Position Auctions with Organic Search

Mehtab Hanzroh
Qidi Hu

December 24, 2025

Abstract

Recent regulatory actions, such as the FTC v. Amazon antitrust case, have raised concerns about the impact of sponsored advertisements on consumer welfare in online search platforms. While theoretical models of position auctions typically predict that sellers are ranked by consumer-match/seller-quality in equilibrium, these models often abstract from the coexistence of sponsored and organic listings. We develop a model in which sellers can appear in both sponsored and organic positions and examine how this affects equilibrium outcomes and consumer welfare. Our model captures a key tradeoff: high-quality sellers value the visibility from sponsored placement but also expect to appear prominently in organic rankings. As a result, under certain conditions, lower-quality sellers may outbid them to obtain the sponsored position - lending some support to the FTC's concern. However, we show that this outcome only arises when all sellers are relatively high-quality, which limits potential consumer harm.

1 Introduction

Search platforms such as Amazon play a central role in online retail by helping consumers navigate large product assortments through ranked lists of search results. These lists typically contain a combination of organic listings and sponsored advertisements, with these sponsored listings frequently occupying the most prominent positions. Empirical evidence shows that these top positions attract a disproportionate share of consumer attention and clicks (Ursu 2018).

Concerns about the mixing of sponsored and organic listings have recently come under regulatory scrutiny. In its ongoing antitrust case against Amazon, the U.S. Federal Trade Commission (FTC) alleges that the platform "degrades the customer experience by replacing relevant, organic [listings] with paid advertisements," thereby extracting rents at the expense of consumers. Yet most theoretical models offer little support for this concern. Canonical models of sponsored search auctions predict that high-quality sellers bid more aggressively and win the top slots when consumers search top-down (Athey and Ellison 2011), while models of organic search assume the platform is aligned with consumer interests and ranks sellers by relevance or match quality. These frameworks typically examine either sponsored or organic listings in isolation, and therefore conclude that the most prominent positions are occupied by the most relevant sellers.

Some work has taken a more integrated view of the two listing types (Xu et al. 2011; Xu et al. 2012). For instance, Xu et al. (2012) model bidding behavior when both sponsored and organic results coexist. They highlight two incentives for bidding for sponsored positions: a promotive effect, whereby firms seek greater visibility, and a preventive effect, whereby they aim to block competitors from gaining it. However, their model does not endogenize consumer search behavior so that they cannot account for consumers anticipating that low-quality firms may win sponsored positions and adjusting their search accordingly. As a result, they do not capture a key strategic trade-off that arises when high-quality firms anticipate favorable organic positions.

This paper identifies and formalizes that trade-off. When sellers can appear in both sponsored and organic positions, high-quality firms face a choice: they may value the additional visibility from winning the prominent sponsored position, but also expect to be assigned a prominent organic position due to their quality. As a result, they may bid less aggressively, creating scope for a lower-quality rival to win the sponsored position - even when that rival is less relevant to consumers. This mechanism has not been explored in prior theoretical work, and it provides a novel explanation for how platform-sponsored rankings can distort the consumer search path. In particular, we show that such distortions can arise endogenously, even when consumers behave optimally and platforms assign organic positions in a way that maximizes user satisfaction. We see this as a key advantage of our model, since if consumers exogenously prefer sponsored positions like in previous literature, it is mechanically easier to find an equilibrium where consumers are harmed by visiting low-quality firms in sponsored positions, which can lead to misguided policy implications.

Based on the framework of Athey and Ellison (2011), we develop a model in which two firms compete for one sponsored and two organic positions on a platform. Consumers sequentially evaluate listings until their need is met, incurring a search cost c per click. A firm's quality q determines the probability that a consumer's need is satisfied upon clicking. The platform ranks organic listings by quality, and firms strategically bid for the sponsored position, anticipating how their placement affects consumer behavior. We consider two informational environments - one with complete information about rival quality, and one with incomplete information. In both settings, we find that a lower-quality seller may outbid a higher-quality one for the sponsored slot, despite being less relevant. Under incomplete information, this occurs only when both sellers are relatively high quality, limiting the welfare loss from the distortion. Under complete information, this occurs when the quality of both firms are sufficiently close, even when they are low-quality, since the higher-quality firm always knows it will receive the top organic position.

The complete information setting serves as a clean benchmark, highlighting how the trade-off shapes bidding incentives, even in the absence of uncertainty. Comparing the two settings reveals how equilibrium bidding behavior is shaped by both strategic positioning and informational frictions. In the complete information setting, bidding functions are asymmetric due to observable differences in quality. In particular, the higher-quality firms account for the fact that they will receive the top organic position. In the incomplete information setting, equilibrium strategies are symmetric, and quality uncertainty changes the equilibrium results.

By modeling position auctions in the presence of organic listings and optimal consumer search, our framework highlights a new mechanism through which sponsored advertisements can distort the consumer experience. High-quality sellers may rationally avoid bidding for the sponsored slot, not because they undervalue visibility, but because they already expect to be found through organic search. This insight helps reconcile mixed empirical findings on the effectiveness of sponsored advertising. Blake et al. (2015) show that sponsored listings mainly influence new users, while most traffic comes from repeat users who are not strongly influenced. The consumers in our model can be thought of as repeat users since the platform knows consumers' match with firms, which in reality occurs because the platform learns repeat consumers' willingness-to-pay and preferences over time through their repeated choices.

Moshary (2025) finds that sponsored listings may cannibalize organic clicks and reduce total sales. In contrast, Yang and Ghose (2010) report complementarity between the two listing types. Our model offers a unifying perspective: when organic listings are informative of quality and consumers search optimally, sponsored listings can distort outcomes by altering the order in which consumers consider alternatives. Apart from the previously mentioned Xu et al. (2011) and Xu et al. (2012) which do not feature endogenous consumer search, to the best of our knowledge we are the only other theoretical model of position auctions with organic search.

2 Model

In this section, we present our model of position auctions, organic search, and optimal consumer search, based on the framework of Athey and Ellison (2011). We consider a continuum of consumers, each of whom has a specific need. A consumer receives a benefit of 1 if the need is successfully fulfilled. To identify firms capable of meeting this need, consumers visit a search platform. For each consumer, the platform displays one sponsored link, which is assigned to firms via a second-price sealed-bid auction, and two organic positions which are allocated to firms by the platform.

We model two competing firms, labeled Firm 1 and Firm 2, each offering a product or service that may satisfy any consumer's need. These firms list their products on the search platform, and compete for the sponsored advertisement position, denoted position s, and are assigned one of two organic positions on the platform, denoted o_1, o_2.[1] The quality of each firm's product determines its ability to meet consumer needs. Specifically, Firm $i \in \{1, 2\}$ satisfies consumer j's need with probability q_i^j.

Each q_i^j is independently drawn from a common, atomless distribution function F with support $[0, 1]$. For tractability, we assume throughout the analysis that $F = U(0, 1)$. The probability q_i^j represents the firm i's match quality or effectiveness in fulfilling consumer j's need. Thus, firms have different match qualities with different consumers so that the same firm is not always the best match across consumers. A firm earns a payoff of 1 each time it successfully satisfies a need, and zero otherwise. Thus, a firm's expected payoff depends on both its placement on the search platform and its underlying quality q_i^j.

Each consumer j can click on a listing, incurring a search cost c_j, in order to acquire information about the firm. A click may either fully reveal whether the firm can meet the consumer's need or provide only partial information about match quality. To keep the

[1] We formally use the term position but also informally refer to positions as listings interchangeably throughout the paper.

model general, we remain agnostic about the exact learning process and instead represent it in reduced form: upon clicking firm i, the consumer's need is satisfied with probability q_i^j. Thus, a *click* is the costly act of gathering information, while a *match* is the probabilistic realization that the firm meets the consumer's need. This formulation nests settings where consumers learn perfectly from a click (full information) as well as settings where they gain only partial or noisy signals about match quality.

