Pure-Strategy Equilibrium in the Generalized First-Price Auction*

Michael Ostrovsky† Andy Skrzypacz‡

August 1, 2022

PRELIMINARY AND INCOMPLETE. PLEASE DO NOT CIRCULATE.

Abstract

We revisit the classic result on the (non-)existence of pure-strategy Nash equilibria in the Generalized First-Price Auction for sponsored search advertising and show that the conclusion may be reversed when ads are ranked based on the product of stochastic quality scores and bid amounts, rather than solely on the latter. Moreover, the expected revenue in the pure strategy equilibrium of the Generalized First-Price Auction may substantially exceed that of the Generalized Second-Price Auction.

1 Introduction

The initial design of sponsored search auctions, implemented in 1997 by the company GoTo (later renamed Overture and subsequently acquired by Yahoo), used a very simple payment rule. Namely, each advertiser, following a click from a user, paid to GoTo the bid that it submitted to the system: e.g., if an advertiser submitted a bid of $1, then that was the amount charged every time some user clicked on that advertiser's ad.

This format, dubbed the Generalized First-Price Auction (GFP), was intuitive and easy to explain to the advertisers, but it suffered from a serious shortcoming. The ads on the results page were sorted purely based on the submitted bids (in the descending order), and so each advertiser, given a particular position, had an incentive to outbid the next highest competitor only by the smallest possible amount – 1 cent. This incentive structure resulted in the non-existence of pure-strategy equilibria, and instead pronounced rapid "cylcing" patterns, in which advertisers would continually outbid each other by 1 cent, then one advertiser would drop her bid by a large amount (when she decided that competing for the top position was no longer worth it, and settling for a lower one would be better) – and then immediately her competitor would drop his bid as well (because he did not want to outbid her by more than 1 cent). This behavior led to a stress on GoTo's ad servers, as well as lowered ad auction revenues (Edelman and Ostrovsky, 2007).

*We are grateful to Sebastián Bauer for excellent research assistance.

†ostrovsky@stanford.edu. Stanford University and NBER.

‡skrz@stanford.edu. Stanford University.

When Google came out with their own sponsored search auction design in 2002, they introduced two major changes. First, instead of ranking ads purely based on the corresponding bids, they ranked them based on the product of each ad's bid and its *quality score*. Initially, quality score was equal to the estimated "clickability" of the ad, i.e., the probability that the ad would be clicked by a user if it was shown in the top position, although in later iterations it started including other factors (such as, e.g., the quality of the landing page). The second change was the introduction of a different payment rule. Instead of paying his own bid, each advertiser would pay the smallest amount he would have to bid to outbid its nearest competitor (given their respective quality scores). With identical quality scores, the payment rule reduced to simply paying the bid of the next highest bidder, and the format was dubbed the Generalized Second-Price Auction (Edelman et al., 2007; Varian, 2007).

The Generalized Second-Price Auction format (GSP) has become enormously successful, becoming the de facto industry standard. Beyond Google, it has been deployed by most of the search engines worldwide. It is used by Amazon for its sponsored products offering (which is a major part of the company's $31 billion per year, and rapidly growing, advertising business[1]), and many other online retailers and marketplaces use it to power their "Retail Media" advertising offerings. DoorDash and UberEats use it to allow restaurants to advertise on their apps and websites[2]. Apple's App Store and Google's Play Store use it to let developers to promote their apps. Travel sites like TripAdvisor, Expedia, Booking, and others use it to allow hotels to promote themselves to travelers. Job search sites like Indeed allow firms to post Sponsored Jobs that receive premium placement when job-seekers search for relevant terms. Instacart is using GSP to power its booming ad auction business.[3]

Thus, until recently, it may have looked like the GSP has "won," while GFP has fallen by the wayside, due to its poor equilibrium properties. In this paper, we show that this conclusion may be premature, and the situation is more nuanced. Specifically, we re-consider the question of equilibrium existence in the Generalized First-Price Auction when we include in the analysis a feature that is realistic (and common) in many real-world settings in which these types of auctions are used—stochastic quality scores. In the original Edelman and Ostrovsky (2007) paper, quality scores were not considered (or equivalently, were constant and equal to 1), but the paper's conclusions on the non-existence of pure-strategy equilibria would have remained largely the same even if the scores were different from 1, but were constant. In practice, however, the scores are stochastic, and may vary widely from one impression to the next (e.g., one ad may perform particularly well with women in New York, while another ad for the same search term may perform particularly well with men in San Francisco). Other real-world features may also make the environment stochastic (e.g., bidders' campaigns may be "paced" if their daily budgets are too low).[4] In such cases,

[1] `https://www.cnbc.com/2022/02/03/amazon-has-a-31-billion-a-year-advertising-business.html`

[2] `https://digiday.com/marketing/the-holy-grail-of-e-commerce-advertising-why-doordash-is-bolstering-its-advertising-offerings/`

[3] `https://www.wsj.com/articles/instacart-goes-deeper-into-digital-advertising-as-grocery-delivery-slows-11630920600`

[4] See Athey and Nekipelov (2012) and Pin and Key (2011) for discussions and empirical evidence of randomness

the GFP strategy "bid 1 cent higher than my opponent" is in general no longer a best response, and the payoff functions become continuous, smooth, and better behaved. As a result, in such environments, we show that the Generalized First-Price Auction may in fact have a pure strategy equilibrium. Moreover, perhaps even more surprisingly, we show that the expected revenue under that equilibrium may substantially exceed that of the Generalized Second-Price Auction.[5] Together with other attractive features of the Generalized First-Price Auction, such as that it is easy to explain to the advertisers (and their managers) and that it does not suffer from potential auctioneer credibility problems (Akbarpour and Li, 2020), our results suggest that it may deserve another look in some settings.

