Consumer Preferences and Optimal Reserve Prices in Position Auctions

Joshua Higbee* Samuel Higbee† Ali Hortaçsu‡

November 3, 2026

Abstract: We study the implications of rich models of consumer preferences on the design of position auctions. Using a natural experiment with reserve prices on a travel meta-search platform, we find that consumers' link-specific preferences are strongly correlated at the hotel level, while bids increase (elasticity 0.122) and the number of bidders per auction decreases (elasticity -0.182) in response to an increase in reserve prices. We estimate a semiparametric nested logit model of consumers' click and booking decisions, compute a doubly robust reserve price policy, and compare the policy implied by a structural model to a nonparametrically optimized reserve price. Our flexible demand model predicts 6% higher revenues and lower consumer welfare losses from the revenue-maximizing reserve price compared to less flexible models. Bidder- and rank-specific reserve prices differ significantly from uniform within-hotel reserve prices but offer only modest revenue improvements.

Keywords: position auctions, advertising, reserve prices, platform design
JEL codes: D12, D44, D47, L81

*The Ohio State University, higbee.25@osu.edu
†UNC Chapel Hill, sdhigbee@unc.edu
‡University of Chicago, hortacsu@uchicago.edu
Working paper draft. We thank Emre Mutlu for excellent research assistance. All errors are our own.

1 Introduction

Trillions of dollars of annual online transactions are shaped by digital advertising, which itself is a multi-billion dollar industry. In the U.S. alone, ad spend via real-time auctions was approximately $100 billion in 2025 (Interactive Advertising Bureau and PricewaterhouseCoopers 2025, 2026). Presumably, at least in the minds of advertisers, spending on digital advertising leads to increased visibility and revenue. However, advertisers often must compete for clicks directly against their product-market competitors through a digital platform, in which consumers are presented with a menu of advertisers' products and choose which, if any, to click. Modern platforms use knowledge about advertiser competition, in the form of economic theory and rich data, to design their own search auction mechanisms.

Despite the importance of consumer preferences in converting visibility to clicks, most studies of online auctions use very basic demand models. Consumer preferences are central to the advertisers' profit maximization problem, as each firm is only paid after successful conversions from clicks to purchases. Firms' willingness-to-pay for each click therefore depends not only on their prominence but also on their substitutability with neighboring products. In determining its own advertising strategy, each bidder must ask: are consumption choices dominated by preferences for product characteristics or relative visibility? The platform's choice of reserve price hangs on the answer to this question, as the reserve price determines whether consumers are exposed to many advertisers with products of varying quality or a few key players with high willingness-to-pay. Without fully accounting for the role of consumer preferences and product-market competition, standard analyses of position auctions cannot capture the effect of reserve prices— and the choice sets they induce— on consumer behavior. Consequently, counterfactual estimates of the optimal reserve price may be incorrect.

We explore the implications of richer demand models for the design of position auctions in the context of a travel meta-search engine, pseudonymized as Platform.com. We begin from the premise, already pervasive in standard demand system analysis, that careful modeling of consumers' preferences and substitution patterns is critical to understand market counterfactuals and optimal pricing decisions. On Platform.com, consumers search for hotel bookings and are presented with a menu of hotels, each of which can be booked through one of several online travel agencies (OTAs). These agencies compete for prominence by submitting bids which determine both the set of agencies offered to consumers and the order in which they are presented. When setting the reserve price, Platform.com must be able to predict whether changing the set of agencies presented to consumers induces consumers to choose higher-bidding agencies still in the choice set, or whether consumers will decide not to book when presented with limited options. Thus, substitutability of agencies is a key

consideration when designing these auctions.

We first use a numerical example to show that properly modeling consumer substitution patterns has important implications: incorrectly modeling consumer choices as completely independent across products will both understate the profitability of higher reserve prices and exaggerate the consumer welfare loss from a smaller choice set. If consumers can always direct their clicks toward close substitutes (for example, reservations at one hotel offered by multiple agencies), then increasing the reserve price to exclude some agencies from the choice set may not lead to a significant reduction in total clicks. Instead, the consumer may even be diverted to a higher-revenue option. After accounting for heterogeneous search costs across listings, the welfare implications of such a change are unclear.

We then present quasi-experimental evidence of high substitutability between agencies on Platform.com. We look at the effects of a limited reserve price change implemented by the platform, in which a selected subset of queries was assigned a new reserve price according to a proprietary formula, while a majority of other queries were unaffected. We estimate an elasticity with respect to the reserve price of 0.076 for bids by bidding agencies and -0.188 for the number of bidders per auction, as low-bidding agencies were excluded by the increased reserve price. Direct measures of the effect of reserve prices on clicks and revenue are noisy and statistically difficult to distinguish from zero, but we cannot rule out economically meaningful positive effects.

We also use the quasi-experimental variation to directly test for nested consumer preferences at the hotel level. IV estimates indicate that excluding a competing agency from the choice set at the same hotel increases the each link's expected clicks by 14.8% on average, while a 1% increase in competitors' prices increases clicks by 2.5% for within-hotel agencies. In contrast, we estimate a small but positive relationship between own-link clicks and the price and availability at other-hotel listings: adding an agency at another hotel is associated with a 0.4% increase in clicks, and increasing the minimum log-price across other hotels by 1% yields a mere 0.3% decrease in clicks. While this reduced-form analysis cannot account for all relevant characteristics of the available hotels within a query, our findings indicate that consumers' preferences are largely hotel-specific. That is, link-level click counts are far more responsive to the presence and prices of competing agencies at the same hotel than for those at other hotels.

Motivated by this empirical evidence, we develop a semiparametric structural model of consumer demand for hotel bookings and estimate it on historical data from Platform.com. In contrast to most analyses of position auctions, our model allows a rich nesting structure to drive substitution patterns between hotels and agencies. This implies a more flexible mapping between travel agencies' strategic decisions and Platform.com's expected profit.

In addition to the nested logit structure, we allow for substantial flexibility for observed determinants of consumer utility and search costs by representing each as neural networks of high-dimensional consumer covariates. On the supply side, we combine the exact invoice formula from this auction mechanism with reduced-form bidding and pricing functions; this allows for substantial heterogeneity in agency strategies without taking a stance on the equilibrium being played in this setting.[1]

We estimate our structural model of consumers' click decisions by exploiting the variation in links displayed to consumers in the wake of Platform.com's reserve price policy change. Our estimates confirm that consumer preferences are highly correlated within hotel nests, implying the click-through rate is much less sensitive to smaller per-hotel choice sets than would be expected absent this structure: approximately 97% of the unobserved variance in consumer utility is determined by hotel, rather than hotel-agency, preferences. Accurately accounting for such preferences also has important implications for estimated search costs: models without flexible substitution estimate large agency-rank effects, on par with the estimated hotel-rank. We also estimate our model using only observational data from before the policy change, and find broadly similar results when using the observational vs. quasi-experimental data.

We next use our model to estimate the optimal reserve price function in these advertising auctions. Misspecified choice models such as binary and simple logits predict that existing reserve prices are approximately optimal. Counterfactual reserve prices that make use of more flexible substitution patterns imply much higher revenue (and at lower welfare costs) than would be expected absent this structure: 15.8% under a nested logit model vs. approximately 10% under either a binary or simple logit click model. However, all models imply roughly comparable changes in expected clicks and consumer welfare at their respective revenue-maximizing reserve prices, with decreases of about 10-11% in each model and outcome. Allowing bidders to endogenously respond to higher reserve prices is an important feature, with nested logit models implying only a 6.6% increase in revenue under optimized reserve prices without bid adjustment. Interestingly, current reserve prices appear to be roughly optimal when using simple or binary logit models without endogenous bid adjustment, with predicted possible revenue increases of only 1.1-1.4%. Further, personal-

1. We take this approach for several reasons. The mechanism used by Platform.com is a blend of generalized first-price and second-price auctions, which to our knowledge has no known equilibrium characterization. Further, at the time of our analysis, Platform.com was still rolling out the full auction mechanism, so many auctions saw a combination of the auction mechanism as well as legacy deal structures, like non-bidding cost-per-click and cost-per-acquisition agreements. Employees of Platform.com have also characterized agency strategies as being largely driven by algorithmic optimization and time-varying strategies. As a result, and as explained in more detail in section 4.2.ii below, we do not model the equilibrium of this evolving multi-mechanism setting.

ized reserve prices that are tailored to individual bidders yield only modest improvements in expected revenue (17.8% vs 15.8%), while rank-specific reserve prices fall in between the uniform and agency specific reserves (16.6%). However, both personalized and rank-dependent reserves imply less severe decreases of clicks and consumer welfare of 8.6% each.

Relative to previous analyses of position auctions, we place special emphasis on the consumer demand model. Given the complexities of typical multi-unit auction settings, this has understandably been simplified by various equilibrium analyses; click-through-rates (CTRs) are typically either separable in bidder identities and ranks or independent across bidders (Varian 2007; Edelman, Ostrovsky, and Schwarz 2007; Edelman and Schwarz 2010; Athey and Nekipelov 2010; Börgers et al. 2013; Sun, Zhou, and Deng 2014; Cavallo and Wilkens 2014; Ostrovsky and Skrzypacz 2022; Ostrovsky and Schwarz 2023; Bergemann, Bonatti, and Wu 2025). Other work has more explicitly modeled the search process by which consumers browse through and select ads (Athey and Ellison 2011; Chen and He 2011), though generally by strict sequential consideration of links (Aggarwal et al. 2008; Giotis and Karlin 2008; Kempe and Mahdian 2008; Gomes, Immorlica, and Markakis 2009; Kuminov and Tennenholtz 2009). In contrast, we model consumers' click probabilities as depending on their underlying preferences for differentiated products.

Other studies of position auctions focus more closely on the role of reserve prices and controlling the set of displayed ads in these settings. Athey and Ellison (2011) shows that reserve prices can raise welfare by curating the set of visible options when consumers are Bayesian, while studies of choice-screen auctions examine outcomes under different auction formats and the resulting incentives (Ostrovsky 2021; Decarolis, Li, and Paternollo 2025). As in Ghosh and Mahdian (2008), Fotakis, Krysta, and Telelis (2011), Hummel and McAfee (2014), and Choi and Mela (2019) (and in other online settings outside of position auctions, like Compiani et al. 2024; Ursu 2018), there are positional externalities among ranks since per-option quality and rank-dependent search costs jointly drive consumer choice. Hummel (2016) models the choice of reserve price along with the optimal number of options to display, allowing click-through rates to depend directly on the total number of ads.

In the broader industrial organization literature, the role of flexible substitution patterns in the demand system is well-documented. Debreu (1960)'s famous "red bus / blue bus" example highlights the potential negative implications of an inflexible demand system; we modify this canonical example to develop intuition for our setting. Subsequent developments, such as the generalized extreme value framework of McFadden (1977, 1980) and the random coefficients logit model of Berry, Levinsohn, and Pakes (1993), allow additional flexibility. This flexibility has strong economic implications for optimal markups: highly correlated preferences within narrow product categories reduce the degree of market power held by

any one firm and may provoke stronger competition for interested consumers (Berry 1994; Bresnahan, Stern, and Trajtenberg 1996). The nested logit structure is attractive because it is computationally straightforward, and has been used in the travel studies in e.g. Betancourt et al. (2022). This has allowed it to be combined with semiparametric methods, including elsewhere in the travel industry (Kim 2025).

Many studies of position auctions differ from ours in their assumptions about agency-side information sets, bidding strategies, and equilibria. Often, theoretical treatments assume full information in an *ex post* equilibrium concept (Edelman, Ostrovsky, and Schwarz 2007; Hsieh, Shum, and Yang 2015). Even with more limited information, it is typical to assume stochastic quality scores around some known average bid (Athey and Nekipelov 2010; Ostrovsky and Skrzypacz 2022). We view this as unlikely in our setting for several reasons. First, due to the high dimensionality of consumer data for each query, bids can be quite complex: these can vary across hotels, length-of-stay, number of guests, source country, and booking window. Second, bidders only see aggregated *own-listing* reports provided by Platform.com instead of granular information across all competing partners and query characteristics. Third, the auction system was relatively new in this time period, with some agencies still switching over from legacy advertising contracts and others displaying strong evidence of active bid experimentation.

Our paper joins others that have implemented, evaluated, or simulated reserve price experiments on an auction platform. Many of these follow Ostrovsky and Schwarz (2023) in implementing a Myerson (1981)-style reserve price per keyword and observing the resulting increase in profits (Topinsky 2014; Dimitrellos 2022). In contrast, Sun, Zhou, and Deng (2014) only simulates results but finds similar potential for platform revenue gains. Of these, Dimitrellos (2022) in particular notes the potential for increased supply to have a positive effect on clicks; our horizontally differentiated product model accounts for this but flexibly estimates this option value through the nesting parameters. We differ by evaluating a natural experiment using our structural model. Further, our agnosticism about bidder equilibria implies our treatments are not calculated in the standard Myerson (1981) method, but rather by numerical optimization.

2 A numerical illustration

We adapt a canonical example and consider a travel platform that advertises for a bus service that can be provided by either a red or blue bus (Debreu 1960). All bus services are purchased if clicked, and the red bus earns profit 1 from consumers, while the blue bus earns profit $\pi \in [0, 1]$. We assume both buses $i \in \{r, b\}$ submit bids $\mathcal{B}_i$ to the travel company, which

operates a generalized second-price auction for the top position. Both buses are displayed as long as their bids meet the platform's reserve price m, an event we denote $d = \mathbf{1}[\mathcal{B}_r, \mathcal{B}_b \geq m]$. All consumers have the outside option of driving their own car.

Consumers' click (and purchase) decisions are represented by a simple nested logit model. Each individual has mean utility 0 for using either bus, as well as for their outside option of driving their own car. However, the less-prominent advertisement incurs a search cost c as long as it is displayed. We represent idiosyncratic preferences using a standard nested logit framework with similarity parameter λ, implying the choice shares for each option are

$$
\begin{aligned}
s_{\text{car}}(c,d) &= \frac{1}{1+\left[1+d\cdot\exp\left(\frac{-c}{1-\lambda}\right)\right]^{1-\lambda}} \\
s_{\text{bus }1}(c,d) &= \frac{1}{\left(1+\left[1+d\cdot\exp\left(\frac{-c}{1-\lambda}\right)\right]^{1-\lambda}\right)\cdot\left[1+d\cdot\exp\left(\frac{-c}{1-\lambda}\right)\right]^{\lambda}} \\
s_{\text{bus }2}(c,d) &= \frac{d\cdot\exp\left(\frac{-c}{1-\lambda}\right)}{\left(1+\left[1+d\cdot\exp\left(\frac{-c}{1-\lambda}\right)\right]^{1-\lambda}\right)\cdot\left[1+d\cdot\exp\left(\frac{-c}{1-\lambda}\right)\right]^{\lambda}}
\end{aligned}
$$

The platform can choose its website design, and therefore determine the search cost c of discovering the second bus, as well as the reserve price m. Note that when $d = 0$, the click (and purchase) shares of the first bus and the car are equal at $\frac{1}{2}$.

We focus on the outcome of the locally-envy free equilibrium of Edelman, Ostrovsky, and Schwarz (2007) that yields the lowest platform profit. This corresponds to the generalized English auction's unique equilibrium in strategies that are continuous in bidders' valuations. In that game, the blue bus (having a lower profit per click) will drop out at bid

$$
\mathcal{B}_b(\pi, c, d, m) = \pi - \frac{s_{\text{bus }2}(c,d)}{s_{\text{bus }1}(c,d)}(\pi - m)
$$

and pay m, the platform's reserve price for the click advertisement. The red bus can always win the first position, and we assume that it bids its value ($\mathcal{B}_r = 1$). The red bus pays $\mathcal{B}_b(\pi, c, 1, m)$ if $d = 1$ and m otherwise.

With this framework in mind, the platform must now determine the optimal reserve price m^*. Note that $d = \mathbf{1}[\pi \geq m]$ since the blue bus will bid all the way up to its valuation. Since both the red bus' payment is increasing in m, and the blue bus' payment is as well while $d = 1$, the platform optimally sets the reserve price as either π or 1. Which reserve is preferred, however, depends on key features of the demand system. We consider both the nested logit demand system and a more typical abstraction used in this setting: constant

click probabilities. Suppose in contrast to the system above that the share of the red bus is the same regardless of whether the blue bus is listed as an option; if it is omitted, its share goes completely to the outside option.