We assume that search costs for organic positions c_j^o are drawn from a common, atomless distribution G with support $[0,1]$. To capture the positional advantage of the sponsored listing, without loss of generality, we assume that the search costs for the sponsored position, c_{sj}, is given by $c_{oj} - \delta$, where $\delta \geq 0$, which means the sponsored position can be easier for consumers to notice and click, reflecting its visual prominence on the search results page. Consumers make search decisions optimally: they continue clicking through available listings until their need is satisfied or until the expected marginal benefit of an additional click falls below their individual search cost c_j.

We also assume that the platform orders firms in the organic positions for each consumer j by their quality q_i^j: the higher-quality firm appears in position o_1 and the lower-quality firm appears in position o_2. Implicitly, we assume the platform observes q_i^j, and think this represents settings like Amazon's search platform well. As Blake et al. 2015 point out, many users tend to be experienced users or those with accounts which have a history of transactions on the platform. Thus, the platform has knowledge about each consumers' willingness-to-pay and preferences, so we think it is reasonable to assume they have knowledge about consumer-firm match.

We allow for sponsored and organic positions to have complementary benefits, by allowing consumers to click onto a firm's sponsored *and* organic listing throughout their search sequence. A consumer who has already clicked on a firm's sponsored or organic listing, and sees the firm again in an organic or sponsored position, clicks onto this second listing with probability $\theta \in [0,1]$. This formulation may capture empirically observed behavioral phenomena such as forgetfulness or the desire to re-verify information, which lead consumers to re-engage with firms they have already seen. Since a second click may yield less incremental information than the first, one could introduce a discount factor $d \in [0,1]$ such that the probability of a match on a second click is $\theta d q_i$. However, because the equilibrium analysis depends only on the product $m = \theta d$, a model with an explicit discount factor simply rescales θ, yielding the same conclusions. We therefore work with a single parameter θ without loss of generality.

Consistent with practice and the existing literature (Athey and Ellison 2011), we assume that the search platform allocates position s via a second-price sealed-bid auction. Note that in practice, platforms accommodate many firms, and thus use a generalized second-price (GSP) auction (Kim and Pal 2025). In our setting with two firms, the GSP auction simplifies to second-price.

Before proceeding, we highlight some simplifying assumptions embedded in the model. Firms are symmetric except for their quality q_i: each firm earns the same profit per successful match, and we abstract from the pricing decisions of firms. By focusing on the probability of need satisfaction rather than on price-setting, we capture environments - common in search platforms as opposed to price comparison sites - where consumers base their search decisions on perceived match quality, and firms compete for visibility via bids.

Allowing firms to set prices would introduce incentives to charge monopoly prices, as in the models of Diamond (1971) and Xu et al. (2011). Thus, our focus on match probabilities simplifies the model while preserving the core strategic trade-offs. We leave extensions that incorporate heterogeneous values conditional on meeting consumer needs for future research.

Finally, the timing of the model is as follows. Firms and consumers draw their qualities and search costs. Firms commit to a bidding strategy $b_i(q_i, q_{-i})$. The platform gives consumers their search results, and consumers follow their optimal search strategy until they have clicked on all positions and/or their need is met. Since consumers draw costs independent of one another, the analysis can be done just at the consumer level. Thus we drop j subscripts in the analysis.

2.1 Complete Information

We begin by analyzing the case in which firms have complete information about each other's product quality. We think of this as a benchmark case which allows us to highlight the key trade-off in our model even in the absence of uncertainty. Without loss of generality, we focus our analysis on the realizations of q_1, q_2 where firm 1 has higher quality than firm 2, implying that firm 1 always occupies the top organic position. Let Eq_{o_1} and Eq_{o_2} denote the expected quality of the sellers occupying positions o_1 and o_2, respectively. Let $H(q_1, q_2)$ denote the joint distribution function of q_1, q_2. Since organic rankings are determined by firm quality, it follows that:

$$Eq_{o_1} = \int_0^1 \int_0^1 \max(q_1, q_2) dH(q_1, q_2), \tag{1}$$

$$Eq_{o_2} = \int_0^1 \int_0^1 \min(q_1, q_2) dH(q_1, q_2). \tag{2}$$

Firms compete in a second-price auction for position s, using a bidding function $b(\cdot)$.

Hence,

$$Eq_s = \int_0^1 \int_0^1 q_{i^*(q_1,q_2)} dH(q_1, q_2), \qquad where \quad i^*(q_1, q_2) = \arg\max_{i \in \{1,2\}} b_i(q_1, q_2). \tag{3}$$

Eq_s is the consumer's expectation of the quality of the firm in position s, based only on the equilibrium bidding strategy and the distribution of quality. In principle, consumers could also infer Eq_s by observing the organic position of the firm in position s. We consider the sponsored listing to be more visible to the consumer, so that the consumer's belief about the quality of the firm in position s is formed irrespective of the organic rankings. For instance, the consumer makes the decision of whether or not to click the sponsored firm before observing the rest of their search results. This could be because they expend effort scrolling past the sponsored advertisement in order to see the rest of the organic rankings. This is certainly more likely with more than one sponsored position, as often occurs in reality, but our model with fewer results is sufficient to highlight the forces that could lead to lower-quality firms winning sponsored positions. It follows naturally that $Eq_{o_1} \geq Eq_s \geq Eq_{o_2}$.

Consumers draw the organic position search cost c_o and thus also realize their sponsored position search cost $c_s = c_o - \delta$. They begin by clicking the position that yields the highest expected payoff and continue sequentially until the expected benefit falls below their respective search costs.

The need for $\delta > 0$ is now made clear. Since $Eq_s \leq E_{o_1}$, consumers never expect a larger payoff from clicking position s unless they face a lower search cost. δ is the parameter which governs how much lower the search cost is for the sponsored listing relative to the organic listings. It is intuitive to think that sponsored and organic listings do not differ greatly in visibility in reality. In this sense, δ governs the irrational preference for the prominent positions given to sponsored advertisements, a behavior often observed by consumers.

Following this logic, the consumer's optimal search strategy can be broken down into two cases. When $\delta \geq Eq_{o_1} - Eq_s$, consumers obtain a larger payoff from clicking position s relative to position o_1, since $Eq_s - c_s = Eq_s - (c_o - \delta) \geq Eq_{o_1} - c_o$. Conversely, when $\delta \leq Eq_{o_1} - Eq_s$, consumers obtain a larger payoff from clicking position o_1. Thus, in the first case, consumers' optimal search order is positions s, o_1, o_2, provided the expected payoffs are positive for any position. In the second case, consumers' optimal search order is o_1, s, o_2.

In any case, which organic positions the consumer will ever search depends on the value of c_o relative to Eq_{o_1} and Eq_{o_2}. They will:

- Search both positions o_1, o_2 if $c_o \leq Eq_{o_2}$,
- Search only position o_1 if $Eq_{o_2} < c_o \leq Eq_{o_1}$,

- Search neither o_1 or o_2 if $c_o > Eq_{o_1}$.

Consumer j will also click position s if $c_s \leq Eq_s$, which corresponds to $c_o \leq Eq_s + \delta$. In the equations that follow, it will be necessary to know whether $Eq_s + \delta$ is greater or less than Eq_{o_1}. If $\delta \geq Eq_{o_1} - Eq_s$, then $Eq_s + \delta \geq Eq_{o_1}$, and vice versa if $\delta \leq Eq_{o_1} - Eq_s$.