We start out by setting up the basic model to illustrate the above features (Section 2). The model is deliberately "stripped down" to its bare essence, to illustrate the key driving forces behind our results and make the derivations as transparent as possible. The simplicity of the model makes the revenue non-equivalence result particularly striking. After proving the existence and uniqueness of the equilibrium of GFP in this basic setting (Section 3) and comparing the revenues of GFP and GSP (Section 4), we explore two extensions of the basic model.

In Section 5, we vary the degree of uncertainty in quality scores, from that of the basic model to the limit case of no uncertainty. Consistent with the intuition, we find that as the amount of uncertainty decreases, the range of other parameter values for which the pure-strategy equilibrium of GFP exists shrinks, completely disappearing in the limit as uncertainty goes away. Revenue comparisons between GSP and GFP also become more subtle as the amount of uncertainty is reduced. These findings suggests that the generalized first-price auction might be particularly attractive in settings in which the amount of uncertainty in quality scores is high (such as, e.g., "display advertising" types of settings, in which user intent is not clear and the ad system is trying to predict user interests based on their highly variable characteristics), whereas the generalized second-price auction might be relatively more attractive in settings in which the amount of uncertainty in quality scores is lower (such as, e.g., in pure "sponsored search," in which ads are shown directly in response to user queries which clearly contain a lot of information and thus make the additional information about user characteristics relatively less important—thus making overall quality scores less variable). This comparative static is consistent with the current state of advertising at such search engines as Google, whose sponsored search system on Google.com (AdWords) is using the generalized second-price auction, while its ad system for external publishers, closer to display advertising (AdSense) has recently shifted to the generalized first-price auction.[6]

In Section 6, we consider another extension to the basic model, where we allow the bidders' per-click values to be heterogeneous. We find that even with this generalization, the equilibrium of the generalized first-price auction continues to exist. The revenue comparison between GSP and GFP, however, becomes more intricate. Section 7 concludes.

---

in quality scores and other auction parameters.

[5] In Section 4, where we present this comparison, we explain why this finding does not contradict the celebrated Revenue Equivalence Theorem of Myerson (1981), despite both formats resulting in identical allocations.

[6] `https://blog.google/products/adsense/our-move-to-a-first-price-auction/`, October 7, 2021.

2 Basic Model

Two advertisers are competing for two advertising slots on an internet page. Each values a click from a user at \$1. The two slots vary in terms of visibility to the user. The better slot has visibility normalized to 1 and the lesser slot has visibility $\alpha \in (0,1)$.

The bidders simultaneously submit bids $\{b_1, b_2\}$. Then a user arrives to the website. The website has information about user characteristics and the potential relevance of each ad to the user, and for each ad estimates the probability that the user will click on the ad conditional on noticing it. The probabilities for the two ads are (proportional to) $\{q_1, q_2\}$. Each q_i is drawn independently from a uniform distribution $[0,1]$.[7]

The auctioneer runs a generalized first-price auction (GFP) with scoring the bids by q_i's.

That is, the publisher after observing q_i's and bids, ranks the bids based on the scores:

$$S_i = b_i q_i.$$

The bidder with the highest score wins the top slot and, conditional on a click, pays b_i. The bidder with the second-highest score wins the lesser slot and pays its bid if the user clicks on his ad.

The expected profit of the winner and loser, conditional on bid b and realized q are:

$$\begin{aligned} U_W(q,b) &= q(1-b), \\ U_L(q,b) &= \alpha q(1-b). \end{aligned}$$

We are interested in characterizing the pure-strategy Nash equilibrium of this game.

3 Existence of a Pure Strategy Equilibrium

Our first main result is:

Theorem 1 *In the generalized first-price auction described above, there exists a unique symmetric pure-strategy equilibrium, with the equilibrium bids equal to*

$$b_1 = b_2 = \frac{2(1-\alpha)}{4-\alpha}.$$

Before we present the proof, we discuss the intuition. When q's are known (or not used at all as in the original auctions), there is no pure strategy equilibrium because at a tie, one of the bidders wants to either deviate to an ϵ higher bid or to zero. Unequal bids cannot be an equilibrium either since then the loser wants to deviate to bid zero, the winner to almost match him, but then the loser should "jump over."

[7] Since all payoffs in the model are proportional to the probability of a click, $q_i = 1$ does not mean that there are users with an estimated probability of click equal to 1. Instead, whatever is the highest probability of click, we normalize it to be 1, without loss of generality.

When bidders do not know q's, the probability of winning changes smoothly in bid. Deviating to an ϵ-higher bid changes the probability of getting the top position only slightly. When there is enough uncertainty about $\frac{q_i}{q_j}$ the best response becomes concave in own bid, and there exists a fixed point.

The proof uses the fact that with the distribution of q_i's having support starting at zero, the ratio of q's ranges from 0 to ∞. It is possible to generalize the existence result for less variation in that ratio, but the existence of a pure strategy equilibrium depends then on α. The smaller is α, the smaller variation in the ratio is needed for existence.