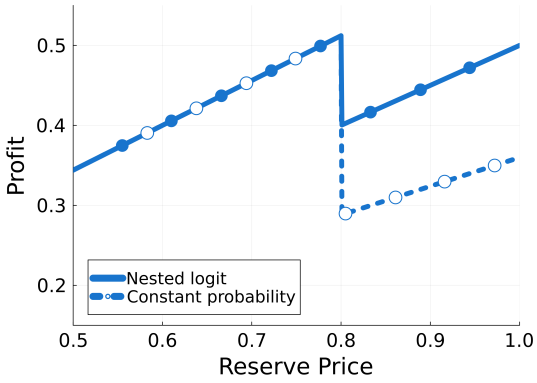
Figure 1 shows platform profits and consumer surplus as a function of the reserve price for various demand systems. With low correlation in within-nest (bus) preferences (panel (a)), including both buses maximizes platform profits. However, increased correlation between bus preferences may make the optimal reserve price either depend on the nature of the demand system (panel (b)), or imply it is always optimal to set a high reserve price and exclude the lower-value option (panel (c)). Figures A1.1 and A1.2 in Appendix A1 provide additional intuition for these cases, showing the minimum π_b for which including both buses is profitable for varied λ and c. These modeling assumptions have significant implications for welfare. Figure 1 shows that welfare loss due to fewer choices decreases significantly as λ increases, and that this loss is dramatically overstated if one does not take into account the substitution to the within-nest option.

This simple example abstracts from other relevant features of the platform's reserve pricing problem. First, per-click values are unlikely to be known to the platform due to information about cost margins or other features that is known only to the competing bidders, so the platform has imperfect information. Second, bidding strategies themselves may vary over time due to random shocks unobserved to both the platform and competitors, so the bidders may not play an *ex post* equilibrium game as depicted here. Third, the presence of additional bidders and any heterogeneity in click-through rates would complicate the bidding functions and yield additional local maxima to be evaluated by the platform; particularly in the presence of high-dimensional covariates, this could be intractable for the platform to solve exactly. Fourth, the platform we study does not use a pure generalized English auction, but rather a hybrid of generalized first and second price auctions with (to our knowledge) no derived equilibrium strategy. Thus, it is necessary to empirically evaluate the game as it is played on the platform.

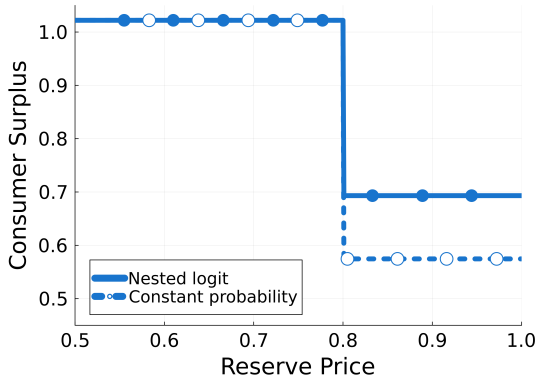
3 Data and reduced-form evidence

We first present an overview of the travel search platform and the natural experiment we use. We then present quasi-experimental estimates of the effect of reserve price on bids, prices, clicks, and revenue. We also present evidence for nested click probabilities in the form of reduced-form effects of choice sets and competitor prices on clicks.

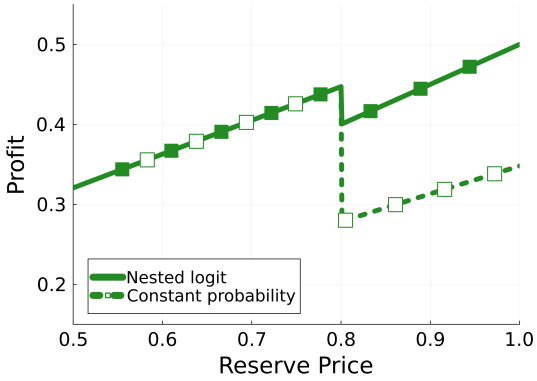
Figure 1: Profits and consumer surplus as a function of reserve prices for different demand systems



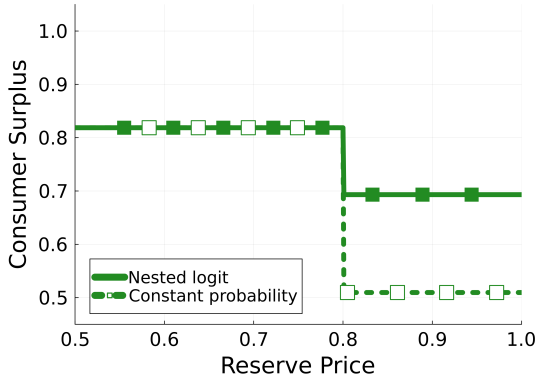
(a) Profit, $\lambda = 0.0$



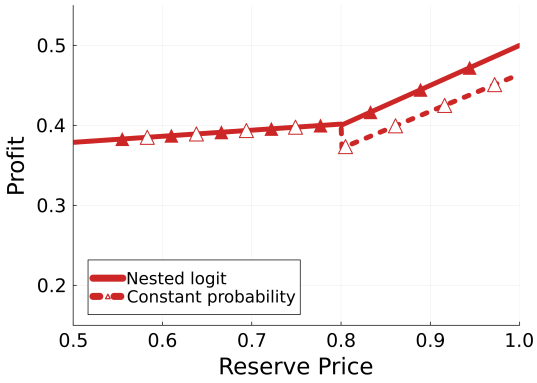
(b) Consumer surplus, $\lambda = 0.0$



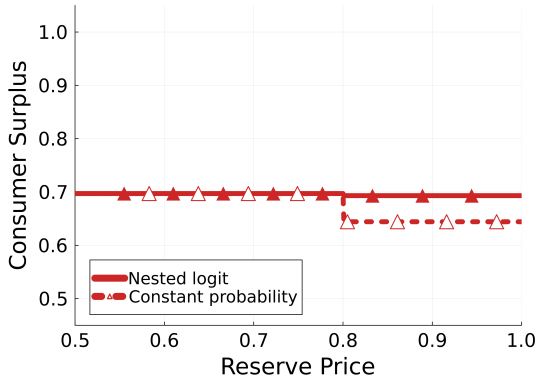
(c) Profit, $\lambda = 0.5$



(d) Consumer surplus, $\lambda = 0.5$



(e) Profit, $\lambda = 0.9$



(f) Consumer surplus, $\lambda = 0.9$

Notes: Each panel displays profit or consumer surplus as a function of different reserve prices, for different values of the nesting parameter λ. Solid lines represent the nested logit demand system. Dashed lines, denoted "constant," represent the alternative choice system in which the red-bus choice share does not vary with the presence of the blue bus.

3.1 Platform and data description

Our data comes from the hotel booking section of a travel meta-search platform, Platform.com. Consumers who visit the platform enter itinerary details, such as the intended travel dates and location, and observe a list of hotels as well as prices offered by competing online travel agencies. These agencies make their price and availability data available to Platform.com, but their page position is determined in part by cost-per-click (CPC) bids offered in a generalized position auction format.[2] The platform aggregates this information and chooses the rankings of agencies and hotels as presented to the target consumer.

Figure 2 presents a generic representation of the platform layout, as if from the mobile app view. Consumers have the option to click directly on a link provided from a general search page, as represented in panel (a), or move to a more specific hotel-detail view as portrayed in panel (b). Hotel ranks on the search page are determined via a proprietary algorithm, but they depend on various consumer-level query features such as location, price, and rating preferences as well as projected revenue per hotel. The ranking of agencies within hotels is likewise determined by a proprietary algorithm that broadly depends on expected revenue per agency listing for that hotel.

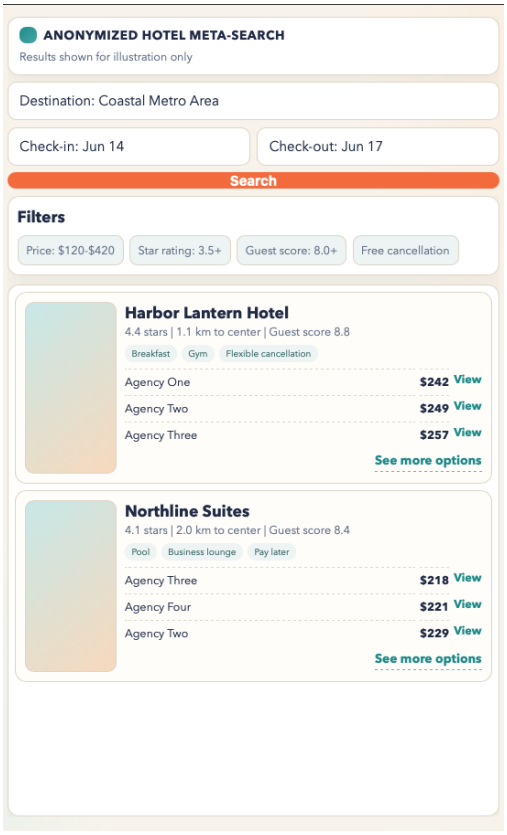
In our data, we restrict attention to hotel booking services offered in Sydney, Australia between July and November 2025, using a random 10% subsample of all queries. We observe link-level (hotel by agency) prices, bids, and ranks, both for agencies within a hotel and for each listed hotel in a query. We also observe consumer characteristics like length-of-stay, booking time, number of guests, and origin market, for each of the queries in our subsample. Most critically for our analysis, we observe which (if any) of the listed results was clicked on by a consumer, and the total number of clicks per link.

We now present several facts about the dataset of interest, beginning with empirical click probabilities as a function of ranks in Figure 3. Panel 3a demonstrates that a large majority of clicks are received by agencies in the top position for their hotel listing, followed by a sharply declining number for subsequent ranks; panel 3b shows that this near-monotonic relationship between agency rank and click probability holds broadly regardless of the number of agencies per auction. Within a given hotel auction, higher ranks are associated with more clicks, with an outside share allocated to the highest-rank agency in any given auction.

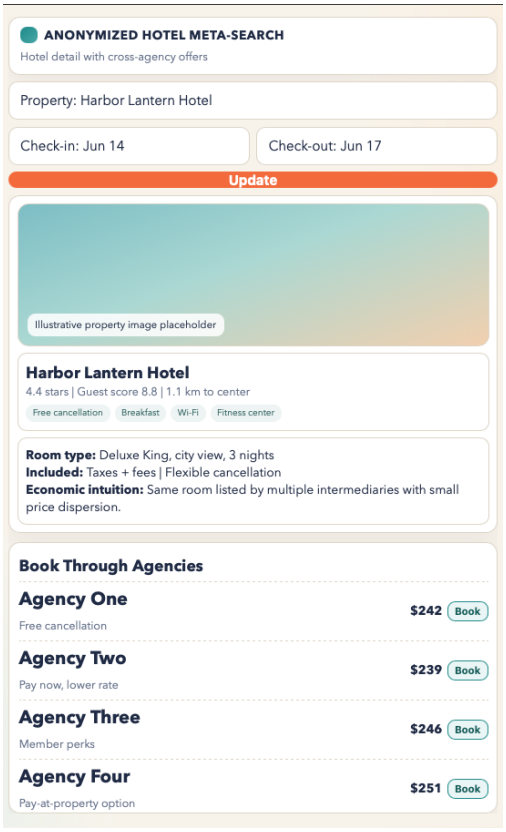
Figure 4 demonstrates the relative size of agencies' expected auction payments, bids, and reserve prices. Panel 4a shows the ratio of expected cost-per-click for different bidder

2. In the period in which we study the platform, not all agencies had switched over to cost-per-click bidding, and instead continued to use legacy business models such as cost-per-acquisition advertising. The ranking mechanism accounts for this mismatch by converting all revenue into an expected cost-per-click, or ECPC.

Figure 2: Search and hotel detail views for a generic travel search platform



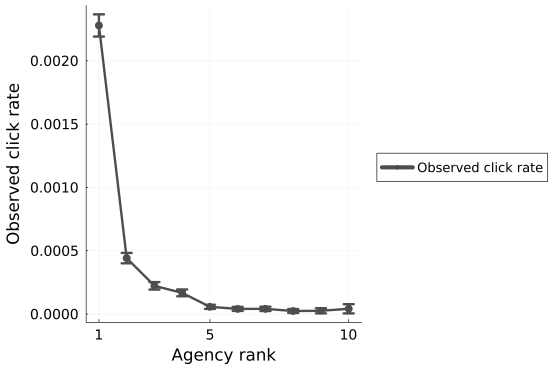
(a) Generic search results view



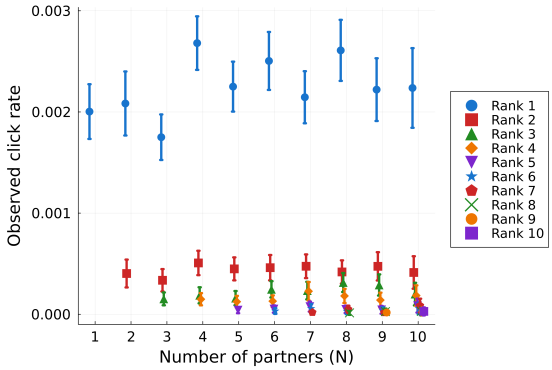
(b) Generic hotel detail view

Notes: These images were generated using GPT-5.3-Codex to represent the approximate web design of our anonymized data source. Panel (a) represents the search results view, and panel (b) represents the hotel detail view. Price and agency information is available for all visible links in both view options, though not all agencies may be immediately viewable from the search results page.

Figure 3: Agency participation and clicks in auctions



(a) Empirical click rate by agency rank



(b) Empirical click rate by agency rank and participation

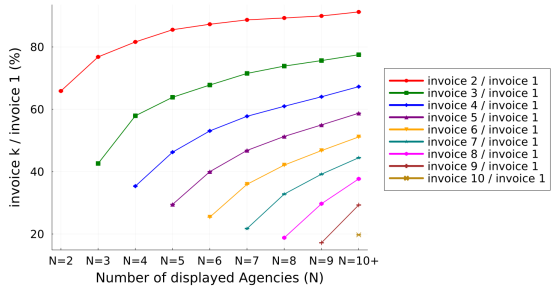
Notes: Panel (a) plots the ratio of observed clicks to observed listings for at each possible agency rank, averaging across all auctions. Panel (b) plots this ratio conditional on N agencies appearing in each auction. Some auctions have more than 10 agencies listed; these were omitted from the figures for clarity.

numbers relative to the first-place bidder's payment, conditional on the number of bidders. Expected payments increase with the number of competing agencies, suggesting endogenous entry into higher-value auctions. However, agencies' invoices vary significantly: the bottom half of bidders' expected payments are under 60% of the expected payment for the first-place bidder. Panel 4b demonstrates how close existing reserve prices (reserve prices) are to various features of the bid distribution. reserve prices are 20% or less than the highest bid in approximately 75% of auctions, and 20% or less than the winning bid in over 60% of auctions. Even the lowest bids are frequently much higher than the reserve price: in over 50% of auctions, the lowest bid is at least twice the reserve price. In sum, neither the reserve price nor additional bidders appear to put much pressure on per-click payments or significantly increase the aggregate click probability of a given hotel.

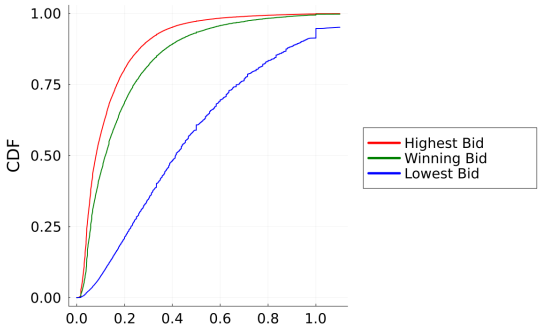
3.2 Effects from a change in reserve price policy

In September 2025, Platform.com changed its reserve price formula to increase reserve prices for specific hotels and queries with higher estimated profitability. While the formula is proprietary, we use detailed link-level data to recover which reserve prices were increased and by how much. Critically, the affected reserve prices included a random component which was held fixed within known query-characteristic groups, which we use to identify the causal effect of the change in reserve prices as a result of this policy change. Figure 5 summarizes the change in reserve prices around the policy change, including the pre-change pilot window

Figure 4: Auction competitiveness and reserve prices



(a) Expected invoice per click vs. first place payment



(b) Distributions of reserve price-bid ratios

Notes: Panel (a) plots the average (with 95% confidence interval) invoice per click for each agency rank, divided by the highest-ranked agency's estimated invoice, conditional on varying numbers of agencies per auction. Panel (b) plots the empirical CDF of the reserve price-bid ratios for the highest, first-place, and lowest bid; the highest and first-place bids do not necessarily coincide due to Platform.com's bid scoring and upranking rules.

in which a similar policy was temporarily tested. On average, this policy involved only a small increase (≈£0.05), but a smaller portion of queries saw a more dramatic increase.