First consider the case where $\delta \leq Eq_{o1} - Eq_s$. In this case, we have that $c_s \leq Eq_s$ corresponds to $c_o \leq Eq_s + \delta \leq Eq_{o_1}$. When firm 1 wins position s, the expected payoffs for firm 1 and 2 are the following. For firm 1 we have:

$$G(Eq_{o_2})(q_1+(1-q_1)\theta q_1)+(G(Eq_s + \delta) - G(Eq_{o_2}))\,(q_1+(1-q_1)\theta q_1)+(G(Eq_{o_1})-G(Eq_s+\delta))q_1. \tag{4}$$

The first term corresponds to the payoff from both the sponsored and position o_1 for consumers with $c_o \leq Eq_{o_2}$. The second term covers the payoff from position s alone for $Eq_{o_2} < c_o \leq Eq_s + \delta$. The third term reflects the benefit from position s for consumers with $Eq_s + \delta < c_o \leq Eq_{o_1}$. Consumers with $c_o > Eq_{o_1}$ do not click any listings. For firm 2 we have:

$$G(Eq_{o_2})(1 - q_1)(1 - \theta q_1)q_2. \tag{5}$$

Here, firm 2 only appears in position o_2 and may be matched only with consumers for whom $c_o \leq Eq_{o_2}$. When firm 2 wins position s instead, the expected payoff for firm 1 becomes:

$$G(Eq_{o_2})q_1 + (G(Eq_s + \delta) - G(Eq_{o_2}))q_1 + (G(Eq_{o_1}) - G(Eq_s + \delta))q_1, \tag{6}$$

and the expected payoff for firm 2 becomes:

$$G(Eq_{o_2})(1 - q_1)(q_2 + (1 - q_2)\theta q_2) + (G(Eq_s + \delta) - G(Eq_{o_2}))(1 - q_1)q_2. \tag{7}$$

Since the auction is second-price, a classic result states that it is weakly optimal for firms to bid their true value. Thus, given the payoffs above, firm 1's bidding function is given by (4)-(6):

$$b_1(q_1, q_2) = [G(Eq_s + \delta)(1 - q_1)\theta q_1]^+, \tag{8}$$

[2] and firm 2's bidding function is given by (7)-(5):

$$b_2(q_1, q_2) = [G(Eq_{o_2})(1 - q_1)q_2((1 - q_2)\theta - (1 - \theta q_1)) + G(Eq_s)(1 - q_1)q_2]^+. \tag{9}$$

We summarize this result in the following proposition:

[2] + means we take the maximum value between 0 and the value inside the square bracket

Proposition 1. *Under complete information, where firms observe each other's quality and* $\delta \leq Eq_{o_1} - Eq_s$*, consumers click on positions in the following order:* o_1, s, o_2*. A consumer clicks a listing if and only if the expected payoff—*Eq_{o_1}*,* Eq_s*, or* Eq_{o_2}*—exceeds their respective search costs* c_o*,* c_s*, and* c_o*. Firms anticipate this search behavior and submit bids according to the functions* $b_1(q_1, q_2)$ *and* $b_2(q_1, q_2)$ *given by equations* (8) *and* (9) *respectively.*

We now turn to the case where $\delta \geq Eq_{o_1} - Eq_s$. The expected payoff for firm 1 when it wins the auction is:

$$G(Eq_{o_2})(q_1+(1-q_1)\theta q_1)+(G(Eq_{o_1})-G(Eq_{o_2}))(q_1+(1-q_1)\theta q_1)+(G(Eq_s)-G(Eq_{o_1}))q_1. \quad (10)$$

The first term captures the benefit from both positions o_1, s for consumers with $c_o \leq Eq_{o_2}$. The second term captures the benefit from position o_1, s for consumers with $Eq_{o_2} < c_o \leq Eq_{o_1}$. The third term captures the benefit from position s only for consumers with $Eq_{o_1} < c_o \leq Eq_s + \delta$. Consumers with $c_o > Eq_s + \delta$ do not click any listings. The expected payoff for firm 2 is:

$$G(Eq_{o_2})(1-q_1)(1-\theta q_1)q_2. \quad (11)$$

In this case, firm 2 only appears in position o_2, and is clicked only when $c_o \leq Eq_{o_2}$. If firm 2 wins the auction, the expected payoff for firm 1 becomes:

$$G(Eq_{o_2})(1-q_2)q_1 + (G(Eq_{o_1}) - G(Eq_{o_2}))(1-q_2)q_1, \quad (12)$$

and the expected payoff for firm 2 is:

$$G(Eq_{o_2})(q_2+(1-q_2)(1-q_1)\theta q_2)+(G(Eq_{o_1})-G(Eq_{o_2}))q_2+(G(Eq_s+\delta)-G(Eq_{o_1}))q_2. \quad (13)$$

Firm 1's bidding function is given by (10)-(12):

$$b_1(q_1, q_2) = [G(Eq_{o_1})(1-q_1)\theta q_1 + G(Eq_s+\delta)q_1 - G(Eq_{o_1})(1-q_2)q_1]^+ \quad (14)$$

and firm 2's bidding function is given by (13)-(11):

$$b_2(q_1, q_2) = [G(Eq_{o_2})(1-q_2)(1-q_1)\theta q_2 + G(Eq_s+\delta)q_2 - G(Eq_{o_2})(1-q_1)(1-\theta q_1)q_2]^+ \quad (15)$$

We summarize this result in the following proposition:

Proposition 2. *In the setting with complete information and* $\delta \geq Eq_{o_1} - Eq_s$*, consumers search in the order* s, o_1, o_2*. A consumer clicks a listing if and only if the expected payoffs* Eq_s, Eq_{o_1}, Eq_{o_2} *exceeds their respective search costs,* c_s, c_o, c_o*. Firms anticipate this behavior* *and bid according to functions $b_1(q_1, q_2)$ and $b_2(q_1, q_2)$ as defined in equations (14) and (15) respectively.*

We return to analyzing these bidding functions in section 3.

2.2 Incomplete Information

We now turn to the setting of incomplete information, where firms do not observe their rival's quality. As in the complete information case, the analysis is divided into two regimes depending on the relative magnitude of δ and $Eq_{o_1} - Eq_s$. When $\delta \geq Eq_{o_1} - Eq_s$, the optimal search order is s, o_1, o_2 and when $\delta \leq Eq_{o_1} - Eq_s$, the optimal search order is o_1, s, o_2.

We first consider the case where $\delta \leq Eq_{o_1} - Eq_s$. The expected payoff for firm i when it wins position s is given by:

$$G(Eq_{o_2})(q_i(q_i + \theta q_i(1 - q_i)) + (1 - q_i)(1 - \frac{1 + q_i}{2})(q_i + (1 - q_i)\theta q_i))+$$
$$(G(Eq_s+\delta)-G(Eq_{o_2}))(q_i(q_i+(1-q_i)\theta q_i)+(1-q_i)(1-\frac{1 + q_i}{2})q_i)+(G(Eq_{o_1})-G(Eq_s+\delta))q_i^2. \tag{16}$$

The first term corresponds to the payoff from both the sponsored and position o_1 for consumers with $c_o \leq Eq_{o_2}$. The second term represents the payoff from position s only for consumers with $Eq_{o_2} < c_o \leq Eq_s + \delta$. The third term captures the payoff from position s for consumers with $Eq_s + \delta < c_o \leq Eq_{o_1}$. Consumers with $c_o > Eq_{o_1}$ do not engage in any search. If firm i does not win position s, its expected payoff becomes:

$$G(Eq_{o_1})(q_i^2 + (1 - q_i)(1 - \frac{1 + q_i}{2})(1 - \theta\frac{1 + q_i}{2})q_i)) + (G(Eq_{o_1} - G(Eq_s + \delta))q_i^2. \tag{17}$$

Given these expected payoffs, firm i's bidding function is given by (16)-(17):

$$b_i(q_i) = [G(Eq_{o_2})(1 - q_i)(1 - \frac{1 + q_i}{2})q_i(\frac{3}{2}\theta + \frac{3}{2}\theta q_i - 1)+$$
$$G(Eq_s + \delta)(\theta q_i^2(1 - q_i) + (1 - q_i)(1 - \frac{1 + q_i}{2})q_i)]^+. \tag{18}$$

Proposition 3. *Under incomplete information, where firms do not observe each other's quality and $\delta \leq Eq_{o_1} - Eq_s$, consumers search in the following order: o_1, s, o_2. A consumer clicks a listing if and only if the expected payoff—Eq_{o_1}, Eq_s, or Eq_{o_2}—exceeds their respective search costs c_o, c_s, c_o. Anticipating this search behavior, each firm submits a bid according to the function $b_i(q_i)$ defined in equation (18).*

We now turn to the case where $\delta \geq Eq_{o1} - Eq_s$. In this regime, consumer j begins their search at position s. The expected payoff for firm i when it wins position s is given by:

$$G(Eq_{o_2})(q_i + (1-q_i)(\theta q_i^2 + (1-q_i)(1 - \frac{1+q_i}{2})\theta q_i))+ \\ G(Eq_{o_1} - G(Eq_{o_2}))(q_i + (1-q_i)\theta q_i^2) + (G(Eq_s + \delta) - G(Eq_{o_1}))q_i \quad (19)$$