Proof. Since the optimization problem is symmetric for the two players, we consider only player 1. Conditional on the two bids being $\{b_1, b_2\}$, the expected profit of bidder 1 is

$$\begin{aligned} EU_1(b_1, b_2) &= (1 - b_1)E[q_1(\alpha + (1 - \alpha)\mathbf{1}_{b_1 q_1 > b_2 q_2})] \\ &= (1 - b_1)\int_0^1 q(\alpha + (1 - \alpha)h(b_1, b_2, q)))dq, \end{aligned} \tag{1}$$

where $h(b_1, b_2, q)$ is the probability the score of bidder 1 is higher than the score of bidder 2, given the bids and $q_1 = q$.

In this expression, $(1 - b_1)$ is the expected profit conditional on a click, $q\alpha$ is the probability of a click in position 2 and the additional probability of a click $q\,(1 - \alpha)$ is only conditional on winning the auction.

Player 1 wins whenever $b_1 q_1 > b_2 q_2$, so

$$h(b_1, b_2, q) = \min\left\{1, \frac{b_1}{b_2} q\right\}.$$

Simplifying the expressions for expected payoffs we get: if $b_1 \geq b_2$:

$$EU_1\left(b_1, b_2\right) = (1 - b_1)\left[\frac{\alpha}{2} + (1 - \alpha)\frac{1}{6}\left(3 - \left(\frac{b_2}{b_1}\right)^2\right)\right], \tag{2}$$

and if $b_1 \leq b_2$:

$$EU_1\left(b_1, b_2\right) = (1 - b_1)\left[\frac{\alpha}{2} + (1 - \alpha)\frac{1}{3}\frac{b_1}{b_2}\right]. \tag{3}$$

A necessary condition for equilibrium is that the FOC holds a the symmetric bidding profile $b_1 = b_2 = b$. By inspection, $EU_1(b_1, b_2)$ is differentiable in b_1 at symmetric bidding profiles and the derivative at such vectors is

$$\frac{\partial EU_1(b_1, b_2)}{\partial b_1} = \frac{2(1 - \alpha) - b(4 - \alpha)}{6b}.$$

The unique solution of the FOC is hence a unique candidate for the pure symmetric equilibrium:

$$\frac{\partial EU_1(b_1,b_2)}{\partial b_1}=0 \text{ and } b_1=b_2=b \iff b_1=b_2=\frac{2(1-\alpha)}{4-\alpha}.$$

To check that this is actually an equilibrium we need to check that these bids are global best responses. Direct inspection of (3) shows that $EU_1(b_1,b_2)$ is concave in b_1 (since it is quadratic with a negative coefficient on b_1). Similarly, for any $b_1 \leq 1$ the expression in (2) is concave in b_1 (note that $\frac{\partial^2 EU_1(b_1,b_2)}{\partial b_1^2}=(1-\alpha)(b_1-3)\frac{b_2^2}{3b_1^4}>0$ and bidding above b_1 is a dominated strategy). Hence, the FOC is both necessary and sufficient for the optimality of the best response of bidder 1. ■

Given the equilibrium bids b, the equilibrium payoff is

$$EU_1(b,b)=(1-b)\left(\frac{1}{6}\alpha+\frac{1}{3}\right)$$

The intuition is that the bidder gets payoff $(1-b)$ conditional on a click. They are half of the time the bidder with the higher score and half of the time with the lower score. Conditional on having the higher q_i, the expected q_i is $\frac{2}{3}$, which is the expected probability of a click in the top position. Conditional on having the lower q_i, the expected q_i is $\frac{1}{3}$, so the expected probability of a click on the lower position is $\alpha/3$. This yields the expression for the expected equilibrium number of clicks.

4 Revenue Comparison: GFP vs. GSP

Our second result compares bidding and revenues between the equilibrium of the GFP we characterized above and the equilibrium in a generalized second-price auction, GSP.

In the equilibrium we have found of the GFP the expected revenues are $b(\frac{\alpha}{3}+\frac{2}{3})$ because the expected q of the winner is $\frac{2}{3}$ and the expected q of the loser is $\frac{1}{3}$. Plugging in the equilibrium bidding strategy we get

$$REV_{GFP}=\frac{1}{3}(1-\alpha)\frac{2(\alpha+2)}{4-\alpha}. \tag{4}$$

In a GSP, the loser pays 0 and the winner i pays

$$P=b_j\frac{q_j}{q_i}.$$

That is the lowest amount the winner could have bid and still at least tied with the loser (with a bid $b_i=P$ the scores of the two bidders are equal).

With two ad slots and two bidders, the equilibrium bidding strategies in GSP are straightforward to find. Each bidder is guaranteed to win the second position and pay 0, so an equilibrium strategy

is to bid per click the full value of the incremental expected number of clicks from the upgrade to the top position:

$$b_{GSP} = 1 - \alpha.$$

To see this, suppose bidder 1 happens to know q_1, q_2, and $b_2 = (1 - \alpha)$. Then his options are:

1. Lose the auction, pay 0, and get the second position (and thus αq_1 clicks), for the expected payoff of αq_1; or

2. Win the auction, pay $b_2 q_2 = (1 - \alpha) q_2$ in expectation, and get q_1 clicks, for the expected payoff of $q_1 - (1 - \alpha) q_2$.

The second option is more attractive to bidder 1 whenever $q_1 > q_2$—and that is exactly the outcome that the bidder will get if he bids $b_1 = (1 - \alpha)$.

The expected revenue in the GSP can be computed as:

$$REV_{GSP} = 2 \int_0^1 q \left(\int_0^q \left((1 - a) \frac{q_2}{q} \right) dq_2 \right) dq = \frac{1}{3} (1 - \alpha)$$

This expression is intuitive: given the realized first and second order statistic of q, $q^{(1)}$ and $q^{(2)}$, respectively, only the winner pays and the expected payment of the winner per impression is

$$q^{(1)} * b \frac{q^{(2)}}{q^{(1)}} = b q^{(2)}.$$

The equilibrium bid is $b = (1 - \alpha)$ and the expected second-highest q is $\frac{1}{3}$.