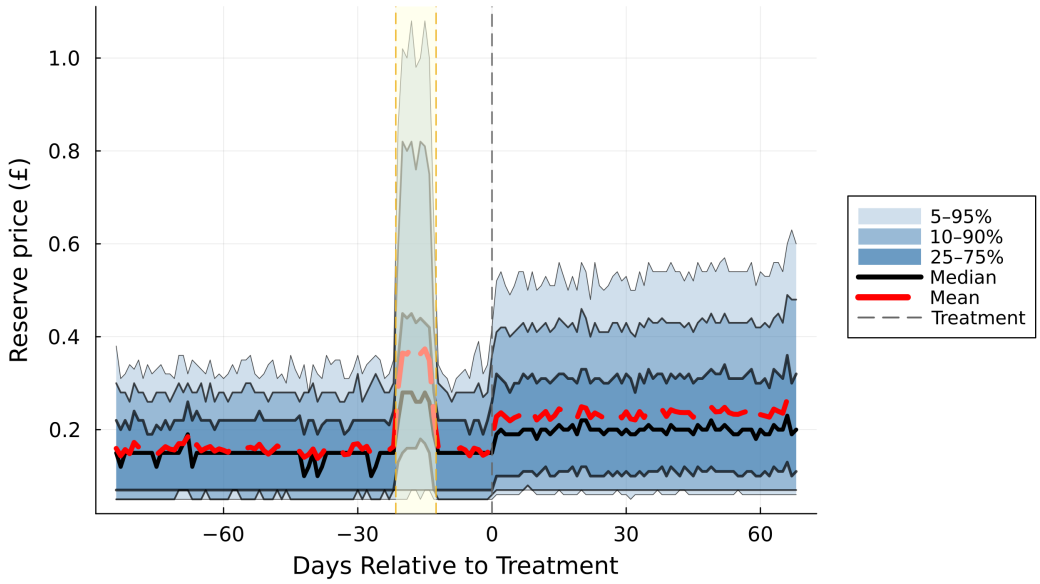
We study the effect of this targeted reserve price increase using a difference-in-differences framework across multiple outcomes of interest. Within the period before and after the policy change (but dropping the pilot window), we estimate the following class of Poisson regressions for general weakly-positive hotel-level outcomes y_{iht}:

$$\mathbb{E}[y_{iht} \mid \text{Data}_{iht}] = \exp\left(\alpha \cdot Post_{iht} \cdot \Delta \ln(\text{reserve}_{iht}) + x'_{iht}\beta + \lambda_t + \eta_j\right) \tag{1}$$

where λ_t is a date-level fixed effect, η_j is a unit-level fixed effect; each unit j is the reserve price randomization unit (hotel by market by device) interacted with itinerary length-of-stay, which is known to enter the reserve price formula. The central variable of interest, $\Delta \ln(\text{reserve}_{iht})$, is constructed by subtracting the previous policy's reserve price from the realized reserve price in the post-treatment period. For bid and price regressions, we analyze agency-hotel-level outcomes and further interact these unit-level fixed effects with agency-specific indicators. Controls x_{iht} include additional fixed effects for binned booking window and number of guests. Since the reserve price only applies to bidders, we report results only for links and queries corresponding to bidding agencies.

In Figure 6 we present the estimated DID coefficients for various outcomes of interest. As predicted in the related generalized second-price auction setting of Edelman and Schwarz

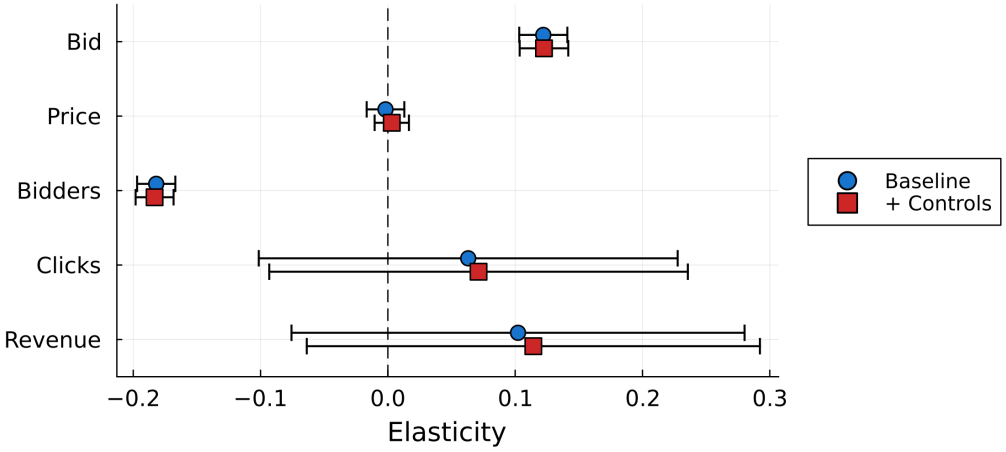
Figure 5: Distribution of reserve prices around policy shift



Notes: This plots the distribution (mean, median, and 3 quantile bands) of the reserve prices observed in Sydney auctions in the weeks before and after the permanent policy change. The pilot window, in which a modified version of this policy was temporarily implemented, is shaded in light yellow approximately 3 weeks before the full rollout.

(2010), increasing the reserve price increases the resulting bids of participating agencies; prices had no corresponding change. The number of bidders per hotel decreased significantly, which is consistent with additional agencies being excluded due to their bids being lower than the new reserve prices in affected auctions. The reserve price coefficients for click and revenue regressions are positive but statistically insignificant at the 5% level. The noisiness of these estimates is unsurprising given the relative rarity of clicks (fewer than 0.1% of links in the data have at least one click) as well as the targeted nature of this policy change, where only a smaller number of queries experienced significant changes in the reserve price.

Figure 6: DID Coefficients



Notes: Each row presents the DID coefficient, taking within-treatment-cell variation as exogenous. The first two outcomes contain agency by treatment-cell fixed effects and present link-level outcomes, while the remaining ones are at the hotel-query level. Standard errors are clustered at the hotel-market-device level. Data from the pilot window is excluded from each regression's data sample, as well as all queries with length-of-stay values which were not exposed to this policy change.

Appendix A2 reports results from various robustness tests under alternative specifications, including splines of $\Delta\text{reserve}_{iht}$ in continuous difference-in-difference specifications (Callaway, Goodman-Bacon, and Sant'Anna 2024). We note that even if the revenue coefficient on $\Delta\text{reserve}_{iht}$ were statistically significant, it alone cannot determine what the optimal reserve price is. A more flexible spline, plotted in Figure A2.1, shows meaningful nonlinearity but large confidence intervals due to the relative lack of high-$\Delta\text{reserve}_{iht}$ auctions despite the policy change (see Figure A2.2). While these results are suggestive, the optimal reserve price and its relationship to consumers' click decisions cannot be determined by this evidence alone.

The results from this natural experiment are intuitively consistent with strong within-

hotel nesting behavior. Figure 6 shows a statistically significant, negative effect of the reserve price on the number of bidders; this is robust to continuous DID specificiations as illustrated in Figure A2.1c. Despite this decrease in choice set variation, clicks and revenue had no statistically significant decrease, as may be predicted in the absence of strong hotel-specific preferences. This cannot be attributed to simple re-ranking by the platform; Figure A2.4 shows that the estimates remain qualitatively similar as in the main specification (though statistically insignificant for both clicks and revenue) when controlling for hotel ranks.

3.3 Evidence for nested click probabilities

We now use the policy change to conduct several empirical tests of nested consumer preferences at the hotel level. The tests we present now are reduced-form evidence of strong correlations in idiosyncratic consumer preferences for available hotel-agency combinations. This complements our additional evidence in Section 4, which uses structural click probability models to show that hotel-specific preferences are both statistically and economically significant.

We regress the number of clicks per link on agency availability and prices for all queries with at least two distinct hotels. In each regression, we control for rich query-level covariates (log-booking window; log-length of stay, and log-guest count) as well as link- and query-level fixed effects (agency, hotel, month, device, rank, hotel-rank, and number of hotels). We run these regressions both on the full sample of all queries and on a restricted sample for only queries with at least one click, since the latter may represent more serious searchers than casual browsers.

In these regressions, we instrument for agency counts per hotel and agencies' listed hotel prices. We use two sets of instruments. First, we use mean log-bid and mean log-price, computed across all auctions on the same day for the same hotel-auction pair, excluding the current observation. This captures variation in bids and prices that are common to a "market" defined by the day-hotel-agency combination. Since bids and prices are set at a level coarser than query-specific attributes, this variation is uncorrelated with query-specific demand shocks. Second, we use the mean difference between own and competitor log-bids and log-prices, computed across all auctions on the same day for the same agency, booking window, length of stay, guest count, device, and market, but different hotels. This provides another measure of shocks common to a market defined along different dimensions, but which are uncorrelated with agency-level strategic decisions. These instruments capture variation in bids and prices relative to other agencies within similar queries, and hence capture variation

in strategies that are also uncorrelated with query-level demand shocks. [3]

We first study direct price competition between agencies across various hotels. We model this competition in a reduced form, regressing clicks on own log prices as well as the minimum log-price at the same hotel (among rival agencies) and the minimum log-price across all other hotels in the same query. For hotel auctions with only one agency, the rival same-hotel log-price is coded as zero and we include an indicator for lack of same-hotel competitors. Own log-price instruments are constructed from mean log-prices within the same hotel-agency-period group, and from mean differences in own log-prices and average log-prices among all other competitors for the same hotel in the same matched agency-period group. Competitors' log-price instruments are constructed from the same comparison groups described above, using average differences between the agency's own log-price and the relevant minimum competing price (whether within the same hotel or across other hotels).

Table 1 illustrates the relationship between clicks per link and the prices of different agency-hotel options. Across all specifications, own prices are negatively related to clicks, with IV elasticities of approximately -2.7 and -3.6 for the full and has-click samples, respectively. Competing same-hotel prices are positively correlated with the number of clicks, with corresponding IV elasticities of about 2.5 and 3.5. While this sign difference is expected of differentiated product demand settings, we note that the difference in magnitudes of the coefficients for same-hotel and other-hotel options: the same-hotel price coefficients are about 10 and 70 times larger than the corresponding cross-hotel price coefficients, suggesting that within-hotel competitors are more salient for consumers' click (and presumably, booking) decisions. The cross-hotel IV coefficients are negative, possibly due to higher-quality competing hotels which cannot be adequately controlled for in this reduced-form specification. However, when conditioning on the "has click" sample this estimate is statistically insignificant.

In Table 2, we study the role of agency availability in consumers' click decisions. In addition to the two comparison-group instrument types defined above, we instrument with the per-hotel policy-induced change in reserve price and per-hotel measures of exposure to the reserve price change. The within-hotel agency-count coefficients are negative across both the full sample and a restricted sample of queries with at least one click, translating to approximate IV semi-elasticities of -0.15 and -0.11 in the full and has-click samples. In contrast, the cross-hotel agency-count coefficients are positive, though smaller in the IV specification for both samples; this suggests that agencies may participate more in auctions for hotels in high-demand regions and for queries with high expected click and booking

3. Both instrument types are similar to Hausman instruments, using narrowly-defined comparison "markets" of consumers with similar queries and results, but defined at the individual query level. Due to data limitations, we use same-week comparisons when same-day comparisons are not available, and set corresponding instruments to the sample average for cases where no daily or weekly average is available.

Table 1: Regressions of Clicks on Own and Other Prices

	Full Sample		Has Click	
Variable	OLS	Price IV	OLS	Price IV
Own log price	-5.47e-4***	-6.01e-4***	-0.0254***	-0.0348***
	(4.07e-5)	(7.27e-5)	(0.0017)	(0.0033)
Rival log price, same hotel	3.98e-4***	5.60e-4***	0.0206***	0.0338***
	(4.42e-5)	(8.19e-5)	(0.0019)	(0.0038)
Minimum log price, other hotel	3.68e-5**	-5.56e-5**	0.0029***	-4.84e-4
	(1.16e-5)	(2.16e-5)	(3.51e-4)	(8.76e-4)
Kleibergen–Paap F	–	3,725.1	–	41.4
Observations	14,783,679	14,783,679	334,380	334,380
Queries	137,310	137,310	2,680	2,680

Notes: The dependent variable is click count, where the sample restricted to multi-hotel queries ($H_q > 1$). "Has Click" restricts to queries with at least one click. The mean outcome is 2.21e-4 in the full sample and 0.0098 in the has-click sample. All specifications include log booking window, log length of stay, log guests, and an indicator for whether a link has no same-hotel competing link; partner, hotel, market, month, device, rank, and hotel-rank fixed effects are also included. Standard errors in parentheses are clustered by query. * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

probabilities. In each specification, the magnitude of the coefficient is larger for the number of agencies at the same hotel than for other hotels, as would be expected when within-hotel competition is more relevant to consumer choices than cross-hotel competition.

Table 2: Regressions of Clicks on the Number of Same- and Other-Hotel Agencies

	Full Sample		Has Click	
Variable	OLS	Count IV	OLS	Count IV
Other agencies, same hotel	-1.73e-5***	-3.27e-5**	-7.81e-4***	-0.0011*
	(4.81e-6)	(1.23e-5)	(1.78e-4)	(4.93e-4)
Agency-hotel options, other hotels	1.19e-6***	9.74e-7***	3.47e-5***	2.14e-5**
	(1.52e-7)	(2.68e-7)	(4.26e-6)	(8.03e-6)
Kleibergen–Paap F	–	3,680.3	–	131.5
Observations	14,783,679	14,783,679	334,380	334,380
Queries	137,310	137,310	2,680	2,680

Notes: The dependent variable is the click count per option-query row. The sample is restricted to multi-hotel queries ($H_q > 1$). "Has Click" restricts to queries with at least one click. The mean outcome is 2.21e-4 in the full sample and 0.0098 in the has-click sample. All specifications include own log price per night, minimum other-hotel price, log booking window, log length of stay, log guests, and partner, hotel, market, month, agency-rank, and hotel-rank fixed effects. Standard errors in parentheses are clustered by query. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4 Model

We first specify the consumer model, followed by the auction platform's rules, the agencies' resulting strategies, and the optimization problem faced by Platform.com.

4.1 Click probability model

We model consumers as viewing many potential options of agency-hotel pairs and choosing the option that maximizes their utility. We assume a consumer books at most one hotel room, but allow for more than one agency-hotel link to be clicked within a search query before making that purchase decision. The baseline model is thus a two-step random utility model: a consumer first decides whether to click a link, and if so, which to inspect, after which the consumer makes a booking decision. Under standard assumptions, this yields a tractable click and booking probability model.

In practice, however, the search process is more complicated than a simple two-step model would allow. We observe that approximately 0.6% of queries have more than one click, but due to data limitations we do not observe the order in which links are clicked. Further, some of these repeated clicks are for the same option, which is not rationalizable under standard search models but directly relevant to a cost-per-click position auction. We therefore extend the baseline model to a multi-click extension which retains this structure for the extensive click margin and for the relative attractiveness of each option, but does not attempt to solve a fully dynamic search problem with endogenous return, stopping, and recall.[4] Instead, we use a parsimonious reduced-form continuation rule for additional clicks. This is best interpreted as a myopic approximation to continued browsing: consumers make an initial click decision, without anticipating the exogenous probability with which they return to the search page and select another option.

4.1.i Utility from search and purchase stage

Each consumer i receives utility from booking hotel h through agency a at time t, and each query it returns a set of possible agency-hotel options $\mathcal{J}_{it}$. Let D_{iaht} denote all observable

4. A fully dynamic search model would require additional data, state variables, and assumptions on the timing and order of clicks, the consumer's ability to return to the results page, the information revealed by inspecting a link, and the stopping rule after each inspection. Alternatively, a sequential search model would generally involve combinatorial search sets or stronger parametric assumptions, especially with the large number of options and rich query-level heterogeneity; simulation-based methods like Moraga-González, Sándor, and Wildenbeest (2023) would require more detailed booking data and impose additional complexity that would impact estimation, particularly in the nested logit setting with a high share of individuals who do not click on a single option. Our model focuses instead on the baseline click probabilities and substitution patterns between link choices.

link-level data, with $D_{it} \equiv \{D_{iaht} : ah \in \mathcal{J}_{it}\}$ denoting the query-level dataset. This data contains endogenous variables: the hotel price p_{iaht}, the bids b_{iaht}, and the resulting agency and hotel ranks r_{iaht} and r_{iht}. All other utility-relevant data is denoted x: individual search-query characteristics are denoted x_{it}, hotel-level characteristics are denoted x_{ht}, and agency-hotel characteristics are denoted x_{aht}; collectively, non-rank and non-price observables are denoted x_{iaht}. Finally, we observe instruments z_{iaht} that do not enter utility directly but are relevant to agencies' pricing and bidding decisions.