If firm i does not win position s, its expected payoff is:

$$G(Eq_{o_2})(q_i^2(1 - \frac{q_i}{2}) + (1-q_i)(1 - \frac{1+q_i}{2})(1 - \theta\frac{1+q_i}{2})q_i)+ \\ (G(Eq_{o_1}) - G(Eq_{o_2}))q_i^2(1 - \frac{q_i}{2}) + (G(Eq_s + \delta) - G(Eq_{o_1})) * 0 \quad (20)$$

Given these expected payoffs, firm i's bidding function is given by (19)-(20):

$$b_i(q_i) = \frac{1}{2}G(Eq_{o_2})(1-q_i)^2 q_i(\frac{3}{2}\theta - 1 - \frac{1}{2}\theta q_i) + G(Eq_{o_1})q_i((1-q_i)\theta q_i - q_i(1 - \frac{q_i}{2})) + G(Eq_s + \delta)q_i \quad (21)$$

We summarize this result in the following proposition:

Proposition 4. *Under incomplete information, where firms do not observe each other's quality and $\delta \geq Eq_{o_1} - Eq_s$, consumers search in the following order: s, o_1, o_2. A consumer clicks a listing if and only if the expected payoff—Eq_s, Eq_{o_1}, or Eq_{o_2}—exceeds their respective search costs c_s, c_o, c_o. Anticipating this behavior, each firm submits a bid according to the bidding function $b_i(q_i)$ as defined in equation (21).*

3 Model Analysis

Having established the theoretical framework under both complete and incomplete information, we now turn to a quantitative analysis to explore the model's implications.

Our analysis focuses on the equilibrium bidding function $b_i(q_i)$ and how it varies across different information structures and parameter values. Specifically, we compare bidding strategies under complete and incomplete information and study how firms respond to two key parameters: θ, which captures the likelihood that a consumer clicks the second listing for a firm after an initial click, and δ, which reflects the consumer's initial bias toward the sponsored listing by representing the relative cost of clicking the sponsored versus organic listings.

One central object in this analysis is the expected quality of the sponsored listing, Eq_s. This value plays a critical role: it enters directly into the consumer's search order, which affects the likelihood of a firm being chosen, and in turn influences each firm's optimal bid. But since bidding behavior also determines which firm appears in the sponsored listing, Eq_s itself depends on equilibrium bidding and must be solved as a fixed point. This interdependence makes analytical solutions intractable, so we employ a numerical fixed-point algorithm to determine Eq_s for given values of θ and δ.

To make it easy to demonstrate the forces captured by our model, we assume that the organic search cost distribution $G = U(0, 1)$. However, the analysis could easily be repeated for other distributions, and simulating our model in these cases reveals that our results are robust to other distributions.

3.1 Complete Information

To build intuition, we begin with the complete information setting, where each firm knows its rival's quality. We fix $\delta = 0.15$ and assume $q_2 = 0.25$, focusing on the case where $q_1 > q_2$ since the analysis assumes firm 1 is higher-quality.[3] Figure 1 illustrates how bidding behavior varies with θ, which governs the likelihood that a consumer revisits a previously clicked firm. The x-axis shows firm 1's quality q_1, and the y-axis shows equilibrium bids. The solid line depicts firm 1's bidding function, and the dashed line represents firm 2.

Figure 1: The Effect of θ on the Complete Information Bidding Function

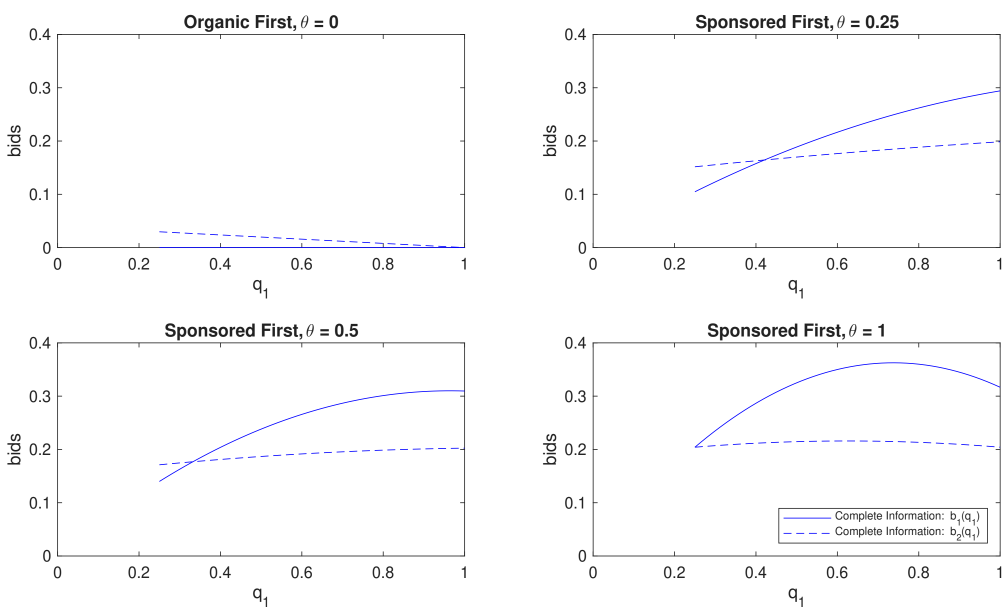

δ is fixed to 0.15.

[3]We use other values of q_2 in Appendix 5.1.

We explore four values of θ: 0, 0.25, 0.5, and 1. Each reflects a different assumption about consumer behavior, from never revisiting a firm ($\theta = 0$) to always doing so ($\theta = 1$). The figure shows that bids increase with θ, as firms place more value on winning the sponsored slot when the possibility of a second click rises. Notably, firm 1's bids rise more steeply than firm 2's, since firm 1 - by virtue of its higher quality and stronger organic position - derives more value from being considered multiple times. As θ increases, so does Eq_s, the expected quality in the sponsored listing, making it more likely that consumers start their search there. When $\delta \geq Eq_{o_1} - Eq_s$, the search order flips and consumers begin with the sponsored listing.

The shape of the bidding function itself reveals an important insight: it is not always monotonic, in particular for high values of θ when consumers search sponsored positions first. Instead, in this case it is concave in q. This inverse-U shape captures a fundamental trade-off. For lower-quality firms, winning the sponsored position can dramatically increase visibility, so they bid aggressively as their quality rises. However, for very high-quality firms, the marginal gain from appearing in the sponsored position diminishes. These firms are already likely to meet consumers' needs in position o_1, which is clicked before o_2. Thus, even though high-quality firms benefit more from potential second clicks, their incentive to bid eventually diminishes - producing the concave curve.

This pattern is clearest when θ is high, as shown in the bottom panels of Figure 1. As θ increases, firms are more certain that the second click will occur but this second click is most valuable for medium-quality firms which increase their bids more than other firms, so the peak of firm 1's bidding function moves closer to the median of possible values of q.

To analyze the effect of changing θ, let us begin with the case where $\theta = 0$. Here, consumers never revisit a firm they have already clicked on, so the value of a second listing is zero. In this case, consumers begin their search from the position o_1 since $\delta \geq Eq_{o_1} - Eq_s$ is not satisfied. Thus, firm 1 - already holding o_1 - has no additional incentive to win the sponsored position and thus bids zero. Firm 2, on the other hand, benefits by winning the sponsored position, where it can attract more traffic than from position o_2, since consumers require a lower search cost draw to click on o_2. As a result, it submits a positive bid.[4]

As θ decreases, consumers are less likely to revisit a firm they've already seen. In this environment, the value of holding both positions o_1 and s falls for firm 1. Since a second click is unlikely, firm 1 gains little from occupying both positions and thus bids less aggressively. In contrast, firm 2 still sees value in moving from the position o_2 to the more visible position

[4]However, if consumers begin their search from position s, the logic shifts. In that case, if firm 2 were to win position s, it would be seen first. Firm 1 might lose some of its initial exposure, causing a decline in profit. To prevent this, firm 1 may also bid positively to block firm 2 from winning the sponsored listing.

s, allowing it to attract more initial clicks. As a result, we observe an expanding region of q_1 where firm 2's bids exceed those of firm 1. However, when q_1 is sufficiently higher than q_2, the probability that a second click leads to a match becomes large enough that firm 1 can still benefit from holding both positions - making it worthwhile for the high-quality firm to outbid its rival.

