))}{\sum_{l \in \mathcal{R}_k} \exp(s_l^C(\alpha, t_l))}, \quad \mathcal{R}_k = \{j_k, \dots, j_J\}, \quad (\text{A.2})$$

where $\mathcal{R}_k$ is the set of remaining hotels from unallocated position k to J . The final factor is one.

Posterior predictions and marginalization. Averaging the order probability under the independent prior gives the normalizing constant

$$Z_\alpha(D) = \int_{\mathbb{R}^J} P_\alpha(\mathbf{q}) \prod_{k=1}^J g(q_k) d\mathbf{q}.$$

For a finite list, finite visible utilities and prices, and $\sigma_{m^{P,C}} > 0$, the integrand is positive and the likelihood is bounded by one. Thus $0 < Z_\alpha(D) \leq 1$ for every $\alpha \in [0, 1]$. Bayes' rule yields the normalized version of (3.8):

$$\pi(\mathbf{t} \mid D, \alpha) = \frac{P_\alpha(\mathbf{t}) \prod_{k=1}^J g(t_k)}{Z_\alpha(D)}. \quad (\text{A.3})$$

For each hotel, integrate out the other predictable components:

$$\pi_j(t \mid D, \alpha) = \int_{\mathbb{R}^{J-1}} \pi(t, \mathbf{t}_{-j} \mid D, \alpha) d\mathbf{t}_{-j}, \quad (\text{A.4})$$

where the scalar t occupies coordinate j . This integration averages over all rival predictions compatible with the list. Because the likelihood compares scores, adding the same constant to all predictable components leaves it unchanged. The prior therefore supplies information about match levels that ordinal comparisons alone cannot determine.

The match revealed by inspection. Conditional on $m_j^{P,C} = t$, the match is $m_j = t + \eta_j$. Convolution with the independent innovation's standard normal law gives

$$\begin{aligned} f_j^C(m \mid D, \alpha) &= \int_{\mathbb{R}} \varphi(m - t) \pi_j(t \mid D, \alpha) dt, \\ F_j^C(m \mid D, \alpha) &= \int_{\mathbb{R}} \Phi(m - t) \pi_j(t \mid D, \alpha) dt. \end{aligned} \quad (\text{A.5})$$

For example, $\Phi(m - t)$ is the probability of discovering a match no greater than m when the predictable component is t . Its integral averages this probability over the posterior predictions. These are continuous match distributions, although they need not be Gaussian. Their moments satisfy

$$\begin{aligned} \mathbb{E}^C[m_j \mid D, \alpha] &= \mathbb{E}^C[m_j^{P,C} \mid D, \alpha], \\ \text{Var}^C(m_j \mid D, \alpha) &= \text{Var}^C(m_j^{P,C} \mid D, \alpha) + 1. \end{aligned} \quad (\text{A.6})$$

At $\alpha = 1$, the likelihood is independent of $\mathbf{t}$, so the predictable-component posterior is its prior $N(0, \sigma_{m^{P,C}}^2)$ and the full-match forecast is $N(0, \sigma_{m^{P,C}}^2 + 1)$. The same unconditional

laws apply when the model assigns no rank signal to a sorted or sponsored position.

Mixing over the weight belief. The session-entry distribution b_s averages the candidate-specific match forecasts:

$$F_j^C(m) = \int_0^1 F_j^C(m \mid D, \alpha) b_s(d\alpha). \quad (\text{A.7})$$

This is (3.9). The mean of the mixture averages the candidate means, while its variance includes both the average candidate variance and dispersion of those means. Adding η_j contributes one to the mixture variance as well. Since visible utility is known,

$$F_j^U(u) = F_j^C(u - v_j). \quad (\text{A.8})$$

Appendix D.4 describes the numerical Laplace approximation and evaluation at the belief mean.

Independence restrictions. The joint posterior (A.11) generally correlates the predictable matches. Marginalizing and adding independent innovations do not remove that dependence. The consumer's independent evaluation of the entry laws (A.16) is therefore imposed separately in Assumption 3(a). It prevents a revealed match from changing other hotels' current forecasts, which substantially simplifies the analysis.

A.2 The Two-Hotel Illustration

Given the consumer's prior and subjective score model, the ordinal ranking changes the probability attached to each possible match. This appendix derives that conditional distribution for two hotels and illustrates how perceived monetization can change the inspection while the displayed order remains fixed.

Information and the probability of an order. Fix one consumer, trip, and session, and suppress their indices *its*. Both hotels $j = 1, 2$ are organic and uninspected, with hotel 2 displayed above hotel 1. The consumer knows the visible utilities v_j , which already include price p_j . Write $D = (\mathbf{r}, W)$, where $\mathbf{r} = (2, 1)$ denotes the observed order and W collects the visible information. Following Section 3.3, her subjective match model is

$$U_j = v_j + m_j, \quad m_j = m_j^{P,C} + \eta_j, \quad m_j^{P,C} \sim N(0, \sigma_{m^{P,C}}^2), \quad \eta_j \sim N(0, 1). \quad (\text{A.9})$$

The predictable components $m_j^{P,C}$, innovations η_j , and ranking disturbances ε_j are independent before observing rank. Inspection reveals m_j . Let $g(t) = \phi(t/\sigma_{m^{P,C}})/\sigma_{m^{P,C}}$ be the predictable component's prior density, with ϕ and Φ the standard normal density and CDF.

First hold a candidate monetization weight α fixed. For possible predictable matches (t_1, t_2) , denote the deterministic subjective scores as $s_j^C = s_j^C(\alpha, t_j)$. The displayed order compares these scores after adding independent unit-scale type-I extreme-value disturbances ε_j . Observing hotel 2 above hotel 1 therefore conditions on a positive noisy score difference. Since $\varepsilon_1 - \varepsilon_2$ has the standard logistic CDF $\Lambda(x) = 1/(1 + e^{-x})$,

$$\begin{aligned} s_j^C &= (1 - \alpha)(v_j + t_j) + \alpha\gamma p_j, \\ \mathbf{r} = (2, 1) &\iff s_2^C + \varepsilon_2 > s_1^C + \varepsilon_1, \\ P_\alpha(t_1, t_2) &\equiv \Pr(\varepsilon_1 - \varepsilon_2 < s_2^C - s_1^C \mid t_1, t_2, W, \alpha) \\ &= \Lambda(s_2^C - s_1^C) = \frac{\exp(s_2^C)}{\exp(s_1^C) + \exp(s_2^C)}. \end{aligned} \tag{A.10}$$

The last line is the two-hotel Plackett–Luce probability. For $\alpha < 1$, this probability increases with the predictable-match difference $t_2 - t_1$, holding visible utility and prices fixed. In other words, a larger t_2 relative to t_1 makes the observed order $\mathbf{r} = (2, 1)$ more likely. The consumer observes the inequality, while both match predictions and the numerical score difference remain unknown. Bayes’ rule in (A.11) weights each possible pair by this order probability. Ranking noise leaves every finite pair possible, so conditioning shifts a continuous distribution over match predictions. Visible utility and commission also affect this probability, so the same position can convey different match information at different prices or weights.

From the order to a distribution over predictions. Bayes’ rule reweights the independent prior by the probability of the observed order:

$$\pi(t_1, t_2 \mid D, \alpha) = \frac{g(t_1)g(t_2)P_\alpha(t_1, t_2)}{\iint_{\mathbb{R}^2} g(q_1)g(q_2)P_\alpha(q_1, q_2) \, dq_1 \, dq_2}, \tag{A.11}$$

where the denominator averages the order probability over all prior predictions and normalizes the posterior to integrate to one.