We begin by specifying the booking-stage utility index, which is given by

$$\mathbf{u}(D_{iaht}) = \delta(x_{iaht}, p_{iaht}) + \xi_{iaht}$$

where ξ_{iaht} is an agency-hotel-time utility component unobserved by the econometrician. Conditional on having clicked link ah, the consumer chooses whether to book that hotel. The utility of booking is $\mathbf{u}(D_{iaht}) + \varepsilon_{iaht}(1)$, while the utility of not booking is $\varepsilon_{iaht}(0)$. We assume $\varepsilon_{iaht}(d)$ for $d \in \{0, 1\}$ is i.i.d. T1EV with mean zero and scale parameter σ_b,[5] which implies the booking-stage value function

$$V_b(D_{iaht}) = \sigma_b \ln\Big(1 + \exp\big(\mathbf{u}(D_{iaht})/\sigma_b\big)\Big)$$

Before having the option to book a hotel, the consumer must first choose whether to click on an agency-hotel link. Doing so incurs a search cost that depends on the hotel and agency ranks:

$$\kappa^h(x_{it}, r_{iht}) + \kappa^a(x_{it}, r_{iaht})$$

We allow these search costs to depend on individual/query-level characteristics such as the device on which search is conducted. There is also a link-specific utility shock

$$e_{iaht}(x_{it}, x_{ht}) = \nu_{iht} + \big(1 - \lambda(x_{it}, x_{ht})\big)\epsilon_{iaht}$$

which takes the nested logit form, where $\lambda(x_{it}, x_{ht}) \in [0, 1)$, ϵ_{iaht} is i.i.d. T1EV with mean zero and unit scale, and ν_{iht} is such that $e_{iaht}(x_{it}, x_{ht})$ is also T1EV mean zero with unit scale. This implies the utility from clicking link ah is

$$u_{iaht}(X_{iaht}) = \underbrace{-\kappa^h(x_{it}, r_{ht}) - \kappa^a(x_{it}, r_{aht}) + V_b(D_{iaht})}_{\mathbf{U}(D_{iaht})} + e_{iaht}(x_{it}, x_{ht})$$

5. Because δ is flexible and we do not observe booking-stage data, we normalize $\sigma_b = 1$.

with the outside option utility of not clicking any link given by $u_{i0t} = \epsilon_{i0t}$, which is also i.i.d. T1EV with mean zero and unit scale. The expected value function of any given choice at the click stage is $\mathbf{U}(D_{iaht})$.

Together, these assumptions and functional forms yield tractable expressions for the expected utility from the set of available agencies per hotel, as well as the set of all search results. The expected utility of clicking one of the links of the agencies $\mathcal{A}_{ih}$ for query i and hotel h is the hotel inclusive value

$$I_h(D_{iht}) = \ln\left(\sum_{a\in\mathcal{A}_h} \exp\left(\frac{\mathbf{U}(D_{iaht})}{1-\lambda(x_{it},x_{ht})}\right)\right)$$

while the expected utility from the set of all possible choices (where individual i sees hotels $\mathcal{H}_i$) is given by the overall inclusive value

$$I(D_{it}) = \ln\left(1+\sum_{h\in\mathcal{H}_i} \exp\Big([1-\lambda(x_{it},x_{ht})]\cdot I_h(D_{iht})\Big)\right)$$

These inclusive values measure the respective expected consumer welfare from a selection of agencies for a hotel and the overall set of options faced by a consumer.

4.1.ii Click and booking probabilities

We now combine the model components above to generate tractable functional forms for the expected click probabilities and click counts. We allow the query-hotel and agency-level utility index δ, as well as the search cost indices κ^h and κ^a, to be nonparametric. The nonparametric nesting parameter $\lambda(x_{it}, x_{ht})$ determines how correlated idiosyncratic utility shocks are for partners that offer booking services for hotel h at time t. We collectively denote all of the possibly-time-varying utility parameters by $\theta_t = \{\delta_t, \kappa_t^h, \kappa_t^a, \lambda_t, \sigma_{b,t}\}$.

First, we specify the booking probability. Under the previous T1EV assumptions, the booking probability of a particular option ah conditional on clicking that link is

$$\begin{aligned} q_{iaht}(D_{iaht};\theta_t) &\equiv \mathbb{P}\big[i \text{ books via } ah \,\big|\, i \text{ clicked } ah,\ D_{iaht}\big] \\ &= \frac{\exp\big(\mathbf{u}(D_{iaht})/\sigma_{b,t}\big)}{1+\exp\big(\mathbf{u}(D_{iaht})/\sigma_{b,t}\big)} \end{aligned}$$

Importantly, clicking on a link does not necessarily imply booking that reservation, and the booking probability does not rely on the booking-stage utility shock.

Next, we write the conditional click probabilities at any stage in the search process.

Again by the maintained assumptions on various utility shocks, the click shares are the nested logit choice probabilities of a single click:

$$q_{iaht}(D_{it};\theta_t) = \frac{\exp\Big(\mathbf{U}(D_{iaht})/[1-\lambda(x_{it},x_{ht})]\Big)}{\exp\Big(I(D_{it}) + \lambda(x_{it},x_{ht})\cdot I_h(D_{it})\Big)}$$

We aggregate these link-specific click probabilities to the probability that at least one "inside" option is clicked, denoted by

$$Q_{it}(D_{it};\theta_t) = \sum_{ah\in\mathcal{J}_{it}} q_{iaht}(D_{it};\theta_t)$$

with "inside" shares denoted by

$$s_{iaht}(D_{it};\theta_t) = \frac{q_{iaht}(D_{it};\theta_t)}{Q_{it}(D_{it};\theta_t)}$$

Aside from their semiparametric definitions, these absolute and relative choice probabilities are standard from the nested logit choice literature.

We now extend this click probability model to a setting with potentially multiple clicks. We assume that the first consideration of a query is free, so that all observations in our dataset correspond to a consumer who chooses whether or not to click on any offered link. Subsequent opportunities to further inspect the options are exogenous and follow a Poisson distribution with mean μ. This myopic search is a reduced-form approximation to the consumer search process, which is more tractable in a setting where consumers face different numbers of options with every possible query. We assume that any "follow-up" searches must focus on the inside options, or the actual links listed on the page. By the Poisson thinning property, the number of additional clicks made to option ah is Poisson with mean $\mu s_{aht}(D_{it};\theta_t)$. Letting C_{iaht} denote the number of clicks for option ah and query it, we can write the expected number of clicks for ah as

$$\begin{aligned}\mathbb{E}[C_{iaht} \mid D_{it}] &= Q_{it}(D_{it};\theta_t)\cdot[s_{iaht}(D_{it};\theta_t) + \mu\cdot s_{iaht}(D_{it};\theta_t)] \\ &= (1+\mu)\cdot q_{iaht}(D_{it};\theta_t)\end{aligned}$$

Summing across all possible options yields the expectation of the total number of clicks C_{it}, or

$$\mathbb{E}[C_{it} \mid D_{it}] = (1+\mu)\, Q_{it}(D_{it};\theta_t)$$

Since μ is constant across queries, the system of relative click probabilities captures which links become more or less attractive when prices, ranks, or observables change. These components can be combined to construct the query-level likelihood contribution, which is derived in Appendix A3. As a consumer welfare measure, we use the overall inclusive value of a consumer who clicks on at most one link, $I(D_{it})$, which applies to over 99% of all consumers in our data.

4.2 Auction rules and bidding

We first describe the details of the auction rules in this setting, followed by a discussion of bidder behavior as it relates to our analysis.

4.2.i Auction rules

This setting differs from typical position auctions by blending a generalized first price auction (GFPA) and generalized second price auction (GSPA). For convenience, we first define the ranking function ρ which translates higher values into a lower (better) ranking. Formally, for a length-A vector of values v, ranks are given by

$$\rho(v) = 1 + \sum_{a=1}^{A} \mathbf{1}[v < v_a]$$

Platform.com computes agency a's payment as a function of both their bid[6] and their relative rank. For any auction (omitting dependence on iht), a set of bids b, and ranks r, the expected cost per click (ECPC) is

$$C_a(b \mid r, m) = \begin{cases} b_a & \text{if non-legacy, } \rho_a(b) > r_a, \text{ and } b_a \geq m \\ \max\left\{\{b_j \mid b_j < b_a\}, m\right\} & \text{if non-legacy, } \rho_a(b) \leq r_a, \text{ and } b_a \geq m \\ 0 & \text{if non-legacy and } b_a < m \\ b_a & \text{if legacy} \end{cases} \tag{2}$$

This ECPC is equal to the invoice when agency a receives a click in that particular auction. Non-legacy agencies with bids below the minimum m are not visible to any consumer. Note that when Platform.com's rankings r align exactly with the raw bid rankings, this collapses to a GSPA.

6. Legacy deals such as cost-per-acquisition, which remain in place for agencies that have not yet switched over to the bidding mechanism, use an estimated bid equivalent as computed by Platform.com.

We now explain the ranking algorithm in these auctions, which depends on both a quality score and the expected payment. The quality score is proprietary and therefore unobserved to customers, agencies, and us; we model it as a function of observable data. In many applications (e.g. Athey and Nekipelov (2010)) this score is modeled as a click-through rate. In our setting, both the click-through rate and the expected payment depend on rankings; the identity of more prominently-ranked agencies (not just the number) can meaningfully affect click shares. To make the score evaluation non-circular we instead use a nonparametric function $\mathcal{Q}$ for the quality score and evaluate the ECPC conditional on each firm being ranked in the first position.[7] Thus, the vector of rankings r is formally determined by

$$r = \rho\Big(C(b \mid 1, m) \cdot \mathcal{Q}(x_{iaht}, p_{iaht}) \cdot \eta_{iht}\Big) \tag{3}$$

where all perturbations η_{iht} to the quality score are i.i.d. from some known distribution F_η; these can arise due to unobservable variation and measurement error in β.

4.2.ii Agency strategies

We do not explicitly model the bidding game in this setting. Instead, we treat bidding functions as having an unknown (to consumers, Platform.com, and us) functional form $\mathcal{B}_{at}(x_{iaht}, z_{iaht}, m_{iht})$, where z_{iaht} contains instruments relevant to agencies' strategic choices but which do not directly impact consumer utility. Similarly, price functions have an unknown structure $\mathcal{P}_{at}(x_{iaht}, z_{iaht}, m_{iht})$. We allow for significant heterogeneity and experimentation by individual agencies at the query-characteristic level.

We take this reduced-form approach for several reasons. First, no existing theoretical model (to our knowledge) studies the equilibrium of this setting's blend between generalized first and second price auctions. We leave this equilibrium solution to future work. Second, even existing theoretical models for the GFPA and GSPA make strong assumptions on the structure and information sets of this game: most importantly, that (average) bids are known in an *ex post* equilibrium. The assumptions of such models are not likely to hold in this setting where bidders are "onboarding" to the auction system and experimenting often. Third, the "dual-price" nature of this auction means there are two margins of adjustment by the agencies: both "upward" (bids paid to the platform) and "downward" (prices charged to consumers). Neither the optimal adjustment margin in a full-information theoretical model, nor the one which will be chosen in practice by rationally-inattentive and incompletely in-

7. This is analogous to the standard quality-adjusted bid ranking, but incorporating the blend between the GFPA and GSPA mechanisms as used in this setting.

formed managers, is obvious *ex ante*. Fourth, conversations with employees of Platform.com suggest that agencies' behavior is guided by relatively simple algorithms, if they employ algorithmic decision-making at all.

With this in mind, we model the agencies' log-price and log-bid functions as the sum of an unknown (to us) strategy function and additive, mean-zero shocks. More formally, latent prices and bids are determined by

$$\begin{aligned} \ln(p^*_{iaht}) &= \mathcal{P}_{at}(x_{iaht}, z_{iaht}, m_{iht}) + \zeta^p_{iaht} \\ \ln(b^*_{iaht}) &= \mathcal{B}_{at}(x_{iaht}, z_{iaht}, m_{iht}) + \zeta^b_{iaht} \end{aligned} \tag{4}$$

The functions $\mathcal{P}$ and $\mathcal{B}$ depend on both key strategic variables and query-level characteristics. We allow both ζ^p_{iaht} and ζ^b_{iaht} to be correlated with the unobserved demand shock ξ_{iaht}, as well as with each other. We assume ζ's distribution is i.i.d. $\mathcal{N}(0, \Sigma_d)$ for different deal types d, either active CPC bidding or legacy. The reserve price rule creates a selection problem through which actively bidding agencies (and their realized prices p_{iaht} and bids b_{iaht}) are only observed when their latent bid b^*_{iaht} meets the reserve price m_{iht}.

In the following sections, we assume these bid and price strategy functions are invariant to changes in the reserve price, and compute counterfactuals assuming bids and prices follow the estimated policy in equation (4). Using the distribution of historical bids is conventional, following Haile and Tamer (2003), Aradillas-López, Gandhi, and Quint (2013), and Coey et al. (2021) among others. Critically, we differ from prior work by including the reserve price m_{iht} as an argument in the bid and price functions rather, allowing us to use quasi-experimental variation in reserve prices to identify agency strategies. As demonstrated in the case of a generalized English auction with known click-through rates (Edelman and Schwarz 2010, section II.A), a change in the reserve price increases the total payment of all remaining bidders. Though our auction mechanism is different, we allow this behavior by including the reserve price m_{iht} as an argument in both the bid and price functions, despite it not being directly observable to bidders. We justify this inclusion by the fact that bidders may participate in millions of auctions per day and receive detailed daily reports at various query characteristic levels (display rank, booking window, etc.), from which they can quickly make accurate inferences about the reserve price.

5 Estimation

We first describe the bid, price, and click model estimation, followed by the implied optimal reserve prices under the various specifications and click models.

5.1 Agency bid and price model

We begin by estimating the agency strategies in equation (4). We use Hausman-style instruments for both bids and prices, from the same "leave hotel-auction out" agency-hotel-period and agency-period comparison groups as in 3.3 above. Here, we use average log bids and prices within the agency-hotel-period, and average differences from competitors' mean values within the agency-period comparison. Periods are days where available, with a weekly fallback where there is no daily comparison group, while any observations without relevant comparison groups have the corresponding IV set directly to the sample average. Each of these four instruments are further interacted with active-bidder status, allowing different strategic choices across deal types. We control for log-reserve price for both bidders and non-bidders, and include fixed effects for each agency, market, booking month, device, day-of-week, and page type. For ease of computation, we use linear specifications for $\mathcal{P}$ and $\mathcal{B}$ to estimate and recover residuals and hotel random effects from the pricing and bidding equations.

Selected coefficients from the first-stage bid and price model are reported in Table 3. While log bids are positively associated with the log reserve price, the total elasticity is small for bidders at only 0.079. This coefficient is similar to the difference-in-differences estimate from Figure 6, despite accounting for the participation margin; a modest but positive change is consistent with theoretical results in related auctions in Edelman and Schwarz (2010). In contrast, but in keeping with Figure 6, log prices hardly vary with the log reserve price. Given agencies' advertising on multiple travel search sites and other costs in providing a room booking, it is reasonable to conclude that reserve prices are not a major factor in determining nightly hotel prices. Bids and ECPCs (for non-bidders) vary similarly with other key query characteristics, with a smaller relationship for log-length of stay among active bidders vs. legacy deals, while prices only vary significantly with the log-number of guests.

5.2 Consumer click model

We next turn our attention to the click probability model. We model the utility function for each individual as depending on average price-per-night, length-of-stay, booking window, country from which the search originated, agency identifiers, and median price-per-night of the hotel (as a proxy for quality). Search costs for hotels and agencies depend on the respective rankings among hotels and agencies for that hotel, as well as the device type and page display type (see Figure 2).