Next, we examine the role of δ, holding $\theta = 0.5$ fixed. Figure 2 plots how bids vary with q_1 as δ changes from 0 to 0.3. Recall that δ reflects the cost difference between sponsored and organic listings. As δ increases, the relative cost of exploring organic listings rises, drawing more consumer attention to the sponsored listing. This shift makes the sponsored listing more valuable, especially to firm 1, which benefits disproportionately due to its higher quality, causing its bidding function to become less concave. Consequently, firm 1 bids more aggressively, and its bid curve dominates that of firm 2 over a wider range of q_1.

Figure 2: The Effect of δ on the Complete Information Bidding Function

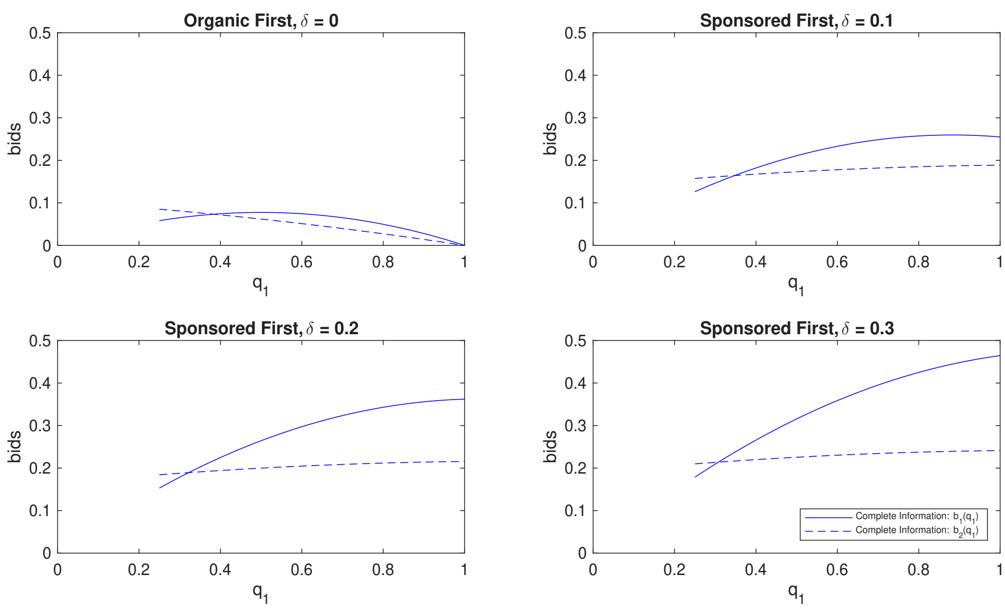

θ is fixed to 0.5.

3.2 Incomplete Information

We now turn to the incomplete information setting, where firms do not know their rival's quality. Figures 3 and 4 show the effect of changing θ and δ respectively on the equilibrium bidding functions. As in the complete information case, we observe a concave bidding function. The same force holds: while higher-quality firms benefit more from potential second clicks, their marginal returns decline once their quality is high enough to reliably satisfy consumers on the first click. Additionally, in the incomplete information case, the higher-quality firm must also weigh the risk of not being the highest-quality firm since in that case their second click would occur in position o_2, only after their higher-quality rival gets a chance to match in position o_1.

Figure 3: The Effect of θ on the Incomplete Information Bidding Function

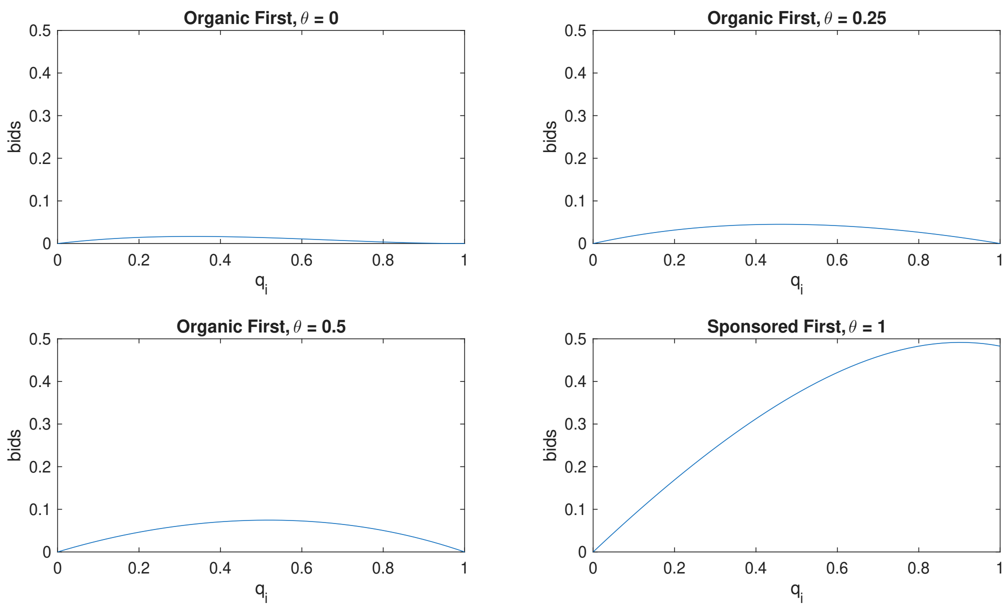

δ is fixed to 0.15.

The effect of changing θ and δ on the incomplete information bidding functions is similar to that of the complete information case. However, notice that when $\delta \geq Eq_{o_1} - Eq_s$ is satisfied so that consumers click on the sponsored listing first, the decreasing portion of the bidding function is at the right end of the quality-domain. So, a higher-quality firm can only be outbid by its lower-quality rival when both firms are relatively high-quality. To directly compare the two information structures, Figure 5 presents bidding functions under both complete and incomplete information, with $\delta = 0.2$ and $\theta = 1$ fixed.[5] The red curves correspond to the incomplete information case, while the blue curves represent complete information, with the solid line for firm 1 (the higher-quality firm) and the dashed line for firm 2. Each panel holds q_2 constant - at 0, 0.25, 0.5, and 0.75 - and plots bids for $q_1 \geq q_2$.

Across all panels, the non-monotonic pattern remains: bids increase with quality before tapering off. In the special case where $q_2 = 0$, firm 2 always bids zero, as it has virtually no chance of being selected. In other panels, both firms bid identically when they share the same quality and gradually diverge as quality increases. Importantly, even though firms

[5] We use other values of δ and θ in Appendix 5.2.