For either hotel’s marginal distribution, integrate out the other hotel’s unknown prediction:

$$\pi_1(t \mid D, \alpha) = \int_{\mathbb{R}} \pi(t, z \mid D, \alpha) \, dz, \quad \pi_2(t \mid D, \alpha) = \int_{\mathbb{R}} \pi(z, t \mid D, \alpha) \, dz. \tag{A.12}$$

From predictions to the match revealed by inspection. Conditional on $m_j^{P,C} = t$, the match is $m_j = t + \eta_j$. Rank leaves the independent innovation unchanged. Averaging its

conditional density and CDF over the posterior predictions therefore gives

$$\begin{aligned} f_j^C(m \mid D, \alpha) &= \int_{\mathbb{R}} \phi(m - t) \pi_j(t \mid D, \alpha) dt, \\ F_j^C(m \mid D, \alpha) &= \int_{\mathbb{R}} \Phi(m - t) \pi_j(t \mid D, \alpha) dt. \end{aligned} \quad (\text{A.13})$$

For example, $\Phi(m - t)$ is the probability that an inspection reveals a match no greater than m when the predictable component equals t . The second integral averages that probability over all predictions still possible after observing the order. This is the continuous conditional match distribution used for inspection decisions.

The innovation η_j preserves the prediction mean and adds its unit variance:

$$\begin{aligned} \mathbb{E}^C[m_j \mid D, \alpha] &= \mathbb{E}^C[m_j^{P,C} \mid D, \alpha], \\ \text{Var}^C(m_j \mid D, \alpha) &= \text{Var}^C(m_j^{P,C} \mid D, \alpha) + 1. \end{aligned} \quad (\text{A.14})$$

At $\alpha = 1$, the ranking does not reflect match and the likelihood no longer depends on predictions, so (A.11) restores their prior $N(0, \sigma_{m^{P,C}}^2)$ and the match prior $N(0, \sigma_{m^{P,C}}^2 + 1)$.

With a session-entry belief b_s over monetization, average the candidate-specific CDFs:

$$F_j^C(m) = \int_0^1 F_j^C(m \mid D, \alpha) b_s(d\alpha). \quad (\text{A.15})$$

This recovers (3.9). As required by Assumption 1, the ranking conditions match beliefs at each candidate without itself updating b_s . Since visible utility is known, the corresponding utility distribution is simply

$$F_j^U(u) = F_j^C(u - v_j). \quad (\text{A.16})$$

Numerical illustration and inspection order. Figure 10 holds the displayed order, prices, and visible utilities fixed while changing the candidate weight. Hotel 2 is more expensive and has lower visible utility. Under low perceived monetization, its first position is strong evidence of a favorable match. Under high perceived monetization, commission helps explain that position and the match loading is small, leaving both posteriors close to their priors. These comparisons use illustrative parameters.

At the low weight $\alpha = 0.05$, observing hotel 2 above hotel 1 raises hotel 2's expected match from the prior mean of zero to 1.302 and lowers hotel 1's expected match to -1.302 . At the high weight $\alpha = 0.95$, the corresponding posterior means are 0.096 and -0.096 , both close to zero. Relative to the common prior match standard deviation $\sqrt{5}$, each hotel's mean shifts by approximately 0.58 standard deviations under the low weight and 0.043

under the high weight, with opposite directions across hotels. Before conditioning on rank, the probability of hotel 2 appearing above hotel 1 is 0.329 under the low weight and 0.517 under the high weight. Moving from the predictable-match distributions in the upper panels to the full-match distributions in the lower panels preserves each posterior mean and increases each variance by one through the independent innovation.




Figure 10. The same ranking under two perceived monetization weights

Notes: Hotel 2 is displayed above hotel 1 in both comparisons. Prices are $(p_1, p_2) = (1, 2)$ and visible utilities, including price, are $(v_1, v_2) = (2, 0.5)$. The price coefficient is $\beta^p = -1$, the commission rate is $\gamma = 0.15$. Ranking noise has unit scale. Predictable matches have independent $N(0, 4)$ priors, independent innovations are $N(0, 1)$, giving $N(0, 5)$ match priors. Gray dashed curves show the priors. Green solid and orange dash-dotted curves condition on the same order at $\alpha = 0.05$ and $\alpha = 0.95$. Colored vertical lines mark posterior means. The numerical integration evaluates (A.11)–(A.13) directly.

These distributional changes can reverse the inspection order. Under the paper's assumptions, equal inspection costs give reservation values $(\zeta_1, \zeta_2) = (2.64, 3.74)$ at the low weight and $(4.18, 2.89)$ at the high weight. The consumer therefore inspects hotel 2 first under the former and hotel 1 first under the latter, despite seeing the same ranking.³³

³³This calculation uses $c(\ell_1) = c(\ell_2) = 0.18$, outside utility $u_0 = 0$, and a terminal session with boundary-shock scale $\tau_B = 1$. Thus $\Psi(u) = \log(1 + e^u)$ and each reservation value solves (3.17) using the corresponding candidate-specific utility distribution. Both indices exceed the no-inspection stopping value $\Psi(-\infty) = 0$.

Appendix B The Consumer's Problem

This appendix establishes Proposition 1, the optimal inspection policy conditional on session-entry forecasts. Section B.1 sets up the history and the decision state. Sections B.2–B.4 establish existence and the transformed search problem. Section B.5 proves the reservation rule. Fix a consumer–trip pair and suppress i, t . All subjective expectations condition on the consumer's known visible utilities, including private pre-search tastes.

B.1 States, Actions, and Transitions

History and entry state. The history $\mathcal{H}_{s-1}$ records earlier lists and actions, inspected identities $\mathcal{K}_{s-1}$, their revealed matches and inspected offers, and the evidence already incorporated in the weight belief. Each inspected offer keeps its price, attributes, and utility. At entry, $z_s = (\bar{u}_s^0, \omega_s)$, where $\bar{u}_s^0$ is the highest inspected utility, or $-\infty$ if $\mathcal{K}_{s-1}$ is empty. The fixed collection ω_s contains the session identity, current list and order, belief b_s , predictive marginals, list and continuation value covariates, feasible continuation actions, and their corresponding forecast functions. The inspectable set is $\text{NI}_s^0 = \mathcal{J}_s \setminus \mathcal{K}_{s-1}$, also including uninspected sponsored hotels.

The organic order conditions match forecasts at the unchanged b_s under Assumption 1. For inspection, Assumption 3(a) imposes independent evaluation of the resulting mixed utility marginals F_{js}^U . These behavioral restrictions aim to simplify the analysis and the computation of the likelihood function. They are not required for the general inspection theorem, which only needs a nondecreasing scalar stop map in Assumption 3(b).

Inspection and the boundary. A within-session history h records the newly inspected identities and revealed matches. At fixed ω_s , inspection decisions use state $(\bar{u}, \text{NI}; \omega_s)$, with $\bar{u}$ the best found utility and $\text{NI} \subseteq \text{NI}_s^0$ the remaining uninspected hotels. Feasible inspection actions are $\{\text{stop}\} \cup \text{NI}$. Inspecting j pays $c(\ell_j) > 0$, reveals U_j , removes j from NI , adds the offer to the inspected history, and replaces $\bar{u}$ by $\max\{\bar{u}, U_j\}$. The consumer can stop before a first inspection or after any subsequent one.