Finally, observe that

$$REV_{GFP} = \frac{1 - \alpha}{3} \frac{2 (\alpha + 2)}{4 - \alpha} > \frac{1 - \alpha}{3} = REV_{GSP},$$

because $\frac{2(\alpha+2)}{4-\alpha} > 1$.

Summarizing the analysis above yields our second main result:

Theorem 2 *The expected revenue in the Generalized First-Price Auction is equal to*

$$REV_{GFP} = \frac{2(1 - \alpha)(\alpha + 2)}{3(4 - \alpha)}.$$

The expected revenue in the Generalized Second-Price Auction is equal to

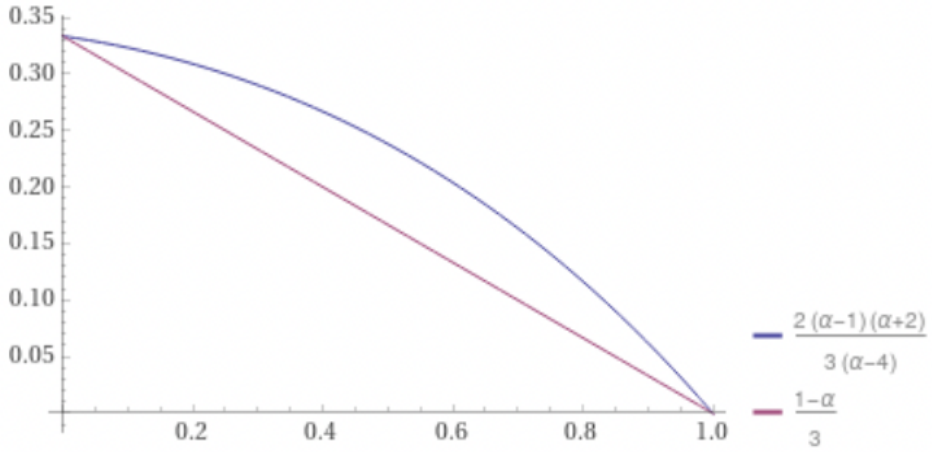
$$REV_{GSP} = \frac{1 - \alpha}{3}.$$

For all $\alpha \in (0, 1)$,

$$REV_{GFP} > REV_{GSP}.$$

Note that as α gets close to 0, the expected revenues from those two formats converge. But for all higher α the expected revenues are not the same, despite both formats resulting in efficient (and thus identical) equilibrium allocations. The difference can be substantial. For example, suppose $\alpha = 0.7$, which is a fairly typical "dropoff" value in ad auctions, meaning that the second slot gets 70% as many clicks as the first one. The the revenue in the generalized second-price auction is equal to 0.1, while the revenue for the generalized first-price auction is equal to $0.1 \cdot \frac{5.4}{3.3} \approx 0.164$—a difference of more than 60%. Figure 1 plots the revenues in the GFP and GSP for all values of α.

Figure 1: Revenues in GFP and GSP as a function of α



At first glance, this result seems to violate the Revenue Equivalence Theorem (RET) of Myerson (1981)—after all, in both formats, we end up with fully efficient (and thus identical) allocations: whichever bidder has the higher relevance q_i gets the top position (because in both GFP and GSP, equilibria are symmetric, and the bidders are submitting identical bids). Since the value per click of each bidder is the same (equal to 1), that is the efficient allocation. And yet the revenues are different!

Of course, there is no contradiction. The purely technical reason is that RET holds subject to the payoffs of the lowest types being equal. And since in our model each bidder has only one informational type (the value per click of each bidder is fixed at 1, and relevance q_i is not known to the bidder beforehand), which is thus automatically "the lowest," RET has no bite. More substantively, one could consider a slightly more general model, where with probability $1 - \varepsilon$, the bidder's value per click is 1, and with probability ε, it is, say, uniformly distributed on $[0, 1]$. In that case, the lowest type would be 0, and its payoff in both equilibria (of GSP and of GFP) would also be 0. But now the equilibrium in GSP will be efficient, while that in GFP will not be—and thus the allocation will not be identical. As a result, RET would not apply. Of course, as ε converges to zero, this model would converge to the basic one above.

5 Pure-Strategy Equilibria with Less Heterogeneity in Click-Through Rates

In our baseline model we have assumed that q_i's are drawn from a uniform distribution $[0,1]$. That implies that the range of possible ratios of $\frac{q_1}{q_2}$ is $[0,\infty]$. In this section we characterize necessary and sufficient conditions for the existence of a pure strategy equilibrium in the GFP when the range of possible ratios is less extreme.

To this end, assume that the probabilities of clocks are proportional to q_i which are drawn from a uniform distribution $[u,u+1]$, for some parameter $u \geq 0$. When $u=0$ we have as a special case our previous model and as u gets larger, there is less and less ex-ante uncertainty about the ratio of the CTRs.[8]

The rest of the model remains unchanged.

Expected Payoffs Given bids (b_1,b_2), the expected payoff of player 1 is still as in equation (1), but the expression for $h(b_1,b_2,q)$ changes. Player 1 wins whenever $b_1q_1 > b_2q_2$, so

$$h(b_1,b_2,q) = \min\left\{1, \max\left\{\frac{b_1}{b_2}q - u, 0\right\}\right\}.$$

That is, when the probability is interior, it is $\frac{b_1}{b_2}q - u$.