We also estimate two alternative types of click probability models to compare with the

Table 3: Click-IV first-stage, selected parameters

	ln(bid (ECPC))	ln(price per night)
Coefficients		
Active bidder × ln(reserve)	0.079	0.015
	(0.004)	(0.002)
ln(length of stay)	0.257	0.028
	(0.003)	(0.003)
Active bidder × ln(length of stay)	-0.080	-0.014
	(0.004)	(0.002)
ln(number of guests)	0.095	0.355
	(0.017)	(0.022)
Active bidder × ln(number of guests)	-0.010	0.019
	(0.017)	(0.007)
ln(booking window)	-0.078	0.009
	(0.002)	(0.003)
Active bidder × ln(booking window)	0.005	0.003
	(0.002)	(0.001)
Variance and covariance estimates		
Hotel-effect variance	0.061	0.245
	(0.004)	(0.012)
Hotel-effect covariance	0.082	–
	(0.004)	
Residual-shock variance: Active bidder	0.138	0.114
	(0.002)	(0.003)
Residual-shock covariance: Active bidder	0.016	–
	(0.000)	
Residual-shock variance: legacy	0.175	0.116
	(0.007)	(0.004)
Residual-shock covariance: legacy	0.058	–
	(0.003)	
Observations	9,156,782	
Queries	98,655	
Hotels	3,020	

Notes: Coefficients are from linear specifications for $\mathcal{B}$ and $\mathcal{P}$, where standard errors clustered by search query and hotel-market-device are in parentheses. Other coefficients, such as IVs, are omitted from this table.

nested logit model: a simple logit and a row-wise binary logit. The simple logit is a special class of the nested logit model; it restricts the nesting parameter λ to be equal to zero, so that idiosyncratic preferences for hotel-specific options are independent. The row-wise binary logit model does not have a satisfying microfoundation for this setting, as it treats the click choice problem as completely independent for each alternative. We include it as a benchmark since it most closely approximates popular data science approaches, as well as the standard "no competitive externalities" model used in much of the related position auction literature.[8]

The purchase utility and search cost functions are non-parameteric, and we select the final model from five different model specifications using holdout validation. The first specification is a simple linear model for each utility and cost component. The remaining four specifications are neural networks (NNs) with either one or two hidden layers (each with a smaller and larger number of nodes). The NNs are each warm-started by fitting each model component to the corresponding linear estimate and batch-training until either a 10% holdout testing loss fails to improve for 10 epochs or the maximum number of training epochs is reached. The final model is chosen from its respective class (binary, simple, or nested logit) based on the log-likelihood value from an additional 10% holdout validation sample, after which the selected model is retrained on the full sample.

Since both prices and bids are endogenous in this setting, we control for the correlation of prices and ranks (which depend on bids) with unobserved demand determinants ξ_{iaht}. We first construct the control function $\mathbb{E}[\xi_{iaht} \mid x_{iaht}, z_{iaht}]$, following Petrin and Train (2010) to write

$$\mathbb{E}[\xi_{iaht} \mid x_{iaht}, z_{iaht}] = \chi_p \zeta^p_{iaht} + \chi_b \zeta^b_{iaht}$$

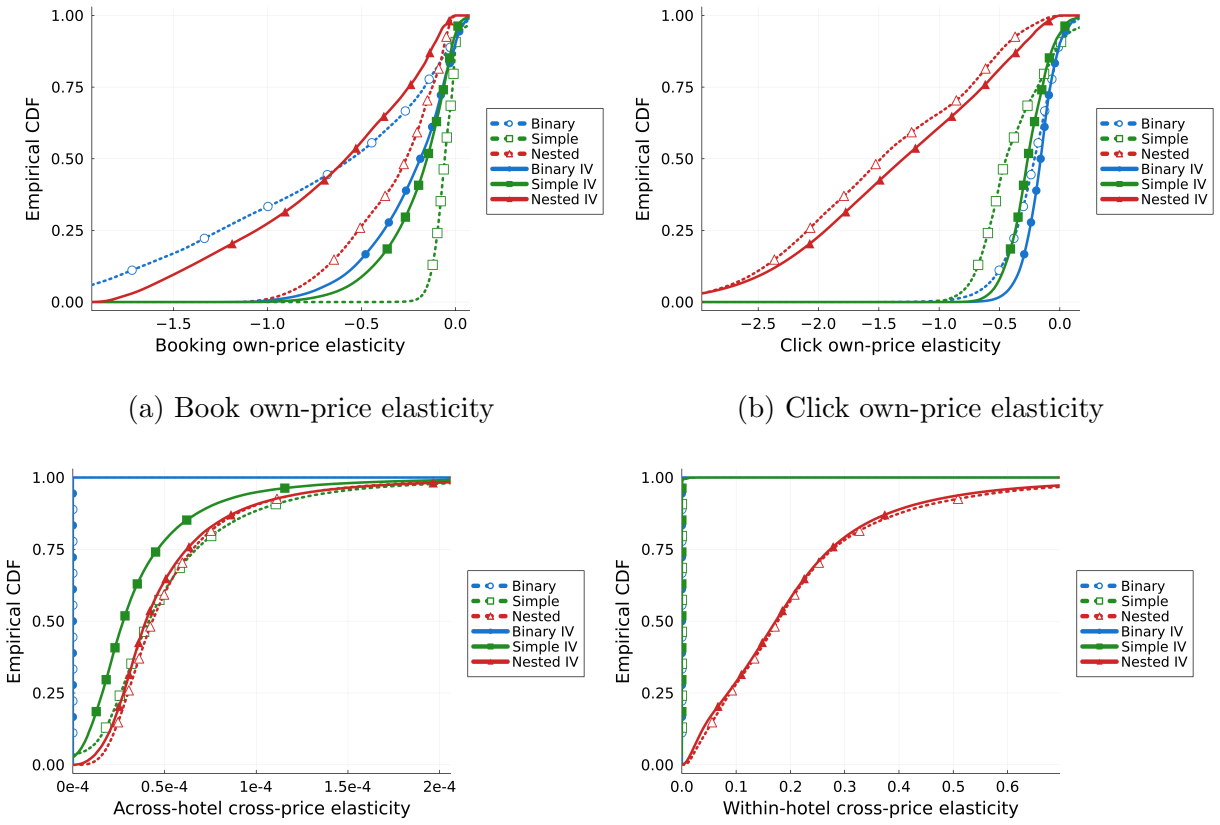
where the predicted residuals $\hat{\zeta}$ come from the agency strategy functions estimated in the previous section. We also estimate and report results for click models in which prices and bids (and the resulting ranks) are treated as exogenous.

Figure 7 plots the CDF of estimated elasticities across each model type. Panel (a) presents own-price elasticities for the booking stage as implied by the click model, though no booking data was used in estimation. Consumers appear relatively inelastic, though this conditions on the set of consumers who have already clicked on the option in question. Own-price elasticities for clicks, as presented in panel (b), show a strong divergence between the nested logit models and both the binary and simple logit models, which suggest that con-

8. When computing first-click consumer welfare for this model, we use sum over the binary choice inclusive values for each of the available links, or $\sum_{a \in \mathcal{A}_h, h \in \mathcal{H}_i} \ln(1 + \exp(\mathbf{U}(D_{iaht}))$.

sumers are highly price inelastic. Cross-price elasticities in panels (c) and (d) are intuitive: the binary logit displays exactly zero cross-price elasticity, while the cross-hotel elasticities for simple and nested logit models are small but nonzero. In contrast, within-hotel elasticities for the nested logit models are large, compared to the near-zero estimates from the simple logit models. This is driven by the strong estimated nesting patterns: the nesting parameter estimate $\hat{\lambda}$ is approximately 0.79 and 0.83 in the nested logit and nested logit IV specifications, respectively.

Figure 7: Elasticity CDFs Across Models and Links



(a) Book own-price elasticity

(b) Click own-price elasticity

(c) Cross-price elasticity (outside hotel nest)

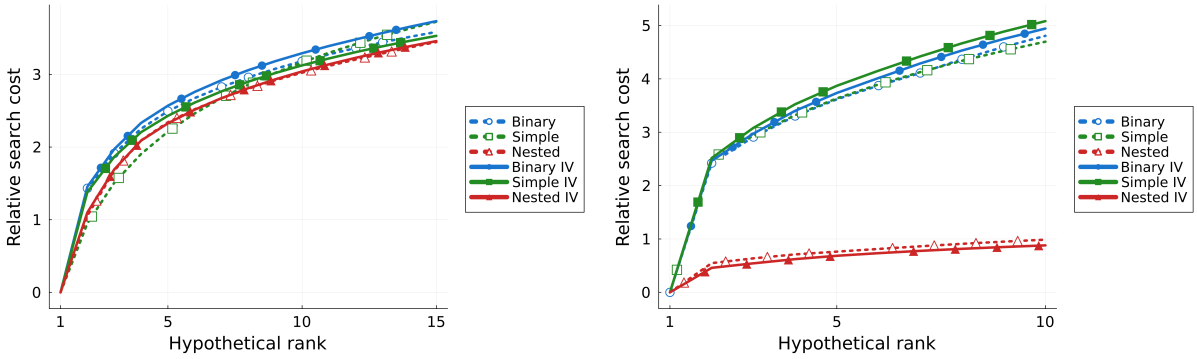
(d) Cross-price elasticity (inside hotel nest)

Notes: Each panel plots the empirical cumulative distribution function (CDF) of link-level price elasticities under each estimated model, evaluated on observed prices and other covariates. Cross-price elasticities for the binary model variants are mechanically zero throughout. All estimated elasticities use point estimates and are not orthogonalized.

The various click models also have important implications for the estimated search costs in Figure 8. Each model largely agrees on role of hotel ranks in search costs in panel (a): lower-ranked hotels are significantly more costly to examine and therefore less likely to be clicked on in a search query. Agency search costs, however, vary significantly between the nested logit and other models. Panel (b) shows that both the binary and simple logit models

estimate agency rank to be roughly as impactful as hotel rank. The nested logit models imply a monotonic but much flatter search cost curve: even the 10th-place agency is cheaper to inspect in a nested logit setting than the 2nd-place agency where consumer search does not exhibit such nesting behavior.

Figure 8: Search Costs Across Ranks



(a) Hotel (normalized) (b) Agency

Notes: Both panels present average search costs as a function of hypothetical rank. Panel (a) plots the hotel search cost component with the rank-1 cost normalized to zero for display purposes, while panel (b) plots the agency search cost component where the rank-1 cost is restricted to be zero. All estimated search costs use point estimates and are not orthogonalized.

5.3 Quality scores for agency rankings

In order to estimate revenue under counterfactual adjustments to the reserve price, we also estimate the quality score $\hat{S}$ in the ranking function (3). We assume the individual-agency quality score perturbations η_{iaht} are i.i.d. T1EV with mean zero and scale $1/\sigma_S > 0$, and estimate the per-agency log-score as

$$S_{iaht} = \ln(C(b \mid 1, m)) + \ln(\mathcal{Q}(x_{iaht}, p_{iaht})) + \eta_{iaht}$$

The T1EV assumption on η_{iaht} implies the score scale σ_S and quality index $\mathcal{Q}$ can be estimated via an exploded logit likelihood. We represent this quality index using a neural network with two hidden layers of 32 nodes each. Since the true score function is proprietary, we report only quality measures. Our neural network estimate for $\mathcal{Q}$ performs well: the Kendall rank correlation coefficient is 0.83, implying that our model places 91.5% of agencies in the correct rank order. Of these, the first agency is correctly ranked in the first position in 76.7% of all auctions.

6 Counterfactual reserve prices

We next explore how the optimal reserve price varies with the different consumer models, as well as under agency-specific and rank-specific reserve prices. Throughout, we allow bidders' prices and bids to adjust to counterfactual reserve prices as estimated in 5.1, with associated re-ranking under Platform.com's estimated quality scores, which we hold fixed.

6.1 Optimal reserve price

Given bids b_{iht} and a reserve price m_{iht}, we can calculate rankings r_{iht}, the resulting click probability q_{iht}, and the expected invoice $C(b_{iht} \mid r_{iht}, m_{iht})$ for all agencies in the auction. Although bids and prices may depend on the reserve price, as specified in equation (4), we assume this relationship is invariant to changes in the reserve price policy. Further, we keep the invoice function in Equation (2) and the ranking rule in Equation (3) fixed.

We now consider the problem of maximizing platform revenue in a one-shot game. The revenue from query it corresponding to reserve price m_{iht} is

$$R_{it}(m_{iht}) = \sum_{a \in A} C_a(b_{iht} \mid r_{iht}, m_{iht}) \cdot q_{ah}(D_{iaht}; \theta_t)$$

Our objective is to select a reserve price rule M_γ to maximize revenue, where $m_{iht} = M_\gamma(x_{iht})$ and $\gamma \in \Gamma$ ranges over a fixed class of policy functions mapping auction-level variables to a reserve price. The unknown, optimal rule is given by

$$\gamma^* = \arg\max_{\gamma \in \Gamma} \quad \mathbb{E}[R_{it}(M_\gamma(x_{iht}))]$$

where the expectation is taken over the distribution of query characteristics. Since both the true click probabilities q_{aht} and the exact distribution of covariates x_{iht} are unknown, directly implementing M_{γ^*} is infeasible. We optimize over a class Γ of low-dimensional representations instead, and select the reserve price function that yields the highest estimated revenue.

We evaluate three specification families for the reserve price:

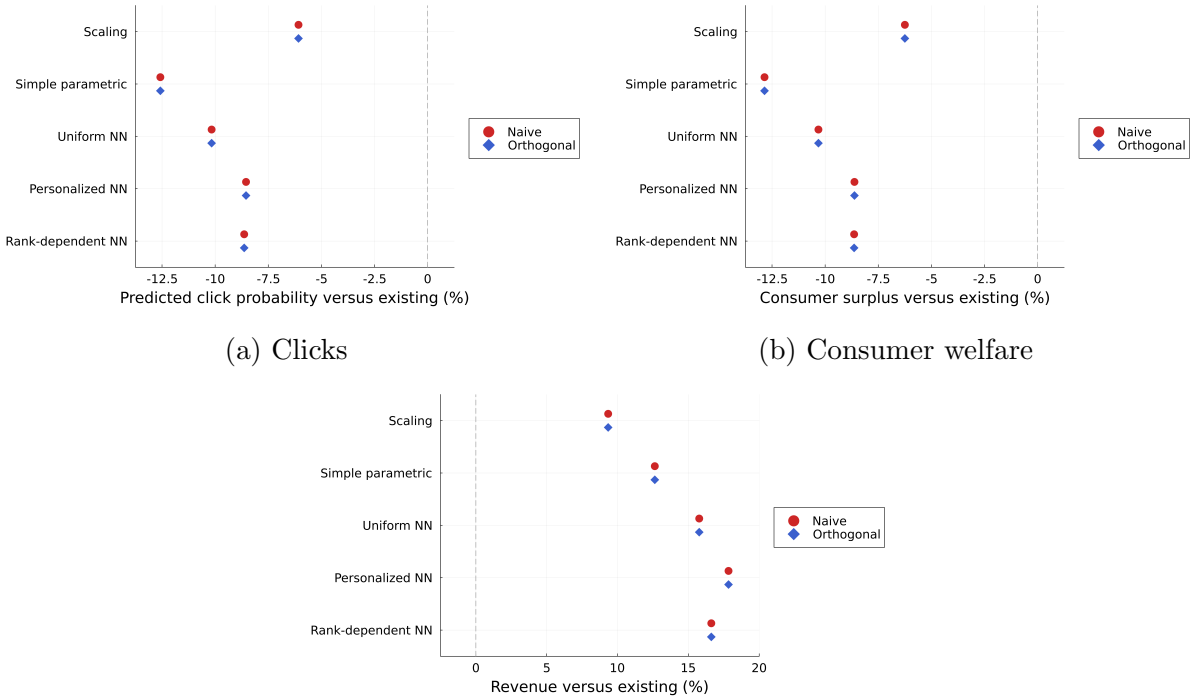
1. **Scaling.** We multiply the existing reserve price formula used by Platform.com by a constant.
2. **Parametric.** We specify the reserve price as a parametric function of auction-level characteristics using Platform.com's proprietary formula, but optimize the coefficients.
3. **Neural network.** We use both 2- and 3-hidden-layer multilayer perceptrons (each layer with 16 nodes).

Finally, we compare uniform-within-hotel reserve prices with those which are personalized,

with identities for the most frequently appearing 25 partners in the data, as well as rank-dependent reserve prices for different positions for each hotel.

Figure 9 summarizes the estimated outcomes from each of the specification families, personalized, and rank-dependent reserve prices, under the maintained nested logit IV model. The naive point estimate for each counterfactual is presented alongside its orthogonalized estimate and corresponding 95% confidence interval. Each counterfactual reserve price estimates a significant revenue increase over the baseline, with the neural network (NN) specifications yielding the highest predicted revenue. Similarly, The following sections examine each in turn.

Figure 9: Outcomes across from reserve price specifications



(a) Clicks

(b) Consumer welfare

(c) Revenue

Notes: All outcomes are evaluated using the nested logit IV model. Naive point estimates are shown in red, above the corresponding orthogonalized estimates in blue (*orthogonalized estimates and their confidence intervals are in progress*).

6.2 Reserve price scaling

We first analyze a simple counterfactual – if the existing reserve price were simply multiplied by a scalar, what is the optimal scalar? This allows us to retain query- and hotel-specific heterogeneity in the reserve price as currently targeted by Platform.com while providing a

transparent and informative demonstration of the importance of substitution patterns in optimizing online auctions.

Figure 10 shows the counterfactual percent changes in clicks and revenue per consumer when scaling the existing reserve price. As shown in panel (a), increasing the reserve price and thereby excluding bidders has a negative effect on overall clicks, though the projected decrease is significantly smaller in the nested logit models because consumers substitute strongly toward within-hotel links by competing agencies. In panel (b), the predicted consumer welfare changes are similar to the corresponding click changes; a reduced choice set makes consumers worse off, though less so when preferences are more strongly driven by hotel availability (under the nested logit models) than idiosyncratic hotel-agency preferences.