Only after stopping does the consumer observe the boundary shocks. At $s < \bar{S}$, the feasible actions are $\{\text{book, quit, def, filter, sort}\}$ with booking removed if no offer has been inspected. At $s = \bar{S}$, continuation is also removed. Each feasible action adds an independent mean-zero EVT1 shock with standard deviation $\sigma_B > 0$ (i.e., with scale $\tau_B = \sqrt{6}\sigma_B/\pi$), independent of inspection rewards and earlier boundary shocks. Its deterministic value is $\bar{u}$ for booking, u_0 for leaving, and $V_s^a(\bar{u}; \omega_s)$ from (3.14) for continuation. The same booking shock applies whichever inspected hotel is selected. Fixed enumerations of hotels and

boundary actions resolve ties and inspection stops at equality in (3.18). Maxima over empty sets are $-\infty$, and empty products are one.

Actual and forecast transitions. The entry belief and predictive marginals remain fixed through the boundary. The continuation functions are also fixed though their values can respond to $\bar{u}$. In particular, the forecast law $\hat{\Gamma}_a(d\omega' \mid \omega_s)$ is not revised using current inspections, newly inspected identities, or an updated belief before that choice. It is distinct from the actual transition law Γ_a .

After a continuation decision, the update (3.10) uses only the set of eligible updating signals $\mathcal{C}_s^{\text{new}}$. A inspected hotel offers an eligible updating signal if it is and organic hotels on default or filtered lists and we only consider the first inspection of each such hotel. Sorted, sponsored, or previously inspected hotel do not supply new evidence for the weight updating. An empty eligible set leaves b_s unchanged, for instance, after zero inspection. The completed history includes all inspected offers regardless whether eligible for learning or not. The actual next list may depend on this history and platform-observed actions. At the next entry, every inspected identity is excluded from NI_{s+1}^0 , and forecasts are rebuilt for uninspected hotels from the new list and b_{s+1} . Actual hotel–trip matches persist. Booking or leaving ends the trip. These transitions require the full history even though the conditional inspection decision uses a smaller summary.

Conditional recursion. Let $B_s(h; \omega_s)$ be the expected payoff of arriving at the boundary after history h , including the boundary maximization and the associated future-payoff forecasts. Scalar sufficiency (of Assumption 3(b)) requires

$$B_s(h; \omega_s) = \Psi_s(\bar{u}(h); \omega_s).$$

The inspection recursion is

$$V_s^I(\bar{u}, \text{NI}; \omega_s) = \max \left\{ \Psi_s(\bar{u}; \omega_s), \max_{j \in \text{NI}} \left[-c(\ell_j) + \int V_s^I(\max\{\bar{u}, u\}, \text{NI} \setminus \{j\}; \omega_s) dF_{js}^U(u) \right] \right\}, \quad (\text{B.1})$$

with $V_s^E(\bar{u}_s^0, \omega_s) = V_s^I(\bar{u}_s^0, \text{NI}_s^0; \omega_s)$. The independent entry laws and continuation functions are inputs to this recursion.

B.2 Regularity and Existence

We impose the following regularity conditions.

Assumption 5 (Regularity).

1. The session horizon and each displayed list are finite. Hotel identifiers admit a fixed enumeration. States, payoff maps, and probability kernels are Borel measurable. Inspection costs are finite and strictly positive. Inspected offers remain available at their revealed utility.
2. The outside option and continuation intercepts are finite. Each assigned continuation forecast is finite, including when no offer has been inspected. Boundary shocks have the independence, scale, and location specified above. The scalar map Ψ_s is measurable, finite, and nondecreasing on $\{-\infty\} \cup \mathbb{R}$, and its transformed rewards are integrable at every fixed entry:

$$\int |\Psi_s(u; \omega_s)| dF_{js}^U(u) < \infty \quad \text{for every } j \in \text{NI}_s^0.$$

3. The weight belief is a probability measure on $[0, 1]$. The normalizing integrals in (3.8) and (3.10) are finite and strictly positive for the histories under consideration, so predictive laws and subsequent beliefs are well defined.

We impose the following integrability conditions on the transformed reward. For the action-level implementation, we assume the following sufficient conditions to hold.

Assumption 6 (Integrability). 1. Every $V_s^a(u; \omega_s)$ be finite and nondecreasing, with a finite no-offer value.
 2. $\mathbb{E}|U_j| < \infty$ for every hotel j .
 3. For every continuation action a and hotel j , $\int [V_s^a(u; \omega_s)]^+ dF_{js}^U(u) < \infty$.

Proposition 2 (Conditional existence). Under Assumptions 3, 5, and 6, the conditional inspection and entry values are finite and measurable, and a measurable optimal inspection policy exists at each entry. With the stated finite action forecasts, the boundary maximization has a measurable maximizing action and finite expected value.

Proof. Fix an entry and write $X_j = \Psi_s(U_j; \omega_s)$ and $\psi = \Psi_s(\bar{u}; \omega_s)$. Monotonicity of Ψ_s implies

$$\Psi_s(\max\{\bar{u}, U_j\}; \omega_s) = \max\{\psi, X_j\}.$$

The terminal transformed reward under any feasible inspection policy is bounded above by $\psi + \sum_{j \in \text{NI}} (X_j - \psi)^+$. Positive costs cannot increase this payoff, and stopping immediately gives ψ . Integrability therefore yields the finite bounds

$$\psi \leq V_s^l(\bar{u}, \text{NI}; \omega_s) \leq \psi + \sum_{j \in \text{NI}} \mathbb{E}[(X_j - \psi)^+] < \infty.$$

We start induction at $\text{NI} = \emptyset$ where the value is ψ . Then induct over the finite remaining set. After inspecting j , the new transformed best is $\max\{\psi, X_j\}$. Its expectation is finite,

and independence and the displayed upper bound control the expected gains from the other hotels (after replacing ψ by $\max\{\psi, X_j\}$). Thus each inspect option in (B.1) has a finite integral. Borel kernels preserve measurability, and a fixed tie enumeration selects a measurable maximizer from the finite action set. This proves the existence for inspection and entry.

At the boundary, the finite maximum of finite deterministic values plus integrable EVT1 shocks is integrable and has a measurable maximizing action. Unavailable actions are omitted. ■

B.3 The Scalar Stop Map

Lemma 1 (Scalar stop map). *In the action-level implementation, finite nondecreasing continuation functions and the stated mean-zero EVT1 shocks give (3.16) exactly. The map is strictly increasing at finite $\bar{u}$ and has a finite no-offer value. The general inspection theorem requires only the nondecreasing scalar map in Assumption 3(b).*

Proof. For deterministic values q_a and independent EVT1 shocks with common location μ_B and scale τ_B , the expected maximum over the feasible actions $\mathcal{A}$ is

$$\mathbb{E} \max_{a \in \mathcal{A}} \{q_a + \varepsilon_a^B\} = \tau_B \log \sum_{a \in \mathcal{A}} e^{q_a/\tau_B} + \mu_B + \gamma_E \tau_B.$$

Mean zero sets $\mu_B = -\gamma_E \tau_B$, cancelling the last two terms. Substituting the booking, outside, and continuation values gives (3.16). Nondecreasing continuation functions and the strictly increasing booking term make the log-sum strictly increasing for finite $\bar{u}$. With no inspected offer, the booking term is omitted, while the finite continuation values and the always-feasible outside option give a finite value. At the final session, only booking and leaving enter. Monotonicity of an integral continuation forecast follows from integrating its nondecreasing value function against the fixed entry law under the conditions in Section B.2. ■

Scalar sufficiency is weaker than fixing each action's function. With EVT1 shocks and a given $\bar{u}$, the expected boundary payoff depends on continuation values through their exponentiated sum. For example, $C_1(h) = K + \tau_B \log p(h)$ and $C_2(h) = K + \tau_B \log[1 - p(h)]$, with $0 < p(h) < 1$, have a constant exponentiated sum while their choice shares vary with history. Assumption 2 instead holds each assigned function fixed through the boundary.