Next we simplify the range of bids we consider. First, bidding so high that $\frac{b_1}{b_2}(u+1) - u > 1$ is never optimal. Second, conditional on bidding such that the probability of winning the top position is zero, the optimal bid is 0. Hence, without loss of generality in solving for the best response we need to only consider bids in the set

$$b_1 \in \{0\} \cup \left[b_2\frac{u}{u+1}, b_2\frac{1+u}{u}\right].$$

Bidding 0 yields a payoff

$$EU_1(0,b_2) = \int_u^{u+1} q\alpha dq = \alpha\left(u+\frac{1}{2}\right).$$

Bidding in the range $\left[b_2\frac{u}{u+1}, b_2\frac{1+u}{u}\right]$ yields payoff:

[8] Note that we can normalize the CTRs by dividing them by $u+1$, so that the model is equivalent to assuming that the distribution of the CTRs is uniform $[\frac{u}{u+1}, 1]$. As $u \to 1$, the uncertainty about the CTRs disappears.

$$
\begin{aligned}
\text{If } b_1 &\geq b_2 \\
EU_1(b_1, b_2) &= (1-b_1)\left[\alpha\left(u+\frac{1}{2}\right)+(1-\alpha)\left(\int_u^{(u+1)\frac{b_2}{b_1}} q\left(\frac{b_1}{b_2}q-u\right)dq+\int_{(u+1)\frac{b_2}{b_1}}^{u+1} q dq\right)\right]. \\
\text{If } b_1 &\leq b_2 \\
EU_1(b_1, b_2) &= (1-b_1)\left[\alpha\left(u+\frac{1}{2}\right)+(1-\alpha)\left(\int_{u\frac{b_2}{b_1}}^{u+1} q\left(\frac{b_1}{b_2}q-u\right)dq\right)\right].
\end{aligned}
$$

Simplifying the integrals we get:

$$
\begin{aligned}
\text{If } b_1 &\geq b_2 \\
EU_1(b_1, b_2) &= (1-b_1)\left[\alpha\left(u+\frac{1}{2}\right)+(1-\alpha)\frac{1}{6}\left(3u^3+3(u+1)^2-2u^3\frac{b_1}{b_2}-\left(\frac{b_2}{b_1}\right)^2(u+1)^3\right)\right]. \\
\text{If } b_1 &\leq b_2 \\
EU_1(b_1, b_2) &= (1-b_1)\left[\alpha\left(u+\frac{1}{2}\right)+(1-\alpha)\frac{1}{6}\left(-3u(u+1)^2+2(u+1)^3\frac{b_1}{b_2}+\left(\frac{b_2}{b_1}\right)^2 u^3\right)\right].
\end{aligned}
$$

When $b_1 = b_2$ the profits simplify to

$$
EU_1(b_1, b_1) = (1-b_1)\left[\frac{1}{2}u(\alpha+1)+\frac{1}{6}\alpha+\frac{1}{3}\right]
$$

FOC and Candidate Equilibrium We start with a necessary condition for a pure strategy equilibrium in GFP is that at the candidate equilibrium the first-order conditions are satisfied. Differentiating the expected profit functions above with respect to b_1 and solving for them at $b_1 = b_2 = b$ we get that the unique candidate for the equilibrium is

$$
b_{GFP} = (1-\alpha)\frac{6u(u+1)+2}{3u(1+2u)(1-\alpha)+6u+4-\alpha}. \tag{5}
$$

It can be shown that this candidate is between 0 and 1 for all u and α.

Necessary and Sufficient Conditions for Existence We now provide necessary and sufficient conditions for the candidate b_{GFP} to be indeed the unique pure-strategy equilibrium of the GFP.

The first necessary condition (and as we show later a sufficient as well!) is that at the bidding profile no bidder wants to deviate to bidding zero. Bidding zero yields payoff $\alpha(u+\frac{1}{2})$ since the average q_i is $(u+\frac{1}{2})$. Therefore the necessary condition is:

$$
\alpha\left(u+\frac{1}{2}\right) \leq (1-b_{GFP})\left[\frac{1}{2}u(\alpha+1)+\frac{1}{6}\alpha+\frac{1}{3}\right]. \tag{6}
$$

Substituting our candidate b_{GFP} we can solve this condition for the critical level of α for any

u: this necessary condition is satisfied when $\alpha \leq \alpha^*$ where

$$\alpha^* = \frac{9u^2 + 12u + 4}{36u^3 + 45u^2 + 21u + 4} \in [0, 1] \tag{7}$$

When $u = 0$ as in our baseline model, $\alpha^* = 1$ and hence this condition is always satisfied. When the heterogeneity/uncertainty of CTRs is smaller, α^* gets smaller and in the limit as u gets large, α^* converges to zero (consistent with the non-existence result when q_i's are known).