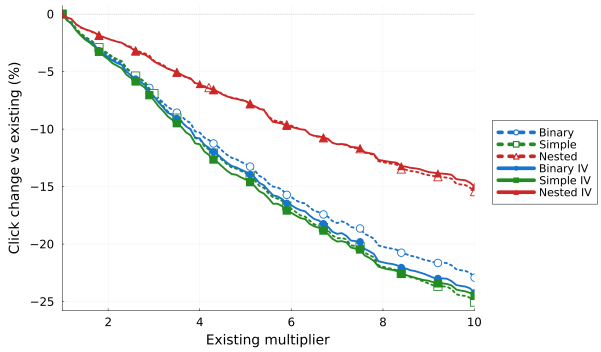
Panel (c) shows the implied changes in revenue under each model: both the binary logit and simple logit models imply that the optimal reserve price used by Platform.com is two to three times the current value, when restricting to scaled current-reserve price counterfactuals. In contrast, the nested logit model implies that the optimal scaled reserve price is roughly four to five times the current value. This is because the nested logit model reveals high substitution across agencies within hotels, so increasing the reserve price directs consumers to high-bidding, high-revenue agencies while only modestly decreasing clicks. Interestingly, Appendix Figure A5.2 suggests that when failing to account for endogenous bid adjustment, the current reserve price appears roughly optimal under either the binary or simple logit models.

6.3 Optimal uniform-within-hotel reserve prices

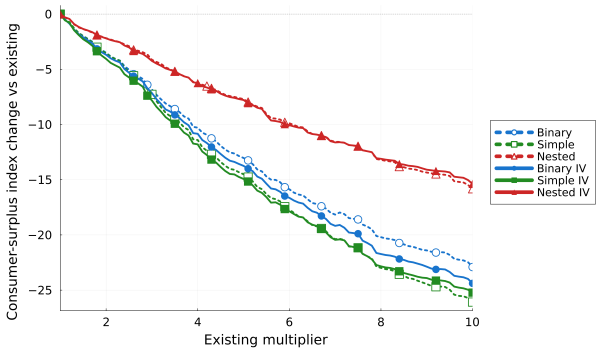
We next compute the optimal per-hotel reserve price in this setting, using two distinct functions of query- and hotel-level data. The first uses Platform.com's proprietary formula and optimizes the coefficients on various auction characteristics. The second is a neural network, which allows for more flexible optimization with respect to the same observed data. In practice, both yield broadly similar results, though the neural network achieves a higher estimated revenue at its optimum at 15.8% increase over the current reserve price's revenue. The parametric specification yields estimated revenue gains from 12.6% at its simplest (as shown in Figure A5.1) to 14.8% with richer covariates. In what follows, we use the optimal uniform neural net reserve price.

Figure 11 plots the counterfactual density of agencies per auction and bidders per auction under the existing and optimal reserve price. There are still a large number of agencies per auction in many cases (panel (a)), though this is mostly attributable to legacy deals for whom the reserve price has no effect. Panel (b) shows that most auctions under

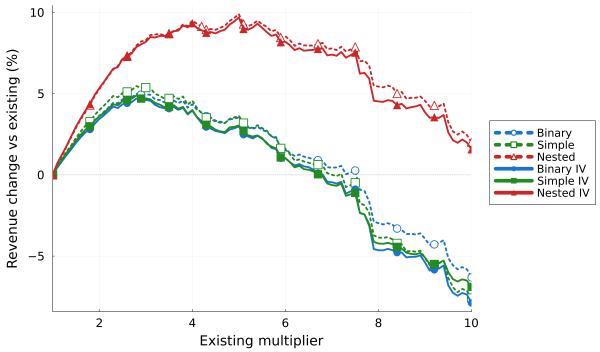
Figure 10: Counterfactual outcomes under scaled reserve price



(a) Clicks



(b) Consumer welfare

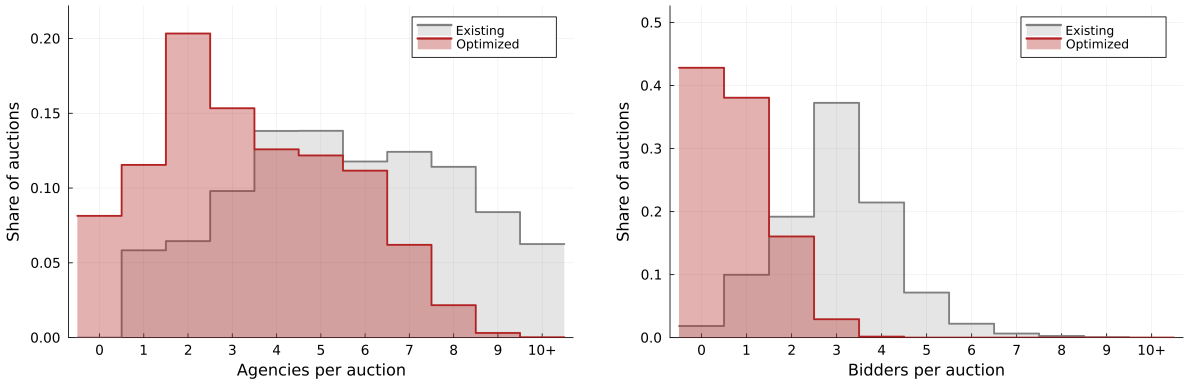


(c) Revenue

Notes: Panels present estimated clicks, consumer welfare, and revenue when multiplying all existing (post-policy-change) reserve prices by a single constant, under each of the estimated model variants. All estimated outcomes for these scaled-reserve price plots use point estimates and are not orthogonalized.

the counterfactual reserve price have at most 3, and often 2 or fewer, bidding agencies. The temporary mechanism split between auctions and legacy deals limits the potential effect of the reserve price policy, since a large number of links are still displayed regardless of the reserve price. However, the revenue-optimal reserve price results in a significant portion of hotels (approximately 8%) that are omitted from search results due to insufficiently high bids by all bidding partners. This omission is a major component of the counterfactual consumer welfare loss, since the large nesting parameter implies the existence of a hotel is relatively more important than the availability of agency links for that hotel. Thus, a soft reserve price or similar "safeguarding" mechanism for the top bidder may be able to mitigate the consumer welfare and click decreases.

Figure 11: Agency and bidder participation under optimal uniform reserve price



(a) All agencies

(b) Bidders only

Notes: Panels are constructed using the nested logit IV model. Estimated agency and bidder counts use point estimates and are not orthogonalized.

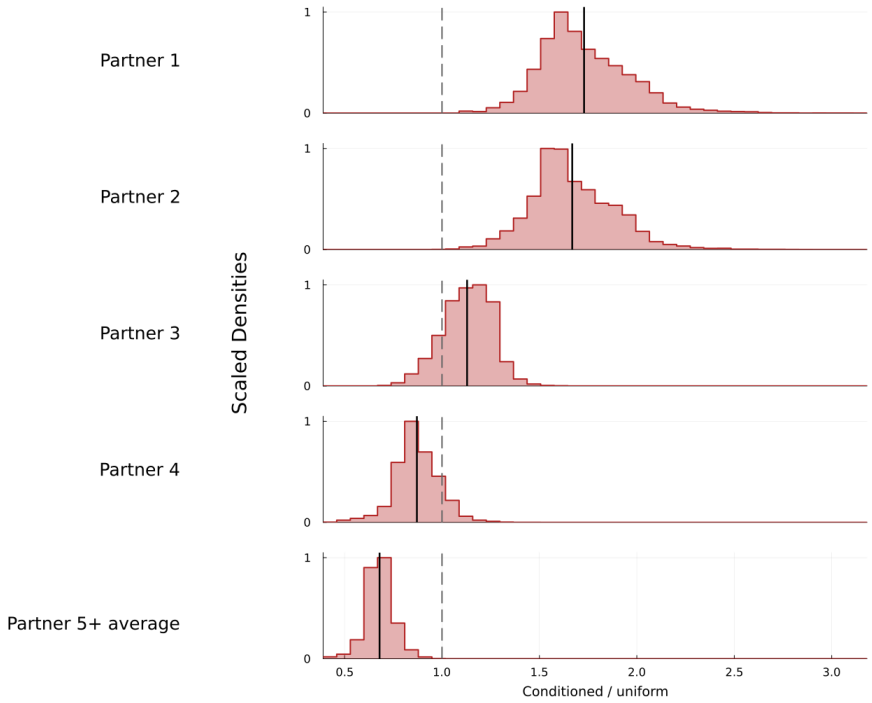
6.4 Personalized reserve prices

Up to this point, we have restricted all reserve prices to be constant across competing bidders for a single hotel. We now relax this to examine the effects of implementing personalized reserve prices across different agencies. Other work has studied personalized reserve prices, notably Ostrovsky and Schwarz (2023) in a multi-position search auction setting. This personalization often comes in the form of a quality-score adjustment, where individual ads face reserve prices determined by their estimated quality or click probability. In contrast, we consider different reserve price functions for individual bidders.

Figure 12 displays the average value of the personalized reserve prices for each of the 4 highest-bidding partners, along with average reserve prices for the remaining lower-bidding partners, normalized as a percentage of the highest-charged reserve price under the counter-

factual. The personalized reserve prices reveal substantial variation in the reserves charged across different partners, but this yields only a modest revenue gain: 17.8% increase relative to the current reserve price, vs the 15.8% increase of the uniform reserve price.

Figure 12: Average personalized reserve price as a percent of highest-charged bidder's reserve price

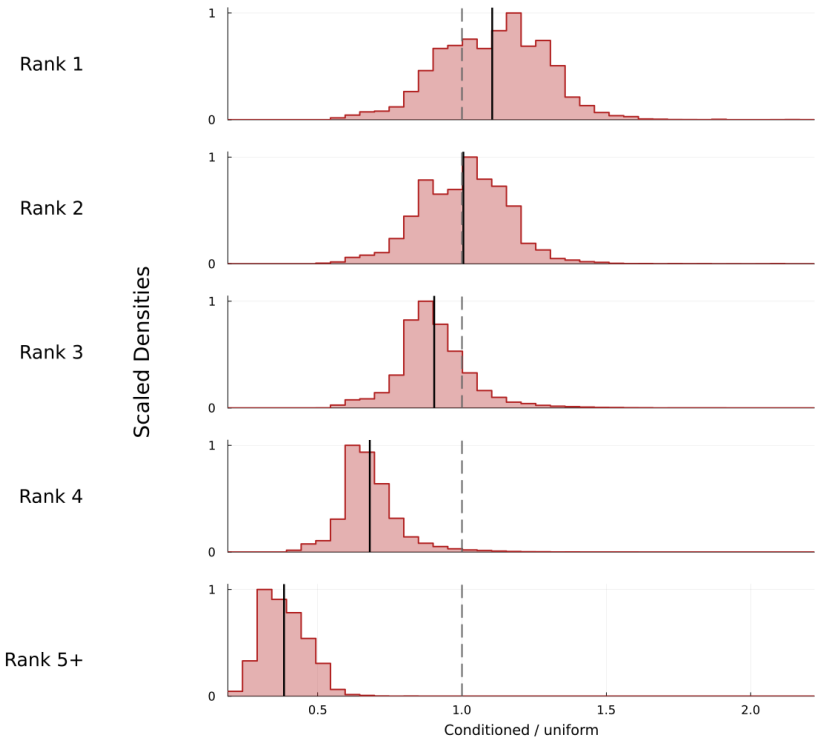


Notes: Panels are constructed using the nested logit IV model. Anonymized partners are sorted by their average implied reserve price for each auction (whether they participated or not in that hotel auction). The last row is the average among all partners whose average reserve prices are not in the top four.

6.5 Rank-dependent reserve prices

Instead of assigning different reserve prices directly to different bidders, we consider rank-dependent reserve prices as studied theoretically in Athey and Ellison (2011). Figure 13 plots the average reserve price applied at each possible agency rank, divided by the corresponding uniform reserve price. More attractive ranks face higher reserve prices, though the first three ranks generally are assigned reserve prices similar to the uniform case. Lower-ranked reserve prices are considerably lower, especially for ranks 5 and worse, where the average reserve is only a quarter of that which is implied by the uniform policy. The resulting revenue gain is very close to that of the uniform reserve price, with a 16.6% revenue gain, despite the relatively large discounting for lower-ranked positions.

Figure 13: Average rank-dependent reserve price as a percent of rank-1 reserve price



Notes: Panels are constructed using the nested logit IV model. Optimal rank-dependent reserves are only calculated for the first four positions, while all lower agency positions are assigned the same reserve price at the same hotel-auction level.

7 Conclusion

This paper studies the design of position auctions when consumers have flexible preferences over the displayed auctions. We illustrate via a simple theoretical model how strong correlation between consumers' preferences for advertised products can increase the revenue-optimal reserve price. We then provide reduced-form evidence of such correlation from a natural experiment on a travel search platform, in which increased reserve prices caused more aggressive bidding, fewer listings per auction, and no statistically significant negative impact on either clicks or platform revenues. We develop and estimate a nested logit model of consumer search, and use this model to evaluate counterfactual reserve prices for these auctions, including with personalized and rank-dependent reserve prices.

Our estimates are economically meaningful: flexibly modeling consumer preferences over advertised products implies significant revenue gains over what is implied by simpler models. Counterfactual reserve prices can increase revenue by 15.8%, with marginal improvements for personalized and rank-dependent reserve prices. This becomes especially important when designing other features of an auction mechanism, such as the pricing rule or optimal weights when adding quality scores. We also note that the counterfactual reserve prices we consider are inherently limited by the mechanism on Platform.com, which combines legacy ad deals with an auction setting, so the "pure auction" effects of optimal reserve prices may be far larger. We leave a more in-depth analysis of multi-mechanism design, and how transitioning between mechanisms affects users' strategies and outcomes, to future work.

The results we present here show that click probabilities in ad auctions, far from being innocuous modeling choices, can have important implications for consumer behavior and platform incentives. Recent regulation of choice screen auctions must appropriately account for both click and bid responses when the set of displayed options changes. Further, the rise of algorithmic bidding systems on ad platforms, such as Google (Decarolis and Rovigatti 2021) and Meta (Tadelis et al. 2023), may increase bidder sophistication and with it the importance of careful auction modeling and design.

References

Aggarwal, Gagan, Jon Feldman, Shanmugavelayutham Muthukrishnan, and Martin Pál. 2008. "Sponsored search auctions with markovian users." In *International Workshop on Internet and Network Economics,* 621–628. Springer.

Aradillas-López, Andrés, Amit Gandhi, and Daniel Quint. 2013. "Identification and inference in ascending auctions with correlated private values." *Econometrica* 81 (2): 489–534.

Athey, Susan, and Glenn Ellison. 2011. "Position auctions with consumer search." *The Quarterly Journal of Economics* 126 (3): 1213–1270.

Athey, Susan, and Denis Nekipelov. 2010. "A structural model of sponsored search advertising auctions." In *Sixth ad auctions workshop,* 15:5.

Bergemann, Dirk, Alessandro Bonatti, and Nicholas Wu. 2025. "Bidding with Budgets: Data-Driven Bid Algorithms in Digital Advertising." *International Journal of Industrial Organization,* 103172.

Berry, Steven T. 1994. "Estimating discrete-choice models of product differentiation." *The Rand journal of economics,* 242–262.

Berry, Steven T, James A Levinsohn, and Ariel Pakes. 1993. "Automobile prices in market equilibrium: Part i and ii."

Betancourt, Jose M, Ali Hortaçsu, Aniko Oery, and Kevin R Williams. 2022. *Dynamic price competition: Theory and evidence from airline markets.* Technical report. National Bureau of Economic Research.

Börgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. 2013. "Equilibrium bids in sponsored search auctions: Theory and evidence." *American economic Journal: microeconomics* 5 (4): 163–187.

Bresnahan, Timothy F, Scott Stern, and Manuel Trajtenberg. 1996. "Market segmentation and the sources of rents from innovation: Personal computers in the late 1980's."

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2024. *Difference-in-differences with a continuous treatment.* Technical report. National Bureau of Economic Research.

Cavallo, Ruggiero, and Christopher A Wilkens. 2014. "GSP with general independent click-through-rates." In *International Conference on Web and Internet Economics,* 400–416. Springer.

Chen, Yongmin, and Chuan He. 2011. "Paid placement: Advertising and search on the internet." *The Economic Journal* 121 (556): F309–F328.

Choi, Hana, and Carl F Mela. 2019. "Monetizing online marketplaces." *Marketing Science* 38 (6): 948–972.

Coey, Dominic, Bradley J Larsen, Kane Sweeney, and Caio Waisman. 2021. "Scalable optimal online auctions." *Marketing Science* 40 (4): 593–618.

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.