B.4 Transforming the Within-Session Problem

Proposition 3 (Continuation-adjusted Pandora problem). *Under Assumptions 3 and 5, fix s, ω_s and write Ψ for the stop map. The inspection value depends on $\bar{u}$ only through $\psi = \Psi(\bar{u})$: $V_s^I(\bar{u}, \text{NI}; \omega_s) = \mathcal{V}(\text{NI}, \psi)$, where*

$$\mathcal{V}(\text{NI}, \psi) = \max \left\{ \psi, \max_{j \in \text{NI}} \left[-c(\ell_j) + \int \mathcal{V}(\text{NI} \setminus \{j\}, \max\{\psi, x\}) dF_{js}^\Psi(x) \right] \right\}, \quad (\text{B.2})$$

where F_{js}^Ψ is the law of $\Psi(U_j)$. Each hotel has a unique finite reservation value solving (3.17).

Proof. For an empty remaining set, $\mathcal{V}(\emptyset, \psi) = \psi$. Suppose the representation holds for all smaller sets. Substitution into (B.1) and the identity

$$\Psi(\max\{\bar{u}, u\}) = \max\{\Psi(\bar{u}), \Psi(u)\}$$

show that the recursion depends on $\bar{u}$ only through ψ . Pushing each law F_{js}^U through $x = \Psi(u)$ gives (B.2). The deterministic transformation preserves the independence imposed for inspection. This is the independent-box search problem of Weitzman (1979), with rewards $\Psi(U_j)$, positive opening costs, perfect recall, and no discounting.

For the reservation solution, set $X_j = \Psi(U_j)$ and $H_j(\zeta) = \mathbb{E}[(X_j - \zeta)^+]$. Integrability of X_j makes H_j finite, continuous, and nonincreasing. For $\zeta_1 < \zeta_2$ with $\Pr(X_j > \zeta_2) > 0$,

$$H_j(\zeta_1) - H_j(\zeta_2) \geq (\zeta_2 - \zeta_1) \Pr(X_j > \zeta_2) > 0.$$

Dominated convergence gives $H_j(\zeta) \rightarrow 0$ as $\zeta \rightarrow \infty$, while Jensen's inequality gives $H_j(\zeta) \geq (\mathbb{E}X_j - \zeta)^+ \rightarrow \infty$ as $\zeta \rightarrow -\infty$. Every finite positive cost therefore has a finite solution. At any solution, $H_j(\zeta) = c(\ell_j) > 0$, so $\Pr(X_j > \zeta) > 0$, and the strict-decrease argument gives uniqueness. The solution is measurable because $H_j(q)$ is measurable at each rational q and continuous and monotone in its argument. ■

B.5 Proof of Proposition 1

Proof. Proposition 2 establishes existence, and Proposition 3 gives the transformed Bellman equation. Fix an entry state and write

$$X_j = \Psi_s(U_j; \omega_s), \quad \psi = \Psi_s(\bar{u}; \omega_s), \quad Y_j = \min\{X_j, \zeta_{js}\}, \quad \Delta_j = (X_j - \zeta_{js})^+.$$

The capped reward is Y_j , and the reservation equation gives $\mathbb{E}\Delta_j = c(\ell_j)$. A feasible policy decides whether to inspect j before seeing its reward. Under inspection independence,

that event is independent of X_j , so

$$\mathbb{E}[\mathbf{1}\{j \text{ inspected}\} \Delta_j] = \Pr(j \text{ inspected}) \mathbb{E} \Delta_j = \Pr(j \text{ inspected}) c(\ell_j).$$

Thus, if $\mathcal{O}$ is the set it inspects, the policy's expected net payoff equals the expectation of the first expression below:

$$\max\{\psi, \max_{j \in \mathcal{O}} X_j\} - \sum_{j \in \mathcal{O}} \Delta_j \leq \max\{\psi, \max_{j \in \mathcal{O}} Y_j\} \leq \max\{\psi, \max_{j \in \text{NI}} Y_j\}.$$

The first inequality follows by subtracting the winning reward's excess, when positive, and any other nonnegative excesses. The last expression therefore bounds every policy's expected payoff.

The reservation policy attains this bound. It inspects the largest remaining index only while that index exceeds the current transformed best. Before its last inspection, no inspected reward can have $X_j > \zeta_{js}$: such a draw would exceed every remaining index and trigger stopping. If the last inspection has positive excess, it supplies the terminal maximum. Its index exceeds the previously inspected best and weakly exceeds every remaining index, so subtracting that sole excess yields the maximum capped reward. If no inspection has positive excess, each inspected reward equals its capped reward Y_j , and the stopping condition bounds every uninspected cap by the inspected maximum. With zero inspections, the same equality holds because all indices are at most ψ . Consequently,

$$\mathcal{V}(\text{NI}, \psi) = \mathbb{E} \left[\max\{\psi, \max_{j \in \text{NI}} \min\{X_j, \zeta_{js}\}\} \right].$$

The costs are already incorporated through the reservation caps in this identity.

Returning to $\psi = \Psi_s(\bar{u}; \omega_s)$ proves the inspection and stopping rules, including optimal stopping at equality. After stopping, maximizing the realized boundary values gives the boundary rule. Conditional on booking, its common action shock leaves the highest-utility inspected offer optimal, and the fixed enumeration resolves hotel ties. Recall preserves this choice across lists. The argument applies separately at each entry, conditional on its assigned forecasts. ■

Index comparisons. At a fixed stop map and cost, first-order stochastic dominance of X'_j over X_j gives $\mathbb{E}[(X'_j - \zeta)^+] \geq \mathbb{E}[(X_j - \zeta)^+]$ at every ζ and hence a weakly higher reservation index. A mean-preserving spread of the transformed reward has the same implication because the positive-part function is convex. Lowering a positive cost also raises the index. A belief change that also changes the stop map requires a separate comparison.

Recall, normalization, and estimation. A hotel first inspected as organic on a default or filtered list supplies eligible weight evidence once, if continuation is chosen. Reappearing sorted or sponsored does not create another inspection or another use of that evidence. A hotel first inspected sorted or sponsored is also recorded, but its known match is never paired with a later organic rank as fresh updating evidence. Uninspected hotels receive new forecasts at the next entry while their actual matches persist.

For fixed action forecasts, a translation $\Psi \mapsto \Psi + d$ shifts every index by d and leaves inspection comparisons unchanged. Changing the primitive shock location throughout a trip with an endogenous number of boundaries is a different operation. The model uses mean-zero shocks at each boundary, with no extra Euler term or adjustment to c_{sess} . Finally, the proof conditions on consumer-known private utility. The researcher's joint integration over persistent tastes, matches, and compatible histories is the separate likelihood calculation in Appendix D.2. The capped-reward identity supplies the net-of-inspection-cost forecast in Appendix D.3. The one-list horizon is an additional numerical approximation to simplify the estimation instead of a premise of the inspection theorem.

Appendix C Data Details

C.1 Data Summary

Release and observation rules. The logs were released by Expedia Group for research use. The users are a random sample of those who clicked at least once during June–July 2021, excluding those who booked more than four distinct properties. Eligibility applies over the release window, so a sampled user’s individual trip or session can contain no inspection. User identifiers are browser cookies. Property and destination identifiers are anonymized. Clicks and bookings are attributed to a search only if they occur after it, before the next search, and within 180 minutes.

Estimation sample. We form sessions and trips as defined in Section 2.2 and select the most searched destination. We keep trips whose first list is the platform’s default list, so that first-session choices inform the consumer’s trip-entry belief before any refinement. The estimation sample uses a maximum of five sessions per trip and five inspections per session, which is where the data become thin. Table 6 reports the resulting 8,792 trips, 13,016 sessions, and 944,127 hotel-card appearances.