The second necessary condition is that the bidder does not want to deviate to a bid that wins for sure, $b = b_{GFP}\frac{1+u}{u}$. Such a bid yields expected payoff $(1 - b_{GFP}\frac{1+u}{u})(u + \frac{1}{2})$. This necessary condition is hence

$$(1 - b_{GFP}\frac{1+u}{u})(u + \frac{1}{2}) \leq (1 - b_{GFP}) \left[\frac{1}{2}u\left(\alpha + 1\right) + \frac{1}{6}\alpha + \frac{1}{3}\right]. \tag{8}$$

Straightforward (though tedious) algebra shows that this condition is satisfied for all u and α

The only remaining condition to check if both players bidding b_{GFP} is a mutual best response is to check if there is a bid b that results in an interior probability of winning that is profitable. To show that not, we compute the second derivative of the expected profit function for $b \in \left[b_2\frac{u}{u+1}, b_2\frac{1+u}{u}\right]$.

$$\begin{aligned}
\frac{\partial^2}{\partial b_1^2} EU_1\left(b_1, b_2\right) &= \frac{1-\alpha}{3b_1^4 b_2}\left(\left(2u^3\right)b_1^4 + \left(u+1\right)^3\left(b_1 - 3\right)b_2^3\right) \text{ for } b_1 \geq b_2 \\
\frac{\partial^2}{\partial b_1^2} EU_1\left(b_1, b_2\right) &= -\frac{1-\alpha}{3b_1^4 b_2}\left(2\left(u+1\right)^3 b_1^4 + u^3 b_2^3\left(b_1 - 3\right)\right) \text{ for } b_1 \leq b_2
\end{aligned}$$

The details of these expressions do not matter, just that the the numerators change sign at most once. That leads to the following observations: In the range $b_1 \leq b_{GFP}$ the expected profit function is either concave or convex or switches once from convex to concave. In the first case (concave), there is no profitable deviation. In the second and third case (convex or first convex and then concave), if there is a profitable deviation then deviating to the bottom of the range is also profitable. But then deviating to bidding zero would be even more profitable. So if deviating to bidding zero is not profitable, deviating to any other bid below b_{GFP} is not profitable either.

Similarly, in the range $b_1 \geq b_{GFP}$ the expected profit function is also either everywhere concave, everywhere convex or first concave and then convex. That means that there is either no profitable deviation or deviating to a bid that wins with probability 1 is profitable. But we have already argued that deviating from the candidate b_{GFP} to a bid that always wins is never profitable.

Combining all these observations we get:

Theorem 3 *With click probabilities proportional to q_i distributed uniformly $[u, u+1]$, there exists a unique symmetric pure-strategy equilibrium in the generalized first-price auction if and only if:*

$$\alpha \leq \alpha^* = \frac{9u^2 + 12u + 4}{36u^3 + 45u^2 + 21u + 4} \in [0, 1].$$

When this condition is satisfied, the equilibrium bids are equal to:

$$b_1 = b_2 = (1-\alpha)\frac{6u(u+1)+2}{3u(1+2u)(1-\alpha)+6u+4-\alpha}.$$

5.1 Revenue Comparisons

We can also extend our analysis of the revenue comparisons between GFP and GSP to this more general range of u. In the GSP, it is still optimal for bidders to bid $b_{GSP} = (1-\alpha)$ since this is the value of improving from position 2 to position 1. As before, only the winner pays and since they pay the minimum amount they could have bid and still win, the expected revenue is

$$REV_{GSP} = (1-\alpha)(u+\frac{1}{3}).$$

The expected revenue in the GFP is

$$REV_{GFP} = b_{GFP}\left(\alpha(u+\frac{1}{3}) + u + \frac{2}{3}\right)$$

because the expected number of clicks in position 2 is $\alpha(u+\frac{1}{3})$ and the expected number of clicks in position 1 is $u+\frac{2}{3}$. Substituting the expression for b_{GFP} we get that

$$REV_{GSP} - REV_{GFP} = (1-\alpha)\,\frac{u\,(u+1) - (3u+1)\left(3u+4u^2+1\right)\alpha}{(6u^2+9u+4) - (6u^2+3u+1)\,\alpha}$$

The denominator is always positive, so the sign depends only on the sign of the numerator. That yields the following result:

Theorem 4 *Suppose click probabilities are proportional to q_i distributed uniformly $[u, u+1]$. Suppose $\alpha \leq \alpha^*$ defined in (7). (The pure strategy equilibrium) expected revenues in the GFP are higher than in the GSP if and only if*

$$\alpha > \hat{\alpha} = \frac{u\,(u+1)}{(3u+1)\,(3u+4u^2+1)}.$$

Note that when $u = 0$, $\hat{\alpha} = 0$ and we get our previous result that the GFP yields a higher expected revenue for all α. However, for $u > 0$, $\hat{\alpha}$ is positive. Hence, for small α the GSP yields higher revenues than the GFP. Moreover we can verify that for all $u > 0$, $\hat{\alpha} < \alpha^*$ so that for every $u > 0$ the range $\alpha \in [\hat{\alpha}, \alpha^*]$ is non-empty. In that range the GFP yields higher expected revenues than the GSP. In other words, GSP tends to dominate when the second position is much weaker than the first, but when they get closer in value (but not too close for the GFP equilibrium to disappear), GFP starts to dominate.[9]

[9]The difference in expected revenues between GSP and GFP is monotone in α when GSP dominates, but can be non-monotone when GFP dominates, but the difference crosses zero exactly once.

6 Heterogeneous Bidder Values

Another natural question is to what extent our results remain valid if the per-click values of the two bidders are different. To answer this question, we go back to the basic model in which the quality scores are distributed uniformly on $[0,1]$, but now assume that the value of the second bidder, v_2, is equal to some value $v \in (0,1)$, while the value of the second bidder, v_1, is still equal to 1.

Notably, the existence result continues to hold.