Debreu, Gerard. 1960. *Individual choice behavior: A theoretical analysis.*

Decarolis, Francesco, Muxin Li, and Filippo Paternollo. 2025. "Competition and defaults in online search." *American Economic Journal: Microeconomics* 17 (3): 369–414.

Decarolis, Francesco, and Gabriele Rovigatti. 2021. "From mad men to maths men: Concentration and buyer power in online advertising." *American Economic Review* 111 (10): 3299–3327.

Dimitrellos, Panagiotis. 2022. *Essays on Sponsored Search Auctions and Online Platforms.* University of Maryland, College Park.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords." *American economic review* 97 (1): 242–259.

Edelman, Benjamin, and Michael Schwarz. 2010. "Optimal auction design and equilibrium selection in sponsored search auctions." *American Economic Review* 100 (2): 597–602.

Fotakis, Dimitris, Piotr Krysta, and Orestis Telelis. 2011. "Externalities among advertisers in sponsored search." In *International Symposium on Algorithmic Game Theory,* 105–116. Springer.

Ghosh, Arpita, and Mohammad Mahdian. 2008. "Externalities in online advertising." In *Proceedings of the 17th international conference on World Wide Web,* 161–168.

Giotis, Ioannis, and Anna R Karlin. 2008. "On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities." In *International workshop on internet and network economics,* 629–638. Springer.

Gomes, Renato, Nicole Immorlica, and Evangelos Markakis. 2009. "Externalities in keyword auctions: An empirical and theoretical assessment." In *International workshop on internet and network economics,* 172–183. Springer.

Haile, Philip A, and Elie Tamer. 2003. "Inference with an incomplete model of English auctions." *Journal of Political Economy* 111 (1): 1–51.

Hsieh, Yu-Wei, Matthew Shum, and Sha Yang. 2015. "To score or not to score? Estimates of a sponsored search auction model." *Estimates of a Sponsored Search Auction Model (February 1, 2015). USC-INET Research Paper,* nos. 15-09.

Hummel, Patrick. 2016. "Position auctions with dynamic resizing." *International journal of industrial organization* 45:38–46.

Hummel, Patrick, and R Preston McAfee. 2014. "Position auctions with externalities." In *International Conference on Web and Internet Economics,* 417–422. Springer.

Interactive Advertising Bureau and PricewaterhouseCoopers. 2025. *IAB/PwC Internet Advertising Revenue Report: Full Year 2024.* Industry report. Prepared by PricewaterhouseCoopers on behalf of the IAB. New York, NY: Interactive Advertising Bureau, April. Accessed April 17, 2025. https://www.iab.com/wp-content/uploads/2025/04/IAB_PwC-Internet-Ad-Revenue-Report-Full-Year-2024.pdf.

———. 2026. *IAB/PwC Internet Advertising Revenue Report: Full Year 2025.* Industry report. Prepared by PricewaterhouseCoopers on behalf of the IAB. New York, NY: Interactive Advertising Bureau, April. Accessed May 14, 2026. https://www.iab.com/insights/internet-advertising-revenue-report-full-year-2025/.

Kempe, David, and Mohammad Mahdian. 2008. "A cascade model for externalities in sponsored search." In *International Workshop on Internet and Network Economics,* 585–596. Springer.

Kim, Sunmin. 2025. "Airline Industry Dynamics 1993–2022." *Available at SSRN 5615551.*

Kuminov, Danny, and Moshe Tennenholtz. 2009. "User modeling in position auctions: reconsidering the GSP and VCG mechanisms." In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1,* 273–280.

McFadden, Daniel. 1977. "Modelling the choice of residential location."

———. 1980. "Econometric models for probabilistic choice among products." *Journal of Business,* S13–S29.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest. 2023. *A framework for the estimation of demand for differentiated products with simultaneous consumer search.* Technical report. Tinbergen Institute Discussion Paper.

Myerson, Roger B. 1981. "Optimal auction design." *Mathematics of operations research* 6 (1): 58–73.

Ostrovsky, Michael. 2021. "Choice screen auctions." In *Proceedings of the 22nd ACM Conference on Economics and Computation,* 741–742.

Ostrovsky, Michael, and Michael Schwarz. 2023. "Reserve prices in internet advertising auctions: A field experiment." *Journal of Political Economy* 131 (12): 3352–3376.

Ostrovsky, Michael, and Andrzej Skrzypacz. 2022. "Pure-Strategy Equilibrium in the Generalized First-Price Auction."

Petrin, Amil, and Kenneth Train. 2010. "A control function approach to endogeneity in consumer choice models." *Journal of marketing research* 47 (1): 3–13.

Sun, Yang, Yunhong Zhou, and Xiaotie Deng. 2014. "Optimal reserve prices in weighted GSP auctions." *Electronic Commerce Research and Applications* 13 (3): 178–187.

Tadelis, Steven, Christopher Hooton, Utsav Manjeer, Daniel Deisenroth, Nils Wernerfelt, Nick Dadson, and Lindsay Greenbaum. 2023. *Learning, sophistication, and the returns to advertising: Implications for differences in firm performance.* Technical report. National Bureau of Economic Research.

Topinsky, Valery. 2014. "Reserve Prices in Asymmetric Auctions." PhD diss., Ph. D. dissertation, Higher School of Economics, Moscow. Available at http ...

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.

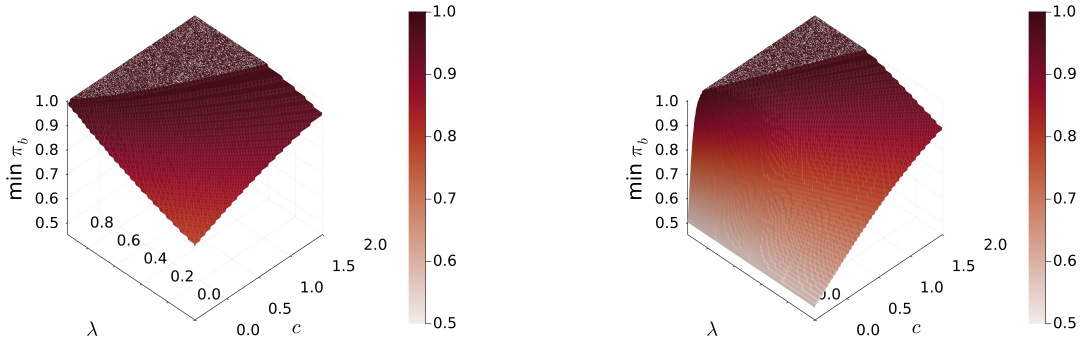
Varian, Hal R. 2007. "Position auctions." *international Journal of industrial Organization* 25 (6): 1163–1178.

Main Appendix

A1 Additional plots describing the numerical example

Figure A1.1 shows the minimum value of π_b for which the platform finds it optimal to set $m^* = \pi_b$, thereby displaying both buses to searching consumers, for different values of λ and c. This minimum is weakly lower for the constant probability model than for the nested logit model.

Figure A1.1: Minimum π_b for which it is profitable to include both firms



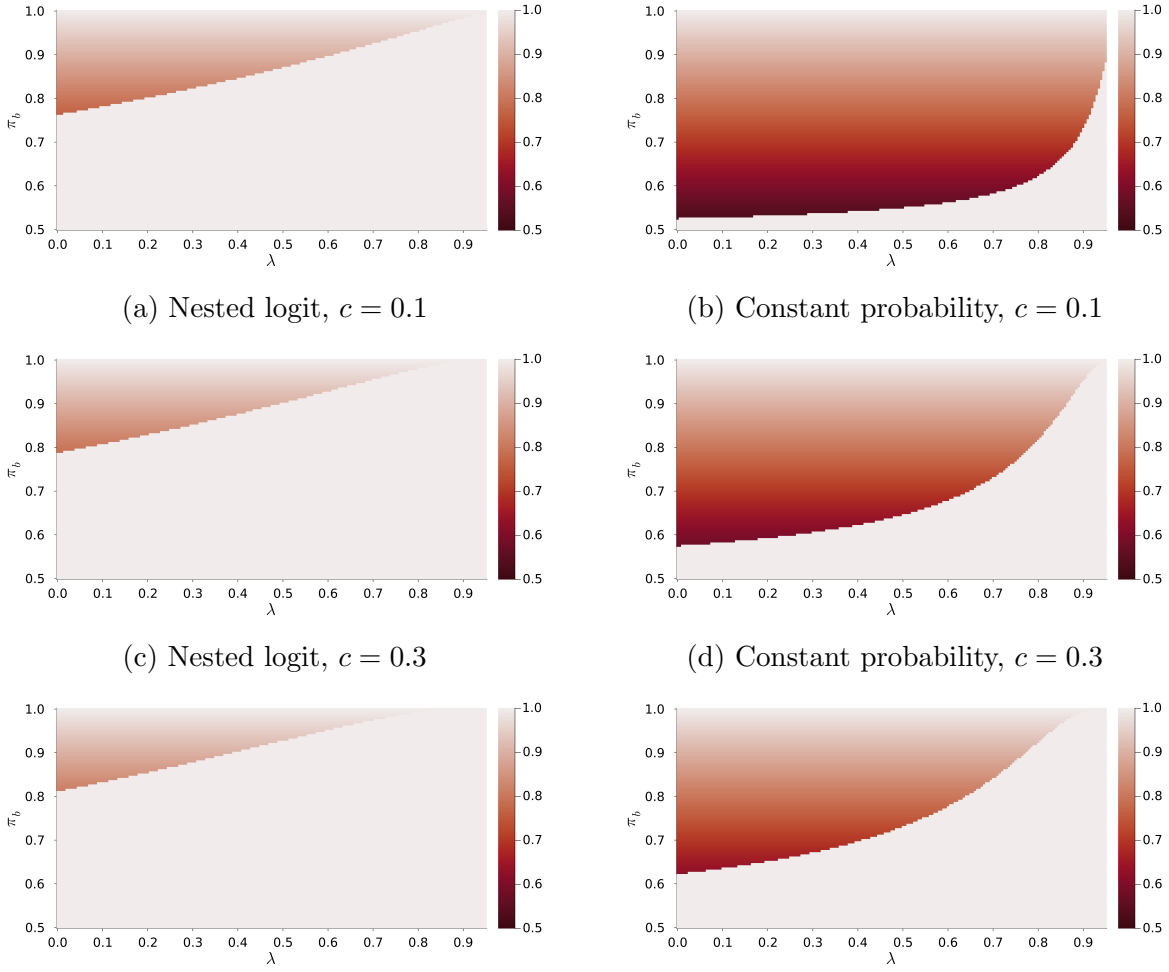
(a) Nested logit

(b) Constant probability

Notes: These plots show the minimum blue bus profit for which it is profitable to set a reserve price that includes both buses, as a function of the reserve price for different values of λ and c. The constant probability plot represents the alternative choice system where the red bus choice share does not vary with the presence of the blue bus.

Figure A1.2 also visualizes the platform problem, though for fixed values of c. Increasing c decreases the range of parameter values for which it is profitable to include both firms via a lower reserve price (in this plot, corresponding to darker shades of red). This set is larger for both lower costs c and for the constant probability model, though in the latter case this boundary is not linear in the nesting parameter λ. Higher within-nest substitution implies weakly higher reserve prices, since the gain through higher per-click payments exceeds the loss due to omitting one link.

Figure A1.2: Optimal reserve price given blue bus profitability and nesting parameter



(a) Nested logit, $c = 0.1$

(b) Constant probability, $c = 0.1$

(c) Nested logit, $c = 0.3$

(d) Constant probability, $c = 0.3$

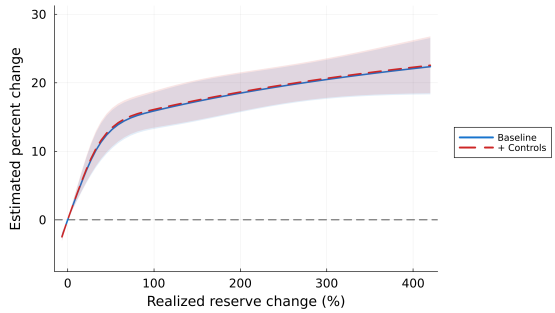
(e) Nested logit, $c = 0.5$

(f) Constant probability, $c = 0.5$

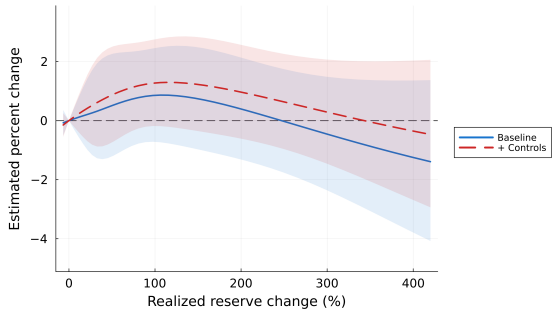
Notes: These plots show the minimum blue bus profit for which it is profitable to set a reserve price that includes both buses, as a function of the reserve price for different values of λ and c. The constant probability plot represents the alternative choice system where the red bus choice share does not vary with the presence of the blue bus.

A2 Reduced form robustness and alternative specifications

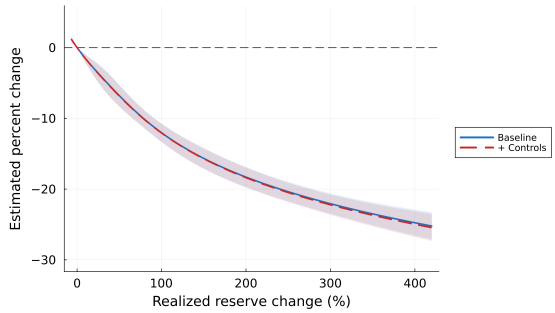
Figure A2.1: Spline estimates of continuous DID regressions, per change in reserve



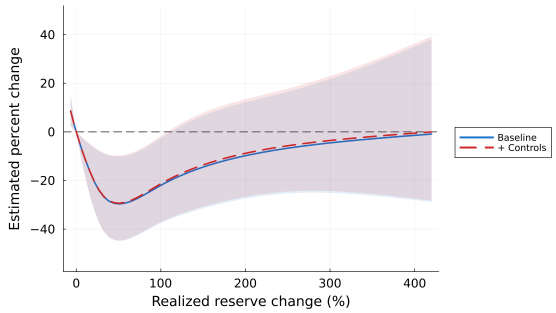
(a) Bid



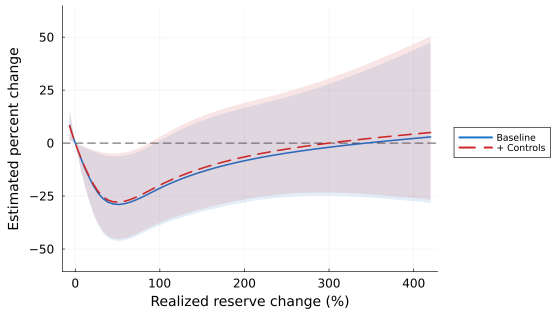
(b) Price



(c) Bidders



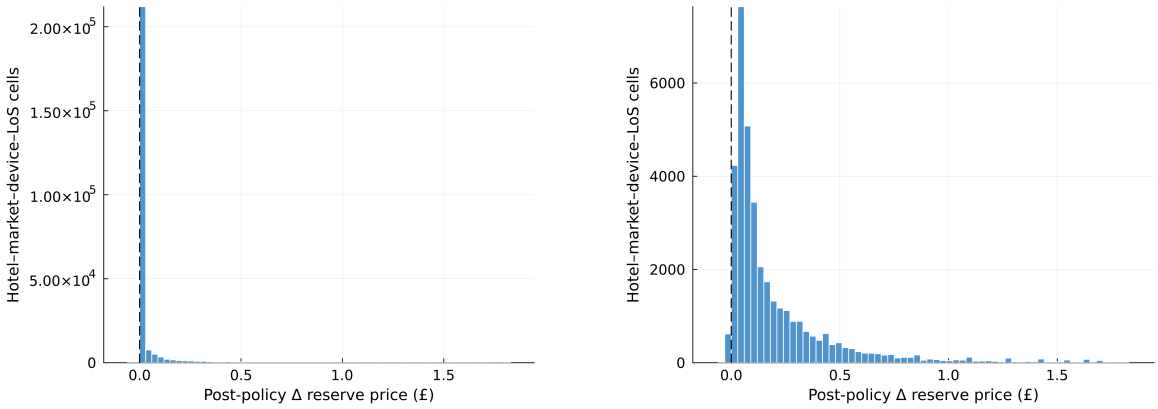
(d) Clicks



(e) Revenue

Notes: Each panel plots estimated spline coefficients for $\Delta\text{reserve}_{iht}$ as a function of the realized change in reserve. Each regression follows equation (1) except by replacing the linear DID term with a 4-knot cubic spline, using knots at the 5th, 35th, 65th, and 95th percentiles of the positive post-treatment $\Delta\text{reserve}_{iht}$ distribution, where each treatment unit is equally weighted. Shaded ribbons show pointwise 95% confidence intervals, using both the baseline specification and the specification with controls.