C.2 The Interface

Figure 11 illustrates the two pages used in the model. The results page (left) displays the list: each card presents the star rating, the guest review score and count, the free-cancellation and discount badges, and the nightly price. Advertisements are labeled and interleaved among the unpaid listings. The filters sit in the left column and the sort menu above the list. The hotel page (right), reached by an inspection, provides more detailed information, including amenities, photographs, and room choices, and is the only place a booking can be made.

C.3 Quitting Evidence by Session Order


Figure 12 separates the quitting evidence in Section 2.3 by session order within the trip. Among sessions that end without a booking, it is more likely that this session also ends the entire trip query when the session contained an inspection than when it did not, regardless of the session order.³⁴

³⁴The fifth session is omitted because the sample admits at most five sessions, so every fifth session ends the trip by construction.


Table 6. Summary statistics for the estimation sample

	Mean	SD	Median
<i>Panel A. Trips (N = 8,792)</i>			
Sessions in the trip	1.48	0.86	1
Two or more sessions (%)	30.7		
Nights	2.96	1.84	3
Days between search and check-in	49.9	58.5	30
Adults	2.06	0.98	2
Children	0.12	0.52	0
Rooms	1.01	0.13	1
Two adults, no children (%)	72.8		
Hotels inspected in the trip	0.66	1.07	0
At least one hotel inspected (%)	39.2		
Ends in a booking (%)	4.3		
<i>Panel B. Sessions (N = 13,016)</i>			
Hotels displayed	72.5	35.3	100
Advertisements displayed	6.5	5.1	4
At least one advertisement (%)	76.5		
Hotels inspected in the session	0.45	0.87	0
At least one hotel inspected (%)	28.2		
List is the platform's default (%)	77.6		
List is filtered (%)	8.1		
List is sorted (%)	14.3		
Session ends with a booking (%)	2.9		
Session ends the trip without booking (%)	64.6		
Next session re-issues the default query (%)	10.1		
Next session filters (%)	8.1		
Next session sorts (%)	14.3		
Mobile device (%)	49.9		
<i>Panel C. Hotel cards (N = 944,127)</i>			
Star rating	3.95	0.82	4
Guest review score (out of 5)	3.94	0.99	4
Number of reviews	5,753	6,392	3,050
Free cancellation (%)	79.7		
Discount badge (%)	79.2		
Price quintile within the list	2.99	1.42	3
Advertisement (%)	9.0		

Notes: Statistics describe the estimation sample constructed from the research release. A session is one displayed list. A trip links sessions for the same query requirements as defined in Section 2.2. Hotel cards are appearances on lists instead of distinct properties because the same hotel can appear on multiple lists. The five session-ending rows distinguish booking, leaving without booking, and the three continuation actions. Percentages use the population in the corresponding panel.



(a) the results page



(b) a hotel page

Figure 11. The results page and a hotel page

Notes: Screenshots of the platform's desktop site. Panel (a) corresponds to the results/list page filled by hotel cards visible to consumers before an inspection. Panel (b) illustrates the detailed hotel page information reached by a click.




Figure 12. Ending the trip after a non-booking session, by session order

Notes: Share of non-booking sessions after which the trip ends, by the order of the session within the trip and by whether the session contained an inspection, in the estimation sample (12,637 non-booking sessions; sessions of order five omitted). Bars are 95% confidence intervals with standard errors clustered by trip.

C.4 Price Construction

The empirical price input in (4.2) is a constructed dollar-price proxy,

$$p_{jts} = c(g_{ts}, q_{jts}) + \delta_j.$$

The observed quintile q_{jts} locates the hotel relative to the other available offers in its list. The stratum–quintile center $c(g, q)$ supplies a dollar level, and the hotel-specific offset δ_j uses additional relative-price information. Both components describe nightly rates.

The source of dollar prices. Expedia released, for a ranking competition on Kaggle ([Expedia Group, 2013](#)), the logs of searches made between November 2012 and June 2013, with the nightly price of every displayed hotel in dollars. We keep the searches for properties in the destination country with the largest amount of transactions, with prices between \$10 and \$2,000 per night, and searches that display at least twenty such hotels, so that quintiles are well defined. 2013 destinations are anonymized as well, so the match to the 2021 data is at the country-query level.

Strata and centers. Strata combine three query characteristics recorded in both datasets: the length of the stay (one to two, three to four, five or more nights), whether the stay includes a Saturday night, and the booking window (within a week, one to four weeks, more than four weeks ahead). We do not split stays of five or more nights by weekend.

For each search impression at each stratum, we split its displayed prices into five quintiles (the 10th, 30th, 50th, 70th, and 90th percentiles) and for each quintile we take the midpoint price. Then the center of quintile q at each stratum is the median across searches. Table 7 lists the fifteen strata with their 2013 searches, the five centers, and the number of 2021 sessions each stratum covers. The centers rise with the length of the stay and the booking window, from \$64 for the cheapest quintile of a short, last-minute stay to \$267 for the priciest quintile of a long stay booked more than four weeks ahead.

Inflation. The centers are in 2013 dollars. Inflating them by the consumer price index for hotels and motels ([U.S. Bureau of Labor Statistics, 2026](#)), 1.13 between June 2013 and June 2021 or 1.29 between the two windows, lowers the agreement with the price filters described below, so the paper uses the 2013 centers as its dollar scale and reads all dollar figures in 2013 prices.

Hotel-specific adjustments. The center gives every hotel in a stratum-quintile cell the same price. The adjustment δ_j of (4.2) places the hotel within the cell from three facts about its own price. First, its price rank on the lists consumers sorted by price, which locates it exactly among the hotels of its quintile on those lists.³⁵ Second, the price bands consumers filtered on: a hotel displayed in a price-filtered list has additional information about its price range.³⁶ Third, a hotel’s characteristics and average quintile across lists provides further information. A hotel with no information has $\delta_j = 0$.

³⁵912 properties appear on at least one price-sorted list, 352 on twenty or more.

³⁶Expedia uses five fixed filter bands: <\$75; \$75-\$125; \$125-\$200; \$200-\$300; >\$300 across years in the United States. 494 properties are seen under two or more bands.

Table 7. Price strata and quintile centers

Stay	Weekend	Booking window	2013 searches	Quintile centers (2013 dollars)					2021 sessions
				1	2	3	4	5	
1–2 nights	no	within a week	23,462	64	82	103	129	169	846
1–2 nights	no	1–4 weeks	14,224	69	90	116	139	185	775
1–2 nights	no	over 4 weeks	10,931	79	104	129	156	200	599
1–2 nights	yes	within a week	33,246	67	86	101	119	154	916
1–2 nights	yes	1–4 weeks	23,143	71	90	110	130	170	1,011
1–2 nights	yes	over 4 weeks	17,761	80	103	127	150	199	992
3–4 nights	no	within a week	4,307	63	84	106	131	179	267
3–4 nights	no	1–4 weeks	5,727	76	100	126	152	199	746
3–4 nights	no	over 4 weeks	7,740	93	124	151	180	241	1,445
3–4 nights	yes	within a week	3,536	68	89	108	129	169	301
3–4 nights	yes	1–4 weeks	5,819	79	102	127	151	199	1,011
3–4 nights	yes	over 4 weeks	8,618	95	125	153	187	247	2,386
5+ nights		within a week	2,064	64	85	105	129	173	195
5+ nights		1–4 weeks	3,657	81	106	132	159	215	373
5+ nights		over 4 weeks	9,001	103	134	162	199	267	1,153

Notes: 2013 searches are those in the companion data for United States properties with prices between \$10 and \$2,000 per night and at least twenty such hotels displayed. Centers are medians across searches of the within-search 10th, 30th, 50th, 70th, and 90th price percentiles, rounded to the dollar. Stays of five or more nights are not split by weekend. The last column counts the 2021 estimation sessions in each stratum, and the fifteen rows sum to 13,016 sessions.