Theorem 5 *In the generalized first-price auction described above, there exists a unique pure-strategy equilibrium.*

Proof. As before, conditional on the two bids being $\{b_1, b_2\}$, the expected profit of bidder 1 is:

$$\begin{aligned} EU_1(b_1,b_2) &= (1-b_1)\,E\left[q_1\left(\alpha+(1-\alpha)1_{b_1q_1>b_2q_2}\right)\right] \\ &= (1-b_1)\int_0^1 q\left(\alpha+(1-\alpha)h_1(b_1,b_2,q)\right)dq, \end{aligned}$$

whereas the profit for bidder 2 is:

$$\begin{aligned} EU_2(b_1,b_2) &= (v-b_2)\,E\left[q_2\left(\alpha+(1-\alpha)1_{b_1q_1<b_2q_2}\right)\right] \\ &= (v-b_2)\int_0^1 q\left(\alpha+(1-\alpha)h_2(b_1,b_2,q)\right)dq. \end{aligned}$$

As before, if we simplify the expressions for expected payoffs, we get that if $b_1 \geq b_2$:

$$EU_1(b_1,b_2) = (1-b_1)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{6}\left(3-\left(\frac{b_2}{b_1}\right)^2\right)\right] \tag{9}$$

$$EU_2(b_1,b_2) = (v-b_2)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{3}\frac{b_2}{b_1}\right], \tag{10}$$

whereas if $b_1 \leq b_2$:

$$EU_1(b_1,b_2) = (1-b_1)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{3}\frac{b_1}{b_2}\right] \tag{11}$$

$$EU_2(b_1,b_2) = (v-b_2)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{6}\left(3-\left(\frac{b_1}{b_2}\right)^2\right)\right]. \tag{12}$$

If in equilibrium $b_1 \geq b_2$ (as will be the case), then equations (9) and (10) are the relevant ones. Advertiser 1's problem is:

$$\max_{b_1}(1-b_1)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{6}\left(3-\left(\frac{b_2}{b_1}\right)^2\right)\right]$$

with FOC:

$$-\frac{\alpha}{2}-(1-\alpha)\frac{1}{6}\left(3-\left(\frac{b_2}{b_1}\right)^2\right)-(1-b_1)(1-\alpha)\frac{1}{6}b_2^2(-2)\frac{1}{b_1^3}=0$$
$$-\frac{b_1^3}{2}+b_2^2\left(\frac{1-\alpha}{3}\right)\left(1-\frac{b_1}{2}\right)=0. \tag{13}$$

Advertiser 2's problem is:

$$\max_{b_2}(v-b_2)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{3}\frac{b_2}{b_1}\right]$$

with FOC:

$$-\frac{\alpha}{2}-(1-\alpha)\frac{1}{3}\frac{b_2}{b_1}+\frac{(1-\alpha)}{3b_1}(v-b_2)=0$$
$$\frac{v}{2}-b_2-\frac{\alpha}{1-\alpha}\frac{3}{4}b_1=0. \tag{14}$$

Equations (13) and (14) fully characterize the equilibrium. Note that equation (14) is a simple linear equation in b_1 and b_2, and using it to express one of these variables in terms of the other, and plugging the resulting expression into equation (13) gives us a cubic equation. Straightforward (though cumbersome) calculations show that it has a unique positive root (details are available upon request). Since both expected utilities are concave in the respective bids for the relevant values, it follows that there exists a unique equilibrium such that $b_1 \geq b_2$.

To rule out the cases in which $b_1 < b_2$ as an equilibrium, we repeat the above computation for equations (11) and (12), and find that for all possible values of $v < 1$ and all possible values of α, all candidate solutions have $b_1 > b_2$, leading to a contradiction (again, the straightforward yet cumbersome calculations are available upon request). This completes the proof. ■

Given equilibrium values (b_1^*, b_2^*), we can compute the equilibrium utilities:

$$EU_1(b_1^*, b_2^*) = (1-b_1^*)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{6}\left(3-\left(\frac{b_2^*}{b_1^*}\right)^2\right)\right]$$
$$EU_2(b_1^*, b_2^*) = (v-b_2^*)\left[\frac{\alpha}{2}+(1-\alpha)\frac{1}{3}\frac{b_2^*}{b_1^*}\right]$$

and the expected revenue for the GFP will be:

$$REV_{GFP} = \int_0^1 \int_{q_2\frac{b_2^*}{b_1^*}}^1 (q_1 b_1^* + \alpha q_2 b_2^*)\, dq_1 dq_2 + \int_0^1 \int_0^{q_2\frac{b_2^*}{b_1^*}} (q_2 b_2^* + \alpha q_1 b_1^*)\, dq_1 dq_2$$
$$= \frac{b_1^*}{2} - \frac{(b_2^*)^2}{6b_1^*} + \frac{\alpha b_2^*}{2} - \frac{\alpha (b_2^*)^2}{3b_1^*} + \frac{(b_2^*)^2}{3b_1^*} + \frac{\alpha (b_2^*)^2}{6b_1^*}.$$

Now, for GSP, the bidding strategies are still straightforward to find. Specifically, it is easy to