Figure A2.2: Histogram of post-treatment changes in reserve

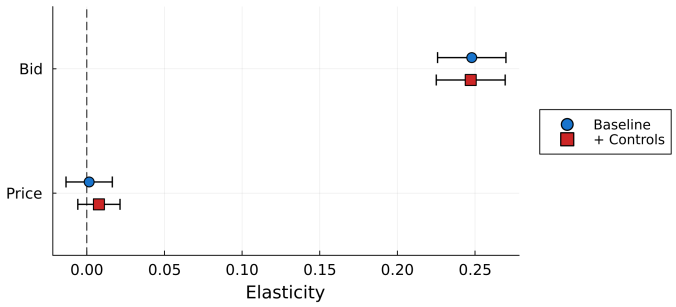


(a) All reserve changes

(b) Nonzero reserve changes

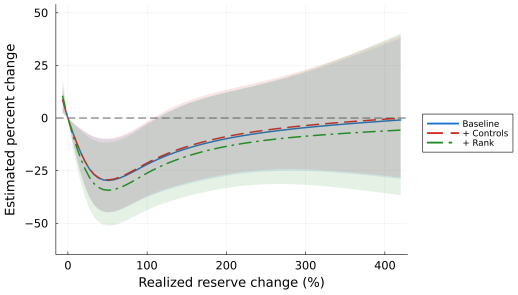
Notes: Both panels display the histogram of reserve changes under the September 2025 policy change for all hotel-query combinations (panel (a)) and for all nonzero reserve changes (panel (b)). Reserve changes are the observed reserve minus the reserve as computed under the previous policy.

Figure A2.3: DID coefficients for bids and prices without agency fixed effects

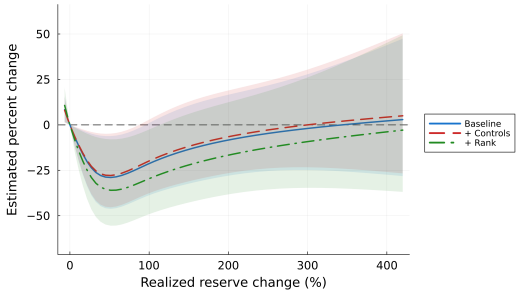


Notes: The regressions for these coefficients follow the structure in equation (1) except with row-level regressions. Fixed effects do not include agency identities; the resulting bid coefficients are approximately three times the size as when regression fixed effects include agency labels. This different specification represents a composition shift in partner participation per auction under the increased reserves.

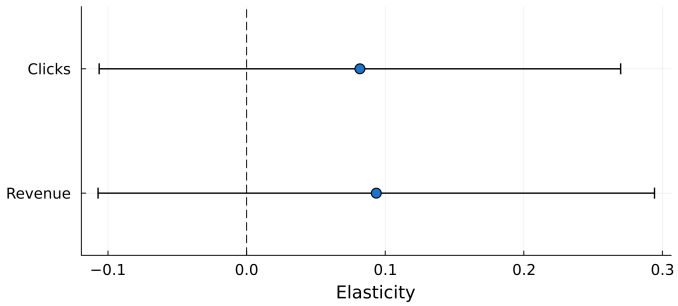
Figure A2.4: Comparison of DID for bids and prices with hotel rank fixed effects



(a) Clicks (splines)



(b) Revenue (splines)



(c) Linear DID

Notes: The regressions for these coefficients follow the structure in equation (1) but using splines for panels (a) and (b), and adding rank fixed effects for panel (c). The estimates are qualitatively similar but not statistically significant after controlling for the position rank of the hotel.

A3 Likelihood derivation

No-click consumers. The likelihood contribution of any no-click consumer is simply the probability that there are no clicks from query it. Thus the contribution to the likelihood for a no-click consumer is

$$\ell^{nc}(\theta_t \mid D_{it}) = 1 - Q_{it}(D_{it}; \theta_t)$$

Consumers with at least one click. With probability $Q_{it}(D_{it}; \theta_t)$, the consumer has at least one click. For the first click, agency-hotel combination ah is selected with probability $s_{iaht}(D_{it}; \theta_t)$. Conditional on this first click occurring, the number of additional "follow-up" clicks is $N_{it} \sim \text{Poisson}(\mu)$. Each subsequent click is i.i.d. from the same distribution, implying the total number of clicks is $C_{it} = 1 + N_{it}$. Thus, the probability mass function of the number of clicks conditional on $C_{it} \geq 1$ is

$$\mathbb{P}[C_{it} \mid C_{it} \geq 1] = \mathbb{P}[N_{it} = C_{it} - 1] = \frac{\exp(-\mu)\mu^{C_{it}-1}}{(C_{it}-1)!}$$

Since all "inside option" clicks are drawn from the same distribution (determined by $s_{iaht}(D_{it}; \theta_t)$), the probability of the observed clicks is given by the multinomial distribution

$$\mathbb{P}[\{C_{iaht}\} \mid C_{it},\ C_{it} \geq 1] = \frac{C_{it}!}{\prod_{ah \in \mathcal{J}_{it}} C_{iaht}!} \prod_{ah \in \mathcal{J}_{it}} s_{aht}(D_{it}; \theta_t)^{C_{iaht}}$$

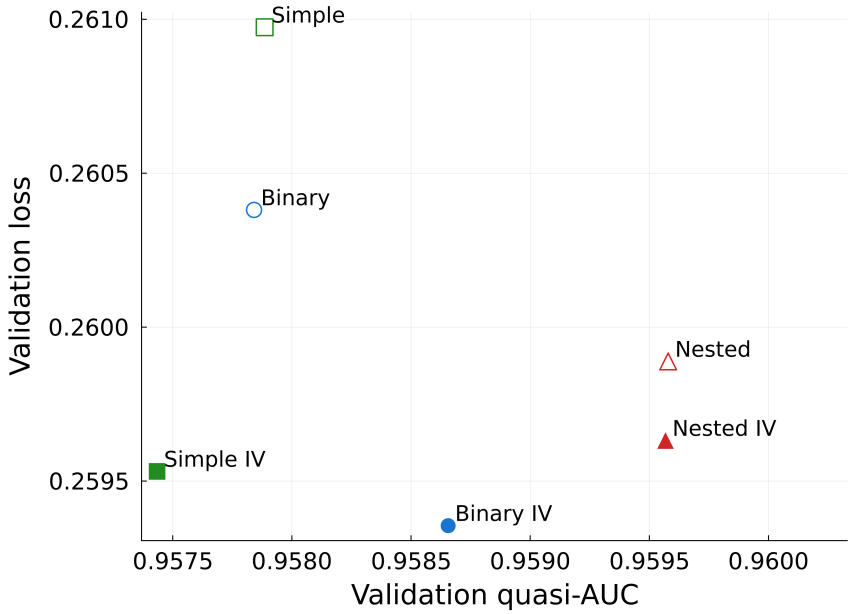
We combine these to obtain the likelihood contribution for any consumer with a click:

$$\ell^{c}(\theta_t \mid D_{it}) = Q_{it}(D_{it}; \theta_t) \cdot \frac{\exp(-\mu)\mu^{C_{it}-1}}{(C_{it}-1)!} \cdot \frac{C_{it}!}{\prod_{ah \in \mathcal{J}_{it}} C_{iaht}!} \prod_{ah \in \mathcal{J}_{it}} s_{aht}(D_{it}; \theta_t)^{C_{iaht}}$$

On the extensive margin (i.e., for any individual with $C_{it} \leq 1$) this choice problem and the substitution patterns are familiar, with the "single-click" likelihood contribution reducing to simply $q_{aht}(D_{it}; \theta_t) \cdot \exp(-\mu)$. As μ limits toward 0, this approaches the standard nested logit likelihood.

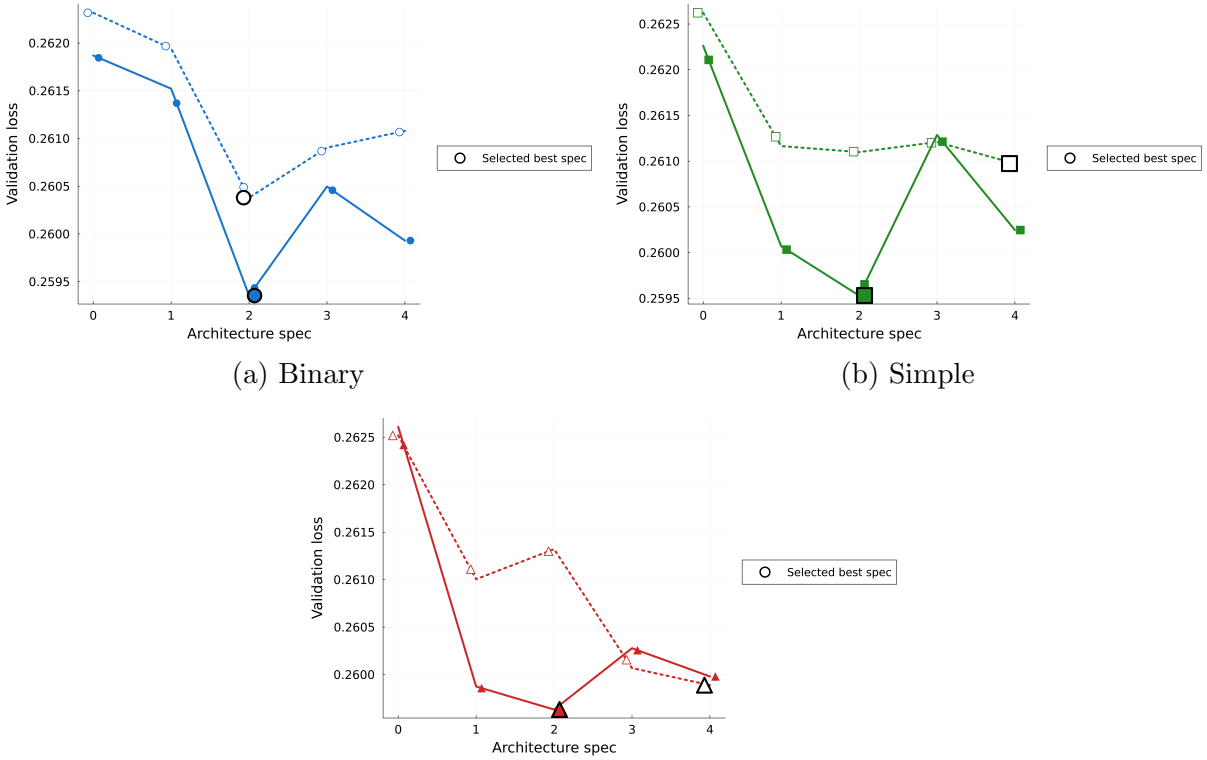
A4 Model selection details

Figure A4.1: Validation loss vs. quasi-AUC for selected logit models



Notes: The validation loss is the negative log-likelihood from the click model stage, omitting the first-stage likelihood for the IV models. The quasi-AUC (area under curve) is the fraction of rows which are distinctly ordered by clicks (e.g., 1 or 2 clicks vs. 0, or 2 clicks vs. 1) where the realized click ordering aligns with the estimated click probability; if all clicks are binary this is the standard AUC measure. Points in the lower-right corner are better under both measures.

Figure A4.2: Model selection summary across each logit model



(a) Binary

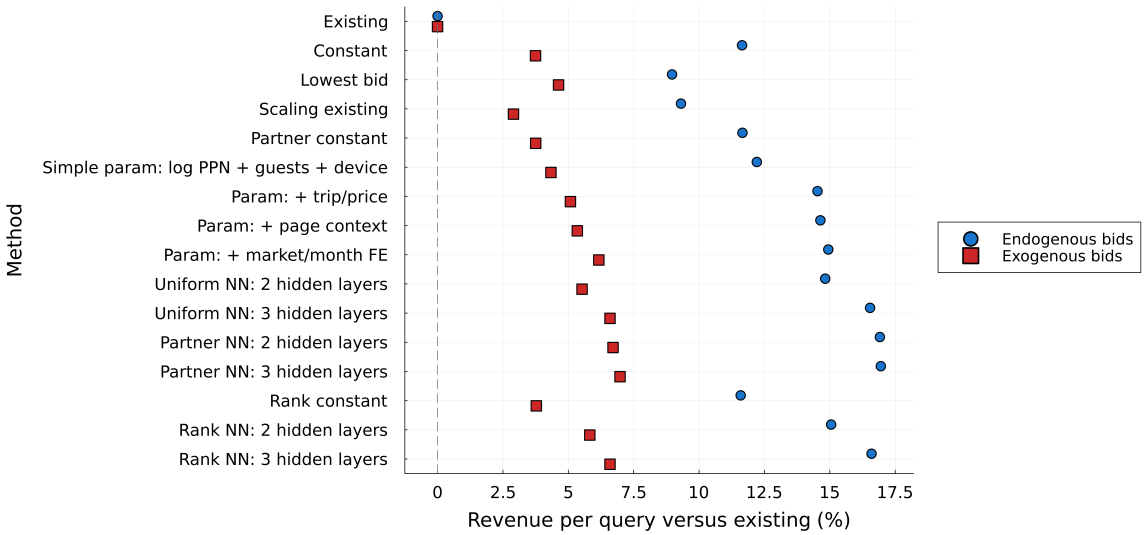
(b) Simple

(c) Nested

Notes: Panels display validation loss in the exogenous-price-and-ranking model (dashed lines) as well as the endogenous/IV model (solid lines). Larger points represent the minimum validation loss across each of the five specifications, which is model 2 for all but the simple IV and nested non-IV models. The model with the lowest holdout validation loss is retrained on the full dataset and used in counterfactuals. Specification 0 models each utility or cost component as linear, while other architecture specifications use neural networks; each of the latter uses the spec 0 model as a warm-starting feature.

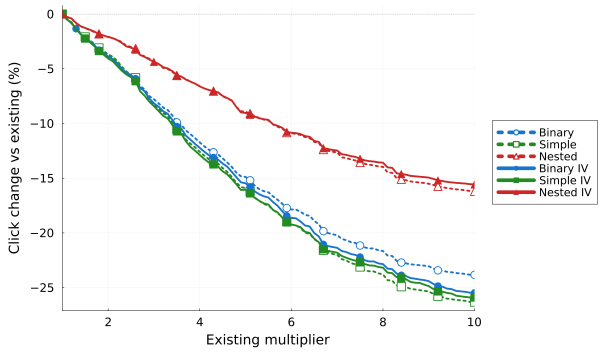
A5 Robustness: structural model and counterfactuals

Figure A5.1: Estimated revenue under reserve specifications

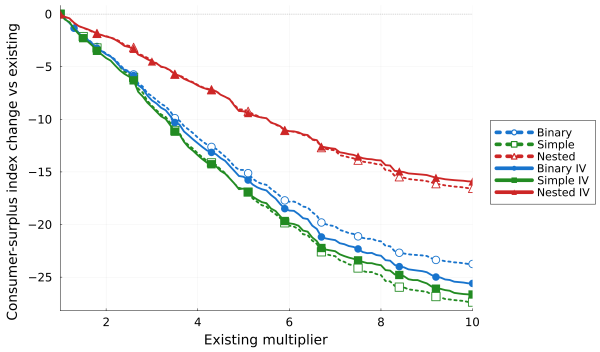


Notes: Percent changes in revenue per query are displayed in red for exogenous bids and blue under endogenous bids, across each tested parameterization. All estimates are point estimates, without orthogonalization or associated confidence intervals.

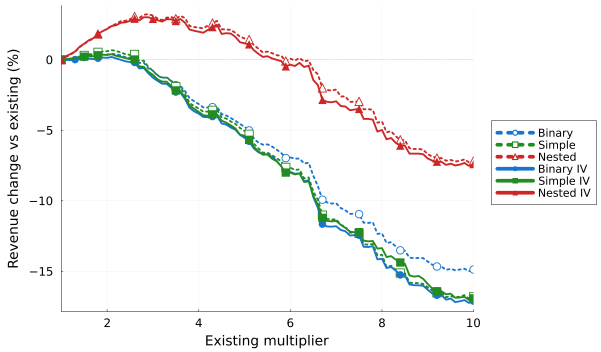
Figure A5.2: Counterfactual outcomes under scaled reserve



(a) Clicks



(b) Consumer welfare



(c) Revenue

Notes: Panels present estimated clicks, consumer welfare, and revenue when multiplying all existing (post-policy-change) reserve prices by a single constant, under each of the estimated model variants. All estimated outcomes for these scaled-reserve plots use point estimates and are not orthogonalized.