Appendix D Estimation Details

This appendix describes the generated ranking function, shared integration of consumer outcomes, and numerical evaluation of continuation and decoding. Independent inspection evaluation, the product likelihood for learning, and entry-fixed forecasts are the behavioral restrictions in Section 3. Finite integration, a one-list forecast horizon, Laplace approximation, interpolation, and mean-weight evaluation are separate computational approximations.

D.1 The Ranking Stage

Features and targets. The ranker’s features include information from (i) the hotel card: star rating, review score and count, cancellation, discount, and price; and (ii) the query: stay length, booking window, number of adults, children, rooms, and check-in month, device, point of sale, consumer location country, and session order within the trip.

After fitting the demand we obtain $\hat{\theta}$ that we use to fit $\hat{m}^P$ with regression (4.4) using the platform-information basis, including hotel identity and prior inspection history. The target $T_{ijts}(\alpha, \hat{\theta})$ combines that prediction with expected visible utility and γp_{jts} .

LambdaMART is a gradient-boosted tree ensemble trained to order the labels within lists (Borges, 2010). The training criterion is Normalized Discounted Cumulative Gain (NDCG) that discounts gains at lower positions and divides by the discounted gain of an ideal ordering defined by the targets $T_{ijts}(\alpha, \hat{\theta})$. This training criterion is distinct from the observed-order likelihood and is commonly used in the computer science literature. We keep the regular random sample split rule: 80% training / 20% evaluation.

The order likelihood (4.10) takes the LambdaMART’s learned output g^R as its index and fixes the ranking-noise scale to $\sigma_R = 1$. Adding a common constant to a list’s indices leaves Plackett–Luce probabilities unchanged. Multiplying the indices by a positive constant preserves their order but changes those probabilities at unit noise. Consequently, good NDCG fit does not calibrate the cardinal differences that the likelihood uses.

Selection and reported fit. The candidate grid of α^P is $[0, 1]$. Each candidate completes its own match prediction, target construction, and ranker training before evaluation in (4.9). Figure 13 reports the ranking-likelihood profile, whose maximum is at 0.78. At the estimated weight, the reported ranker removes 58% of the misordered pairs of a random order and 81% of its NDCG loss. Table 6.4 of Kaye (2026) implies reductions of 45% and 66%, respectively, using the Kaggle 2013 dataset.




Figure 13. Fit of the ranking stage by candidate weight

D.2 The Trip Likelihood

Compatible orders and inspected histories. The observed session outcome records an inspected set, a boundary action, and a booked identity when applicable. A set of at most five new inspections has at most $5!! = 120$ permutations. The collection $\Pi(\mathcal{C}_s)$ contains orders compatible with the recorded identities and the completed entry history, and an impossible order receives zero weight through (4.11). The empty set has one empty order. Previously inspected identities are excluded from new inspection opportunities and remain available at their revealed matches, prices, and attributes.

For each persistent draw Λ and compatible history, the policy is evaluated chronologically:

1. start from the inspected set and best utility at entry;
2. evaluate the highest-index and stopping comparisons;
3. keep each revealed offer;
4. evaluate the boundary and identity factors at the resulting best utility.

Stopping at entry is allowed. With zero new inspections, booking can select an earlier inspected offer. At the final session continuation is infeasible. With no inspected offer, leaving is the only boundary action. Fixed enumerations resolve inspection and booking ties, stopping is chosen at equality, maxima over empty sets are $-\infty$, and empty inspection and learning products equal one.

Consumer uncertainty and researcher integration. An outer researcher draw fixes the consumer's price coefficient, private pre-search tastes, and persistent hotel-trip innovations. The price coefficient is $\beta_{p,i} = \beta_p + \sigma_p \eta_{p,i}$ with $\eta_{p,i} \sim N(0, 1)$. The same applies to each $\tilde{e}_{ijt}$ and ν_{ijt} . These draws determine realized utility when an offer is inspected and the

evidence carried into a later session. They are shared across inspection, boundary, booking, and session contributions.

The consumer knows her visible utility but not an uninspected hotel's realized match. Her reservation calculation therefore integrates a subjective forecast. For outer draw d , let $v_j^{(d)}$ be known visible utility and let $M_{jk}^{C,(d)}$, $k = 1, \dots, K$, be equally weighted nodes from the subjective entry match law used in computation. The numerical reservation equation is

$$c(\ell_j) = \frac{1}{K} \sum_{k=1}^K \left[\Psi_s^{(d)} \left(v_j^{(d)} + M_{jk}^{C,(d)} \right) - \widehat{\zeta}_{js}^{(d,K)} \right]^+. \quad (\text{D.1})$$

These inner nodes approximate the expectation in (3.17).

Let $(\Lambda^{(d)}, h_s^{(d)})$ and normalized weights $w_s^{(d)}$ approximate the conditional researcher law μ_s . The numerical outcome contribution is

$$\widehat{d}_s = \sum_{d=1}^D w_s^{(d)} \sum_{\pi_s \in \Pi(C_s)} I_s^{(d)}(\pi_s) A_s^{(d)}(a_s \mid \pi_s) B_s^{(d)}(j^* \mid \pi_s). \quad (\text{D.2})$$

All three factors use the same draw and path, and the log criterion uses $\log \widehat{d}_s$ after integration. Computation uses Gauss–Hermite quadrature for price heterogeneity and common simulation draws for matches.

Learning. After observing Y_s , the researcher weights every draw and compatible order by $G_s = I_s A_s B_s$. For $d_s > 0$, the normalized end-of-session law can be written

$$\mu_s^+(\mathrm{d}\Lambda, \mathrm{d}h_s, \pi_s) = \frac{G_s(Y_s, \pi_s \mid \Lambda, h_s) \mu_s(\mathrm{d}\Lambda, \mathrm{d}h_s)}{d_s(Y_s \mid \mathbf{r}_s, w_s, H_{s-1}^R)}.$$

Normalization sums over π_s as well as integrating the draw and history. Transitioning the completed histories and conditioning on the next observed list produce μ_{s+1} .

Within each history that continues, the consumer separately updates b_s by (3.10). Eligible evidence consists only of first organic inspections on default or filtered lists. Sorted, sponsored, and previously inspected hotels do not supply any weight updating evidence. The update follows the continuation decision. Until that decision, entry beliefs, marginal forecasts, and continuation functions remain fixed.

D.3 Continuation Values and the Lists a Tool Returns

A one-list-ahead forecast. The entry-fixed restriction in Assumption 2 specifies when forecasts are revised. Computation adds a horizon approximation: when valuing continuation

now, the next list is treated as terminal, with only booking and leaving at its boundary. This approximates $\widehat{V}_{s+1|s}^E$ inside $\widehat{\mathcal{W}}_s^a$ in (3.14). The approximation omits the option of lists beyond that forecast list.³⁷

Let $\mathcal{J}$ be its forecast inspectable set, excluding identities already inspected when the current entry forecast is formed. For its independent predictive utility laws, define the terminal stop map

$$\Psi_{\bar{S}}(u) = \tau_B \log[\exp(u/\tau_B) + \exp(u_0/\tau_B)], \quad \tau_B = \sqrt{6}\sigma_B/\pi.$$

Let ζ_j solve the reservation equation with this terminal map and the forecast slot cost. Appendix B.5's capped-reward identity then gives

$$W(\mathcal{J}, \bar{u}) = \mathbb{E} \left[\max \left\{ \Psi_{\bar{S}}(\bar{u}), \max_{j \in \mathcal{J}} \min \{ \Psi_{\bar{S}}(U_j), \zeta_j \} \right\} \right]. \quad (\text{D.3})$$

This is the forecast value of search already net of inspection costs and including its terminal boundary payoff. The reservation caps incorporate those costs, so they are not subtracted again. With no new inspectable offers, W equals the terminal stop value of the revealed best offer and outside option.