Figure 4: The Effect of δ on the Incomplete Information Bidding Function

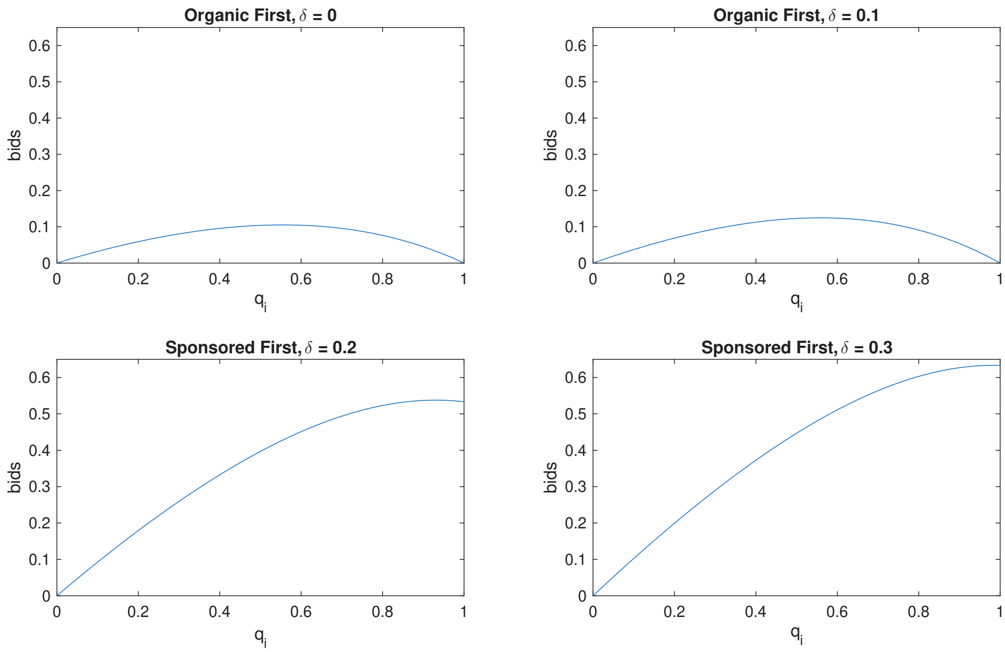

θ is fixed to 1.

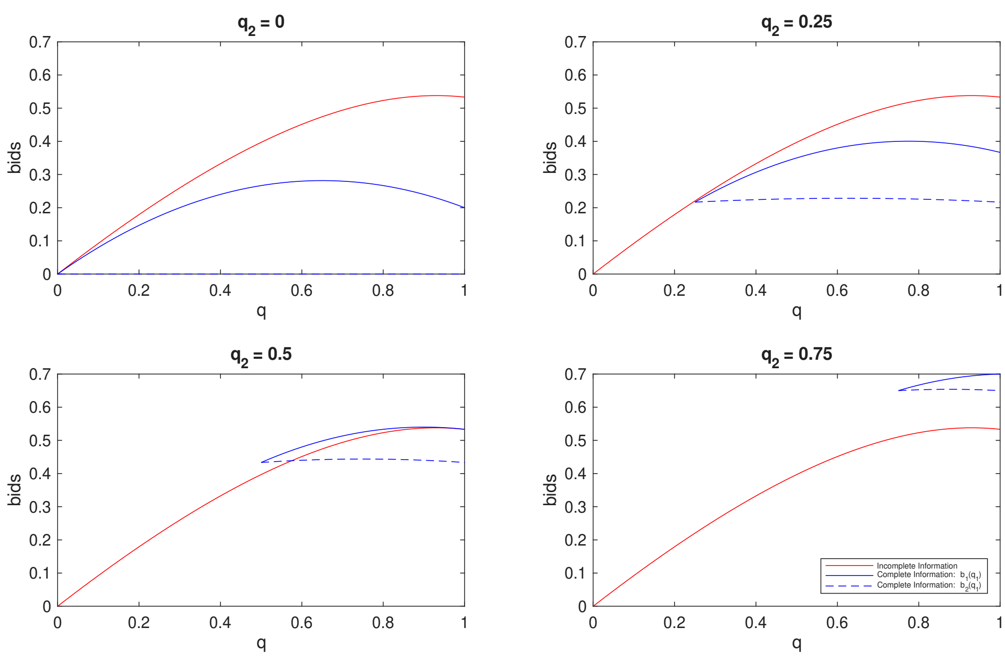

Figure 5: Comparison of the Bidding Function between Complete and Incomplete Information Cases

in the incomplete information case do not observe their opponent's quality, their behavior closely mirrors that in the complete information case.

One subtle difference arises from how firms form expectations about their competitor. Under incomplete information, firms behave as though competing against an "average" firm, with expected quality 0.5. This leads to slightly more aggressive bidding under incomplete information when the rival is weaker ($q_2 \leq 0.5$), and more conservative bidding when the rival is stronger ($q_2 \geq 0.5$). As a result, firm 1 tends to bid more in the incomplete information case when q_2 is low, and less when q_2 is high. Thus unlike other results in the auction literature, firms do not always bid higher under uncertainty. This is because high-quality firms have a high payoff even when they do not win position s.

3.3 Firm Behavior

Figure 6 illustrates how the parameter δ influences various relevant quantities in equilibrium when $\theta = 1$ in the complete information case. The horizontal axis represents the value of δ, while the vertical axis displays a set of expected quality measures derived from the model. The graph includes several key lines that help visualize consumer behavior and firm outcomes as δ changes.

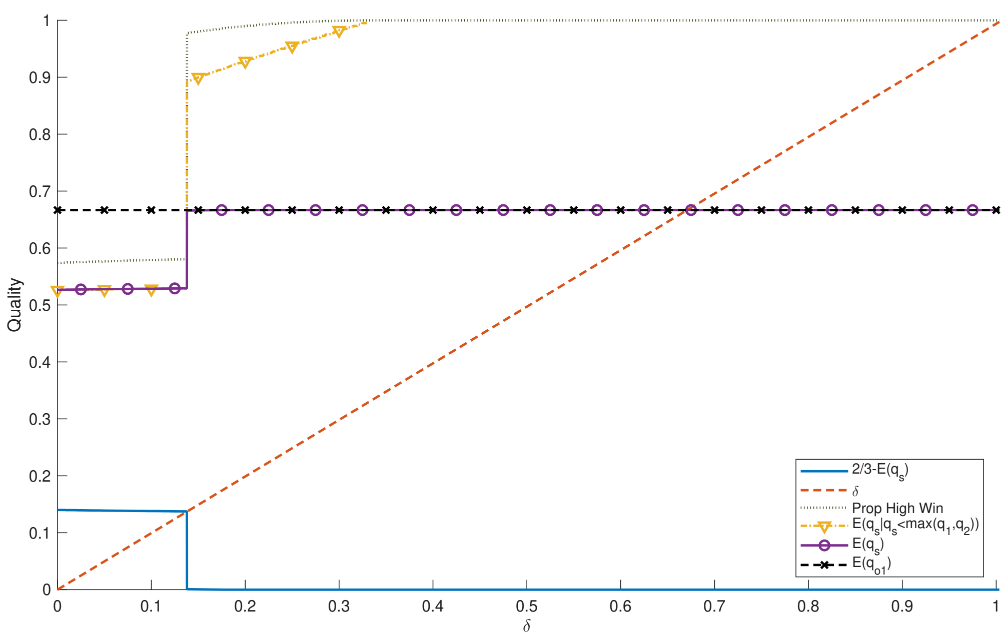

Figure 6: The Relationship Between δ and Expected Qualities

The blue line tracks the difference $\frac{2}{3} - Eq_s$. This line intersects with the red-dashed 45-degree line - which represents the value of δ - at around $\delta = 0.14$. Comparing these two reveals the consumer optimal search order. To the left of the intersection, $\delta < Eq_{o_1} - Eq_s$ where consumers begin with position o_1 , and to the right they begin with position s.

The black-dashed horizontal line denotes the expected quality of position o_1, Eq_{o_1}, which remains constant at $\frac{2}{3}$ due to the uniform distribution assumption. The purple line shows the expected quality of the firm occupying position s, Eq_s. Below the 0.14 threshold, consumers search position o_1 first. As the bidding function graphs show, in this case high-quality firms do not bid high to win position s. Thus the purple line closely follows the yellow line, which represents the expected quality of the winning firm conditional on that firm being the lower-quality of the two competitors. After the threshold, position s is held by the higher-quality firm in a majority of quality draws. At this point, Eq_s gradually converges with Eq_{o_1}, and the two lines coincide for higher values of δ.

A particularly insightful feature of the graph is the green-dotted line, which shows the probability that the higher-quality firm wins the auction for position s. Prior to the $\delta = 0.14$ threshold, this probability is around 0.5. However, once consumers begin their search with the sponsored listing, this probability increases sharply, indicating that a lower-quality firm only occasionally wins the auction. As the bidding function graph shows, this only occurs when both firms' quality draws are relatively high. Thus, even when the lower-quality firm wins the sponsored slot, the expected quality of the winning firm remains relatively high. This is reflected by the yellow line, which shows that the expected quality conditional on winning the auction remains above 0.89 after the threshold. Although the allocation is not strictly efficient, the resulting welfare loss is limited. In essence, while some consumers may encounter a suboptimal first listing, the quality gap is not large, and in equilibrium consumers face a reasonably high standard of match quality overall.