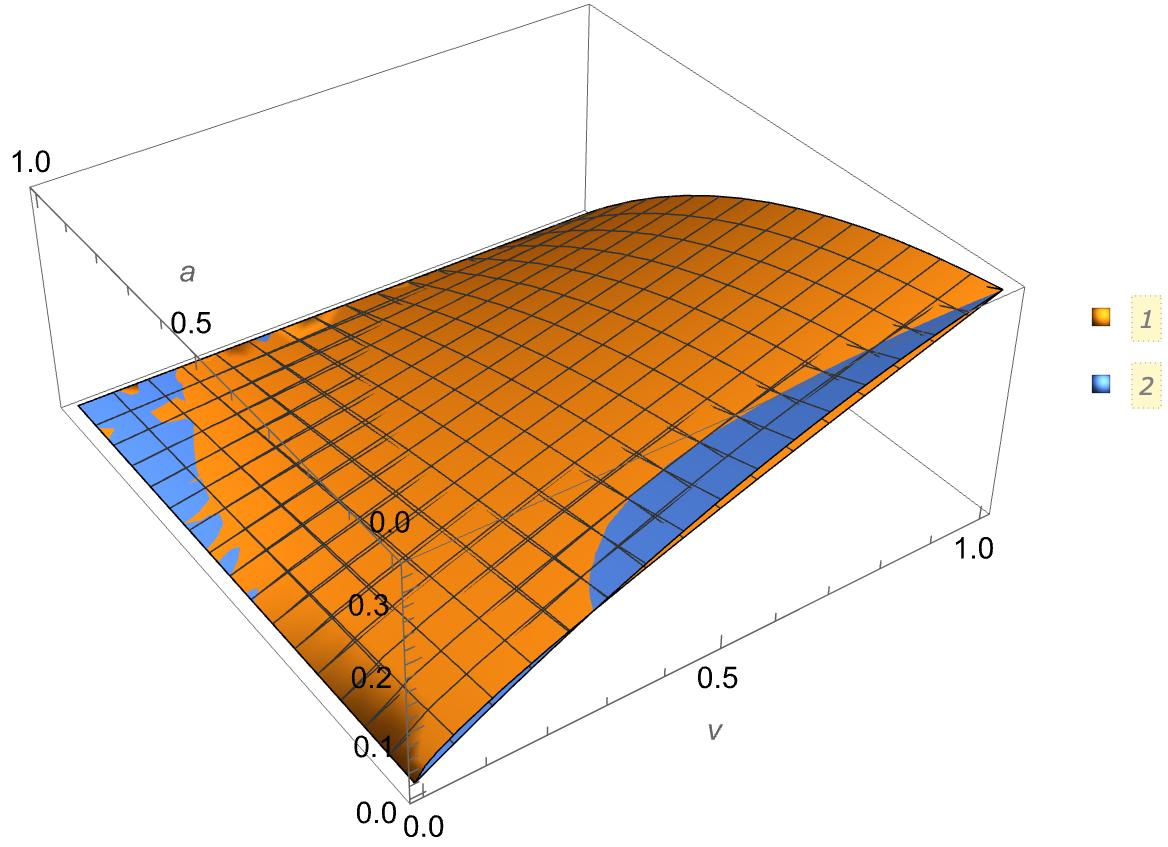
see that advertiser 1 should bid $b_1 = 1-\alpha$ as before, whereas advertiser 2 should bid $b_2 = v(1-\alpha)$. The argument is identical to the one in Section 4.

The expected revenue in the GSP will be given by:

$$\begin{aligned} REV_{GSP} &= \int_0^1 \int_{q_2 v}^1 (1-\alpha) v \frac{q_2}{q_1} q_1 dq_1 dq_2 + \int_0^1 \int_0^{q_2 v} (1-\alpha) \frac{q_1}{q_2} q_2 dq_1 dq_2 \\ &= \frac{(1-\alpha)v}{2} - \frac{(1-\alpha)v^2}{6}. \end{aligned}$$

The relationship between the revenues in GFP and GSP is now more complex. Figure 2 plots the revenues in the two auction formats for the values of bidder 2's value v and the visibility of the second slot α between 0 and 1. As we saw earlier, when $v = 1$, the revenue in GFP is always weakly greater than that in GSP. However, for lower values of v that is not always the case. For instance, for relatively high (but lower than 1) values of v and relatively low values of α, the revenue in GSP is higher, although the difference is small.

Figure 2: Revenues in GFP and GSP as a function of α and v



7 Conclusion

In this paper, we revisit the classic result on the (non-)existence of pure-strategy Nash equilibria in the Generalized First-Price Auction and show that the conclusion may be reversed when ads are ranked based on a product of stochastic quality scores and bid amounts, as is commonly done in practice—provided the uncertainty in these quality scores is large. Moreover, the expected revenue in the pure strategy equilibrium of the Generalized First-Price Auction may substantially exceed that of the Generalized Second-Price Auction.

Our model is deliberately simple and streamlined, to illustrate the essential driving forces behind our results. Of course, in practice, sponsored search auctions usually have more than two positions and more than two bidders; distributions of quality scores can be quite complex; and other practical features may matter for the results. Nevertheless, the overall intuition of stochastic quality scores "smoothing out" payoff functions in the Generalized First-Price Auction and thus making the existence of a pure-strategy equilibrium more likely, will continue to hold in these more general settings. Combined with other desirable features of the Generalized First-Price Auction that we mentioned in the Introduction, we believe that our conclusions suggest that it may deserve a careful consideration in many practical settings.

References

Akbarpour, M. and S. Li (2020). Credible auctions: A trilemma. *Econometrica 88*(2), 425–467.

Athey, S. and D. Nekipelov (2012). A structural model of sponsored search advertising auctions. working paper.

Edelman, B. and M. Ostrovsky (2007, 02). Strategic bidder behavior in sponsored search auctions. *Decision Support Systems 43*, 192–198.

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

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Pin, F. and P. Key (2011). Stochastic variability in sponsored search auctions: Observations and models. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, New York, NY, USA, pp. 61–70. Association for Computing Machinery.

Varian, H. R. (2007). Position auctions. *International Journal of industrial Organization 25*(6), 1163–1178.