Forecast list laws. For a re-issued default query, each simulated list draws without replacement as many hotels as the current list from those displayed in the market during the observation window, with their card information. Hotels are ordered by the score the consumer attributes to the platform at the current belief mean, with a fresh ranking disturbance. The simulated order supplies match forecasts using the same decoder as an observed platform-ordered list. A filter uses the same draw but keeps the share of hotels that filters keep on average in the data, preserving the platform order within that set. A sort instead orders the draw by price, assigns the unconditional match forecast, and keeps displayed-slot inspection costs.

Averaging (D.3) over these simulated lists approximates the continuation forecast:

$$\widehat{\mathcal{W}}_s^a(u; \omega_s) = \int \widehat{V}_{s+1|s}^E(u, \omega') \widehat{\Gamma}_a(d\omega' \mid \omega_s) \simeq \mathbb{E}_{\widehat{\Gamma}_a}[W(\mathcal{J}', u)],$$

with the future list's other forecast inputs suppressed in W . Current list-opening cost, tool intercepts, same-tool shifters, and current-list covariates enter separately in (3.14).

³⁷De los Santos and Koulayev (2017) use the same one-session-ahead valuation and give its rationale: every continuation action remains available at every session, so what is lost by not looking beyond the next list is the option value of a third list, which is small in our context with most consumers search for 2–3 sessions only.

D.4 Computing Rank Decodes

Candidate-specific Laplace approximation. For each weight α , the posterior in (3.8) combines the organic-order likelihood with the subjective Gaussian prior for $m^{P,C}$, which has no closed form. We approximate the posterior by a Laplace approximation. Let t_α^* be the vector of predictable-match posterior modes and let H_α be the negative log-posterior Hessian, the Laplace approximation is the multivariate normal law with mean t_α^* and covariance H_α^{-1} .

Marginalizing this normal approximation and adding the independent innovation gives

$$m_j \mid D_s, \alpha \sim N(t_{\alpha,j}^*, (H_\alpha^{-1})_{jj} + 1).$$

The added one is the innovation variance in the full match. For example, the corresponding approximate density uses the standard normal density φ with standard deviation $\sqrt{(H_\alpha^{-1})_{jj} + 1}$. An uninspected sorted or sponsored hotel instead receives the unconditional full-match law $N(0, \sigma_{m^{P,C}}^2 + 1)$. An inspected hotel keeps its known match and offer.

Belief mixture, mean evaluation, and interpolation. The full theoretical forecast is $F_j^C(m) = \int F_j^C(m \mid D_s, \alpha) b_s(d\alpha)$. Computation evaluates the candidate decode at $\bar{\alpha}_s = \int \alpha b_s(d\alpha)$ and linearly interpolates between decodes at the two neighboring grid weights.

Candidate-specific densities remain separate for learning. For each grid weight, form the product of its normalized marginal densities over $\mathcal{C}_s^{\text{new}}$, then multiply by the current weight probability and normalize across weights as in (3.10).

Grid decodes can be stored for repeated evaluation. Reuse requires agreement of every decoder input: the order, non-order list information, candidate weight, parameters, and consumer-known visible utilities, including private tastes and price heterogeneity. Within the likelihood, these inputs belong to the same persistent draw and chronological history used for the choice factors.

D.5 A Conceptual Joint-Likelihood Alternative

Our current estimation procedure selects ranking fit after fitting the demand. A joint-likelihood estimator would choose the common parameters to maximize the conditioned joint probability of rankings and consumer outcomes. We tried this joint approach and found that it was computationally much slower than the current two-stage procedure, and that it can produce similar results if we penalize the ranking component appropriately in the likelihood.

Appendix E Counterfactual Details

This appendix specifies the comparison outcomes, distinguishes surplus measures, and presents the additional counterfactual results.

E.1 Regimes and Comparisons

Write $r \in \{(M, E), (N, E), (M, T), (N, T)\}$ for the four regimes: monetized (M) or neutral (N) ranking, each with estimated entry beliefs (E) or transparency (T). (M, E) represents the status quo. The actual monetization weight is 0.78 under M and zero under N . The inherited-belief specification E retains the estimated trip-entry distribution and subsequent updating. Estimated preferences, search parameters, recovered costs, the commission contract, and the sponsored-placement mechanism are common to the regimes.

Let $\rho \in \{\text{fix}, \text{eq}\}$ denote the price regime. Fixed-price outcomes use the same constructed nightly rates in every policy, while equilibrium outcomes use the conditional price vector solved for each r . The common population contains N_T trips, including noninspectors and nonbookers. For outcome X , Table 5 reports

$$100 \left(\frac{X_r^\rho}{X_{M,E}^\rho} - 1 \right).$$

Totals and their per-trip counterparts have the same percentage changes because N_T is common. Adding transparency to neutrality instead compares $X_{N,T}^\rho$ with $X_{N,E}^\rho$. Comparing ranking rules under transparency uses $X_{M,T}^\rho$ as the reference for $X_{N,T}^\rho$. The formulas below suppress ρ .

E.2 Bookings, Search, Prices, and Commission

Let B_{it} indicate an on-platform booking, J_{it} be the booked hotel when $B_{it} = 1$, and S_{it} be the number of lists opened before the trip ends. Let $\mathcal{C}_{its}$ contain hotels newly inspected in session s , and let ℓ_{ijts} be the displayed slot at that inspection. This set includes all new inspections, not only the subset eligible for weight learning. Define

$$I_{it} = \sum_{s=1}^{S_{it}} |\mathcal{C}_{its}|, \quad C_{it}^I = \sum_{s=1}^{S_{it}} \sum_{j \in \mathcal{C}_{its}} c(\ell_{ijts}), \quad (\text{E.1})$$

where I_{it} is the inspection count and C_{it}^I the inspection expenditure in utility units. Both are zero on a trip without inspection. Regime totals are $\sum_{(i,t)} B_{r,it}$ for bookings, $\sum_{(i,t)} I_{r,it}$ for

inspections, and $\sum_{(i,t)} C_{r,it}^I$ for inspection expenditure. Dividing by N_T gives their per-trip counterparts. Pooled cost per inspection is

$$\bar{C}_r^I = \frac{\sum_{(i,t)} C_{r,it}^I}{\sum_{(i,t)} I_{r,it}},$$

in utility units per inspection.

Search-intensity measures use different denominators. Mean sessions and continuations per trip are, respectively,

$$\frac{1}{N_T} \sum_{(i,t)} S_{r,it}, \quad \frac{1}{N_T} \sum_{(i,t)} (S_{r,it} - 1).$$

Inspections per session divide total inspections by total sessions. For inspections per active trip and per active session, the respective denominators are

$$\sum_{(i,t)} \mathbf{1}\{I_{r,it} > 0\}, \quad \sum_{(i,t)} \sum_{s=1}^{S_{r,it}} \mathbf{1}\{|C_{r,its}| > 0\}.$$

These denominators select trips or sessions with an inspection and can change across policies.