4 Conclusion

In this paper, we analyze a model of sponsored and organic search. Consumers who have a "need" visit a search platform to find firms that can meet their need. The platform shows consumers a ranking of the firms who differ in their ability to meet the consumer's need, or in other words, differ in quality. They also compete for a prominent sponsored position that is sold via second-price sealed-bid auction. Consumers can learn whether a consumer will meet their need by clicking on one of its listings in either organic or sponsored positions, each time paying a cost c.

We allow for the sponsored listing to receive more attention from consumers since they face a lower cost c_s relative to the cost of clicking an organic position c_o. This incentivizes high-quality firms to win the sponsored listing since they are confident in their ability to meet the consumer's need. At the same time, they face a trade-off since they would also be confident in meeting the consumer's need should the consumer eventually click on their organic listing, and high-quality firms are ranked higher in organic listings. We analyze models of complete and incomplete information, where the firms either know or do not know their rival's quality, and thus may not or may need to make an inference about their chance of obtaining a high-ranking organic listing. We show that this trade-off exists in both settings, and that under certain conditions, a lower-quality seller could outbid its higher-quality rival for the sponsored listing when consumers click on sponsored positions first.

These results lend some support to the FTC's concern. The equilibria in our model under which consumers are harmed are specifically those where consumers click on sponsored positions occupied by lower-quality sellers first. The conditions we derive for such equilibria are revealing about whether the FTC's concern is valid. We find that consumers must give sufficiently more attention to sponsored positions than organic positions. Assuming that platforms rank firms to the best of their knowledge by consumer-firm match, organic rankings are most informative to the consumer about the firms they should consider. It is only when consumers irrationally give attention to the sponsored listing without considering organic listings that they can be harmed. Indeed, there is support in the marketing literature that shows this is true (Ursu (2018)). In addition, our incomplete information setting shows that even though lower-quality rivals may outbid the high-quality firm, this only occurs when both firms are relatively high-quality. Intuitively, the sponsored listing may cause the consumer to purchase a second-best product, but only in scenarios where second-best is not much worse than first-best. This is because for very low-quality firms, even if they win the sponsored listing, they cannot be confident in matching with the consumer, and thus do not wish to bid to win sponsored positions.

However, future work must consider other important and realistic settings. First, it is natural to extend the model to allow for more than two firms and/or more than one sponsored position. In reality, platforms like Amazon host many firms, and consumer harm may be worse if the lowest-quality of these firms win sponsored positions. Allowing for more firms and/or sponsored positions would allow the model to speak to these concerns. Furthermore, there may be strategic interactions between a platform that steers consumers, for instance, and the equilibrium outcomes in sponsored positions. As a result, other platform organic rankings including those that account for steering, may result in a more harmful outcome for consumers. In addition, thus far we have assumed that the platform has full-information about consumer-firm match. If there is some uncertainty about this, for instance because consumers' willingness-to-pay is unknown, then the equilibria in the model will likely change. In the future, we hope to derive direct testable implications, specifically those which an empirical model of consumer search and sponsored search auctions could predict, thereby verifying the predictions of our model. We leave these important extensions for future work.

5 Appendix

5.1 Complete Information Bidding Function with Different q_2

Figure 7: The Effect of θ on the Complete Information Bidding Function

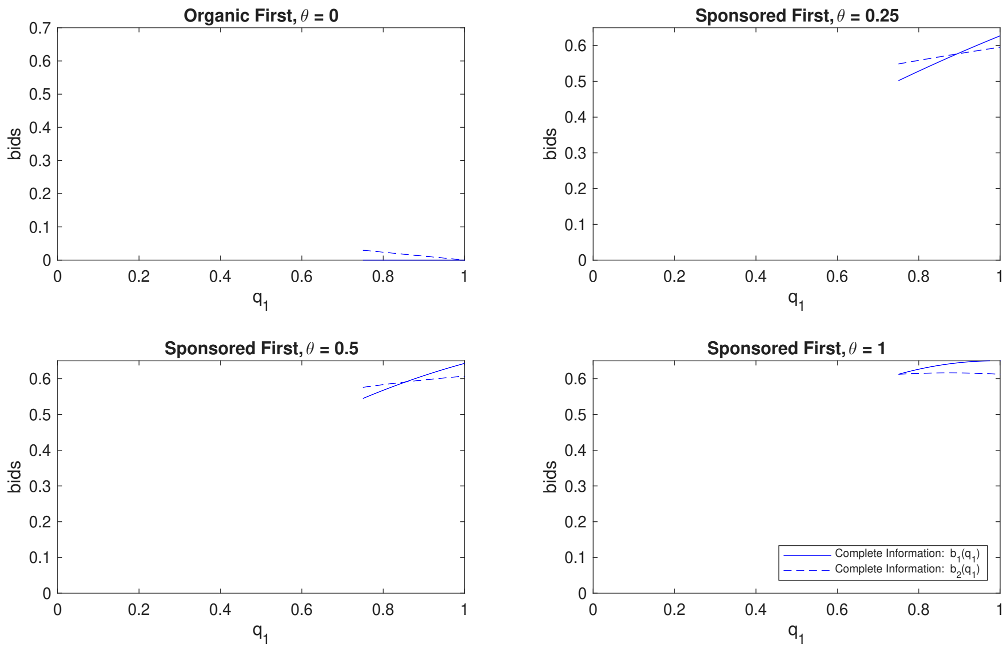

δ is fixed to 0.15, and q_2 is fixed to 0.75.

Figure 8: The Effect of δ on the Complete Information Bidding Function

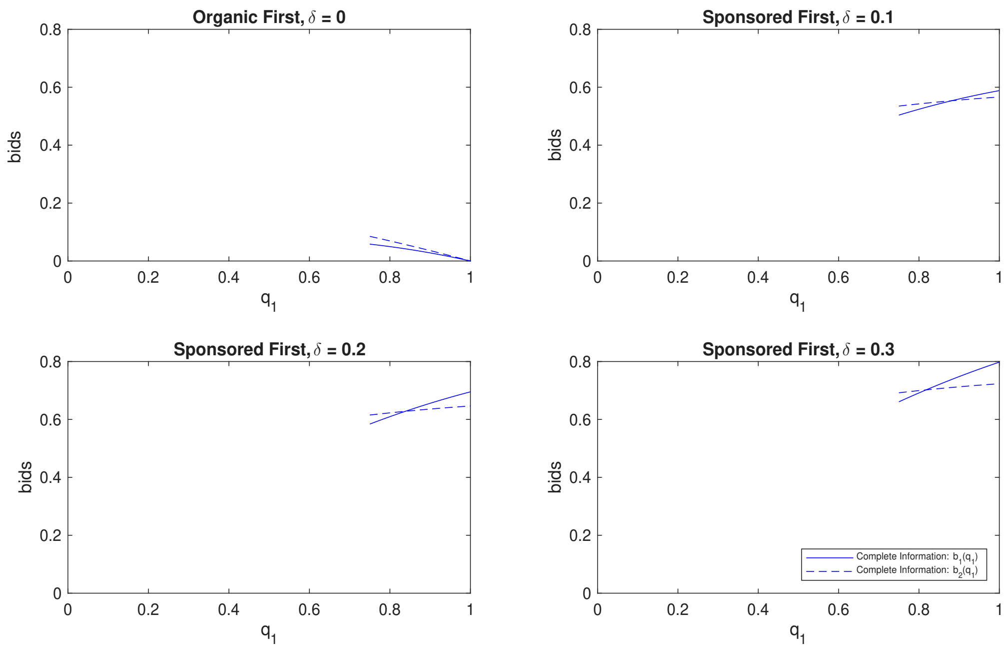

θ is fixed to 0.5, and q_2 is fixed to 0.75.

Figures 7 and 8 reproduce Figures 1 and 2, fixing q_2 to 0.75 instead of 0.25. The effects of changing θ and δ on the bidding functions are qualitatively identical.

5.2 Bidding Function Comparison

Figure 10 reproduces Figure 5, fixing θ to 0.5 instead of 1. Here the comparison between the bidding functions are qualitatively similar. However, in this case θ is not high enough to cause the bidding function to be decreasing in the right end of the quality-domain, so in the higher-quality firm never gets outbid. This can still happen in the complete information case since the higher-quality firm knows exactly what it is giving up by not winning position s. Figure 9 keeps θ at 1, but changes δ to 0.1 instead of 0.2, which means the consumer's search order switches to position o_1 first. In this case, firms always bid lower under uncertainty since the high-quality firm is confident in obtaining position o_1 and the risk of being outbid by a higher quality rival outweighs the benefit of having two chances to match.

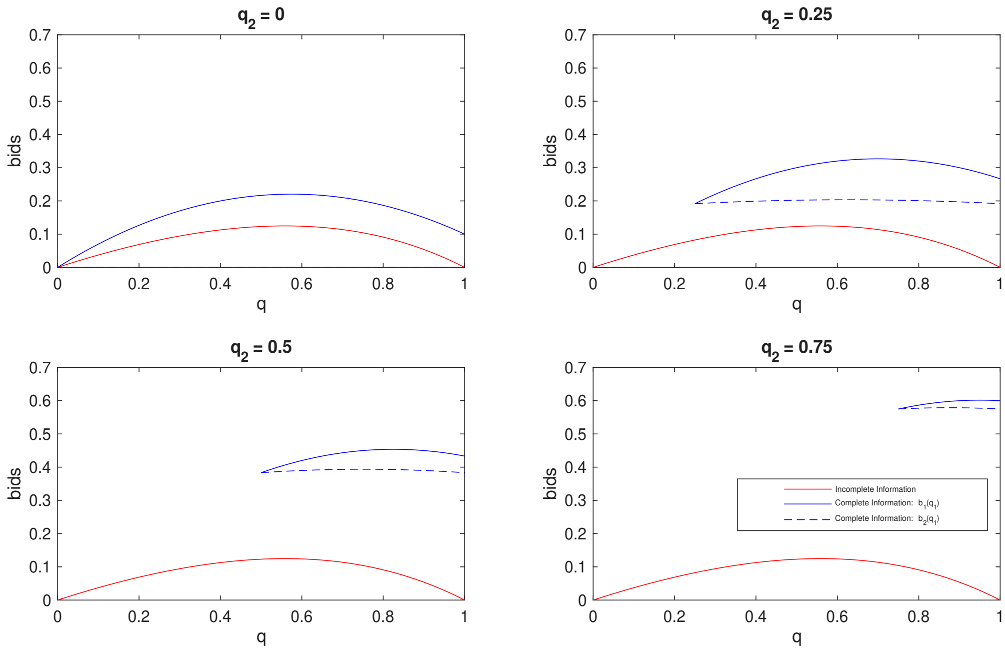

Figure 9: Comparison of the Bidding Function between Complete and Incomplete Information Cases

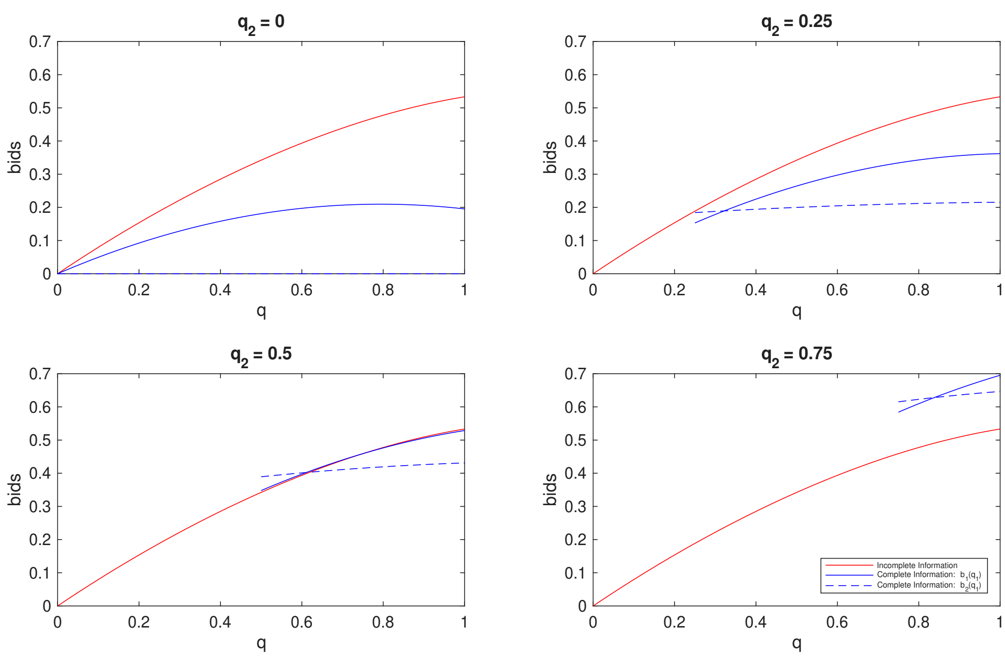

Figure 10: Comparison of the Bidding Function between Complete and Incomplete Information Cases

References

Athey, Susan and Glenn Ellison (Aug. 2011). "Position Auctions with Consumer Search*". In: *The Quarterly Journal of Economics* 126.3, pp. 1213–1270. ISSN: 0033-5533. DOI: 10.1093/qje/qjr028. eprint: https://academic.oup.com/qje/article-pdf/126/3/1213/5246144/qjr028.pdf. URL: https://doi.org/10.1093/qje/qjr028.

Blake, Thomas et al. (2015). "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment". In: *Econometrica* 83.1, pp. 155–174. DOI: https://doi.org/10.3982/ECTA12423. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA12423. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA12423.

Diamond, Peter A (1971). "A model of price adjustment". In: *Journal of Economic Theory* 3.2, pp. 156–168. ISSN: 0022-0531. DOI: https://doi.org/10.1016/0022-0531(71)90013-5. URL: https://www.sciencedirect.com/science/article/pii/0022053171900135.

Kim, Dongwoo and Pallavi Pal (2025). "Nonparametric Estimation of Sponsored Search Auctions and Impact of Ad Quality on Search Revenue". In: *Management Science* 0.0, null. DOI: 10.1287/mnsc.2023.02052. eprint: https://doi.org/10.1287/mnsc.2023.02052. URL: https://doi.org/10.1287/mnsc.2023.02052.

Moshary, Sarah (2025). "Does Sponsored Search Advertising Augment Organic Search? Evidence from an E-Commerce Platform". In: *Management Science* 0.0, null. DOI: 10.1287/mnsc.2024.04792. eprint: https://doi.org/10.1287/mnsc.2024.04792. URL: https://doi.org/10.1287/mnsc.2024.04792.

Ursu, Raluca M. (2018). "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions". In: *Marketing Science* 37.4, pp. 530–552. DOI: 10.1287/mksc.2017.1072. eprint: https://doi.org/10.1287/mksc.2017.1072. URL: https://doi.org/10.1287/mksc.2017.1072.

Xu, Lizhen et al. (2011). "Price Competition and Endogenous Valuation in Search Advertising". In: *Journal of Marketing Research* 48.3, pp. 566–586. DOI: 10.1509/jmkr.48.3.566. eprint: https://doi.org/10.1509/jmkr.48.3.566. URL: https://doi.org/10.1509/jmkr.48.3.566.

— (2012). "Effects of the Presence of Organic Listing in Search Advertising". In: *Information Systems Research* 23.4, pp. 1284–1302. DOI: 10.1287/isre.1120.0425. eprint: https://doi.org/10.1287/isre.1120.0425. URL: https://doi.org/10.1287/isre.1120.0425.

Yang, Sha and Anindya Ghose (2010). "Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?" In: *Marketing Science* 29.4, pp. 602–623. DOI: 10.1287/mksc.1090.0552. eprint: https://doi.org/10.1287/mksc.1090.0552. URL: https://doi.org/10.1287/mksc.1090.0552.