Let P_{it} be the booked hotel's price in dollars per room-night. The average price and commission are

$$\bar{P}_r = \frac{\sum_{(i,t)} P_{r,it}}{\sum_{(i,t)} B_{r,it}}, \quad R_r = \gamma \sum_{(i,t)} P_{r,it}. \quad (\text{E.2})$$

Equation (E.2) states commission on the booked nightly-rate basis used with booking demand in (3.19). The rate is a nightly price on Appendix C.4's constructed 2013-dollar scale. Its mean conditions on booking and can change through booking composition as well as hotel repricing. A hotel's or price quintile's booking share likewise uses the regime's total bookings as its denominator. The commission rate in (E.2) is contractual $\gamma = 0.15$.

E.3 Consumer Surplus

A booked hotel contributes its realized utility $U_{i,J_{it},t}$ from (3.1), including price, the persistent pre-search taste, and the revealed match. A trip without a booking contributes the

outside option $u_{0,it}$. Terminal outcome utility and the two reported means are

$$D_{it} = \begin{cases} U_{iJ_{it}}, & \text{if } B_{it} = 1, \\ u_{0,it}, & \text{if } B_{it} = 0, \end{cases} \quad (\text{E.3})$$

$$S_r^{\text{out}} = \frac{1}{N_T} \sum_{(i,t)} D_{r,it}, \quad S_r^{\text{inspect}} = \frac{1}{N_T} \sum_{(i,t)} (D_{r,it} - C_{r,it}^I). \quad (\text{E.4})$$

Table 5 labels these “outcome surplus” and “surplus net of inspection costs.” Both are utility levels per trip, including nonbookers, under the stated outside-option normalization.

We are computing the surplus measure including continuation outcomes as follows. Let a_{its} be the action chosen at the boundary of session s . Every boundary before the realized terminal session, $s < S_{it}$, selects a continuation action. Its current deterministic flow, excluding the forecast of later payoff, is

$$g_{its}(a) = -c_{\text{sess}} + \kappa_a + \theta_a \mathbf{1}\{a = k_s\} + b'_a x_s^{\text{list}}.$$

The intercept is a relative continuation value shifter. The reuse term adds value when its coefficient is positive and subtracts value when negative, as in (3.14). The complete realized path payoff is

$$\mathcal{U}_{it} = D_{it} - C_{it}^I + \sum_{s=1}^{S_{it}-1} g_{its}(a_{its}) + \sum_{s=1}^{S_{it}} \varepsilon_{a_{its}s}^B. \quad (\text{E.5})$$

This counts terminal utility once, each new inspection cost once, the flow from each continuation once, and the selected shock at every reached boundary, including the terminal one. The first list incurs no additional continuation cost. Price, the pre-search taste, and the match already enter D_{it} and are not added or subtracted again.

An important remark is that the continuation forecast $\widehat{\mathcal{W}}_s^a$ in (3.14) helps determine the chosen action but does not act as an additional realized reward in (E.5). Later utilities and costs are recorded when they occur on the path. Adding the subjective forecast to them would count future payoffs twice. Likewise, the capped-reward forecast in (D.3) is already net of its forecast inspection costs. Those costs are not subtracted again when that forecast is used in the decision rule. Its entry-fixed, one-list-ahead valuation remains distinct from realized payoff under actual subsequent choices.

Separate the list-opening cost from the other deterministic continuation shifters by

defining

$$G_{it} = \sum_{s < S_{it}} [\kappa_{a_{its}} + \theta_{a_{its}} \mathbf{1}\{a_{its} = k_s\} + b'_{a_{its}} x_s^{\text{list}}], \quad E_{it}^B = \sum_{s \leq S_{it}} \varepsilon_{a_{its}s}^B.$$

Averaging the path accounting gives

$$\mathcal{W}_r = \frac{1}{N_T} \sum_{(i,t)} \mathcal{U}_{r,it} = S_r^{\text{inspect}} - c_{\text{sess}} \frac{1}{N_T} \sum_{(i,t)} (S_{r,it} - 1) + N_T^{-1} \sum_{(i,t)} G_{r,it} + \frac{1}{N_T} \sum_{(i,t)} E_{r,it}^B. \quad (\text{E.6})$$

Expected payoff averages this identity over the regime's realized paths.

E.4 Paths Across Monetization Weights and Common Beliefs

Changing the actual weight. Figure 14(a) varies the actual monetization weight from zero to one while keeping the estimated consumer's belief and subsequent updating. Bookings and outcome surplus are expressed as percentages of their status-quo levels within the same price regime. Solid lines hold prices fixed and dashed lines re-solve hotel prices. Along the displayed path, both outcomes decline with monetization in both price regimes, with the largest changes concentrated between weights of 0.4 and 0.8. The estimated weight of 0.78 lies in this steeper region. Below a weight of 0.4 the commission revenue is relatively too weak to reorder the top of the list, above it the ranking is more likely to hand the top of the list to the priciest quintile, and the platform's estimated high weight sits in the steep part of the displayed path.

This shows that the economic cost of monetization is governed by the recommendations displaced at the margin, not by the parameter's magnitude in isolation. A small weight change may be important where many relevant offers are close to exchanging positions and unimportant elsewhere.

The same belief under both rules. Panel (b) instead varies the perceived weight used in both ranking regimes and compares neutrality with monetization at fixed prices. At each point, the specified belief is held the same across the two rules. For example, a perceived weight of $x = 0.4$ understates the monetized rule's actual weight $0.4 < 0.78$ and overstates the neutral rule's ($0.4 > 0$). It is important to note that this is not transparency. Consumers are not given a different perceived weight corresponding to each rule and thus can have biased beliefs under both rules. The comparison isolates the effect of the belief on the relative performance of the two rules, holding the actual weight and prices fixed.

Moving right on the horizontal axis means consumers perceive more monetization. Their belief becomes closer to the monetized rule (thus relatively less overtrust), and




Figure 14. Bookings and outcome surplus across weights and beliefs

Notes: "Surplus" denotes outcome surplus S^{out} throughout. Panel (a) varies actual monetization with the estimated consumer belief and its subsequent updating. Each outcome is a percentage of its (M, E) level within the same price regime. Solid lines use fixed prices and dashed lines re-solved prices. Panel (b) holds prices fixed and compares neutrality with monetization, with the belief held at the same value of the horizontal axis in both ranking regimes.

farther from neutral-ranking (thus relatively more skeptical). The booking gain from neutral-ranking relative to the monetized ranking declines along the displayed path: it is 29% near a belief of zero, 27% at the estimated entry belief, and 25% at 0.4. Each gain uses bookings under monetization at that same belief as its reference. The outcome-surplus gain also declines on the displayed path.

This shows that a consumer who attributes less information to rank gains less from an improved ordering. Making the ranking neutral genuinely improves the recommendation as the hotels at the top now suit the consumer's needs rather than pay the platform. But that improvement reaches the consumer's behavior only through her willingness to follow the list. A trusting consumer inspects near the top, so she is directed to the better hotels and books more. A skeptical consumer reads little into any rank and searches on other grounds, so a better ranking redirects her less.

From the left to the right of the horizontal axis, the consumer's belief is less aligned with the neutral ranking but more aligned with the monetized one. Under the monetized ranking, a more skeptical consumer inspects less and books less. Neutrality is supposed to expand trade. However, without a belief change, the more skeptical consumer wastes the opportunity for increased trade and such misplaced skepticism costs a good ranking much more than what it costs a bad one, which reduces the gap in bookings between the two

rankings and decreases the gain from neutrality. So trust and neutrality are complements as the effect of neutrality is larger when consumers trust more.

To summarize, greater perceived monetization discounts the informational content of prominent positions under both rules. Consumers then respond less to the better recommendations supplied by neutrality, so its advantage is smaller. This comparison illustrates how the interpretation of rank conditions the return to improving the ordering, separately from the policy experiment that changes the interpretation along with the rule.