Dynamics of the Second Price

Julian Romero*and Eric Bax†

October 17, 2008

Abstract

Many auctions for online ad space use estimated offer values and charge the winner based on an estimate of the runner-up offer value. By awarding the ad space to the winner, the auctioneer produces more information about the value of the winning offer but not about the runner-up offer. Since price is based on an estimate of the runner-up offer value, failing to improve the estimate can harm revenue for the publisher and auctioneer. Also, failing to improve the estimate of the runner-up value can cause a less valuable offer to be repeatedly selected over a more valuable one, harming fairness for advertisers. This paper explores the dynamics of the second price and shows that learning the second price can improve revenue and fairness.

1 Introduction

This paper focuses on auctions to place ads in online content. Each opportunity to show an ad is called an *impression*. Advertisers buy impressions; publishers sell impressions. An exchange manages the auctions. The exchange stores bids and selection criteria for impressions from advertisers. The exchange also stores selection criteria for ads from publishers. The exchange hosts an auction for each impression offered by a publisher on the exchange. In the auction for each impression, the exchange uses selection criteria from advertisers and publishers to determine which ads are eligible for the ad slot, rewards the ad slot to an eligible ad that maximizes expected revenue for the publisher, and sets the price for the impression. The winning advertiser pays the price, which is usually partitioned between the publisher and the exchange provider. For background on auction theory and designing auctions, refer to Krishna (2002) and Milgrom (2004) For more information on auctions for internet advertising see Edelman, Ostrovsky & Schwarz (2007), Feng, Bhargava & Pennock (2004), and Lahaie (2006).

Most auctions to place ads in online content allow performance-based pricing, where the advertiser pays the price only if displaying the ad leads to an action. Otherwise the advertiser pays nothing. So, to compute expected revenue for each ad, the auctioneer

*Division of the Humanities and Social Sciences, Mail Code 228-77, California Institute of Technology, Pasadena, CA 91125. jnr@hss.caltech.edu

†Yahoo! Marketplace Design, PO Box 60543, Pasadena, CA 91116-6543. ebax@yahoo-inc.com

multiplies advertiser bid by probability of action. Since the actual probability of action is unknown, the auctioneer estimates it.

Most online auctions are second-price auctions, meaning that the charge to the winner is based on the estimated value of a runner-up offer. In a second-price auction, misestimating the value of the runner-up offer can harm revenue and fairness. If the value of the runner-up offer is underestimated, then the auctioneer under-charges the winning advertiser, harming revenue. If the value of the runner-up offer is overestimated, then either the auctioneer over-charges the winning advertiser, or, if the value of the runner-up offer is estimated to be greater than the value of the most valuable offer, the auctioneer fails to award the auction to the most valuable offer.

When an ad wins an auction, the ad is displayed. By observing whether the displayed ad prompts an action, the exchange gathers information to update its estimate of the value of the winning offer. However, since the ad for the runner-up offer is not displayed, the auction produces no new information about the value of the runner-up. This causes misestimates of the value of the runner-up offer to persist.

This paper is organized as follows. Section 2 defines an auction and a sequence of auctions, with estimates of the values of winning offers updated as the winner's ads are shown. Section 3 examines the dynamics of the second price over the sequence of auctions. Section 4 describes an asymmetry in the evolution of the second price that causes underestimates of the second price to persist longer than overestimates, reducing revenue for publishers and the auctioneer. Section 5 analyzes the impact on revenue and fairness of using some impressions to learn the second price. Section 6 concludes with a discussion of possibilities for future work.

2 Definitions and Notation

This section introduces an auction for impressions and a process for holding a sequence of auctions. The sequence of auctions is the basis for exploring the dynamics of the second price later in this paper. Also, this section defines the notions of revenue and fairness used later to analyze the effectiveness of learning the second price.

2.1 Auction

The auction outlined here is second-price auction with a single winner, making it like auctions for display advertising. In contrast, auctions for text advertising (which includes search advertising) produce a ranked list of multiple winning offers rather than a single winner. Though we limit our analysis to single-winner auctions, the observations about dynamics and the solutions also apply to multi-winner auctions, where the unlearned prices are for the offers not ranked high enough to be winners. For more information about the effect of multiple winners on strategy and pricing, refer to Lahaie (2006).

The auction selects as the winner an offer that maximizes expected revenue per impression. The auction is second-price; a runner-up offer determines the price charged to the winner. The term *offer* refers to an advertiser bidding to show an advertisement on an impression.

Using expected revenue as a basis for comparison allows advertisers with different price types to participate in the same auction. Common price types include CPM (cost per *mille*, or thousand impressions), CPC (cost per click), and CPA (cost per action). Expected revenue per thousand impressions is sometimes called eCPM.

Let b_i be the bid for offer i:

- For a CPM offer, b_i is the price the advertiser is willing to pay to show its ad.
- For a CPC offer, b_i is the price the advertiser is willing to pay for a click on its ad.
- For a CPA offer, b_i, is the price the advertiser is willing to pay for an action/conversion.

Let p_i be the predicted probability of a payout for offer i:

- For a CPM offer, p_i is 1 .
- For a CPC offer, p_i is the predicted probability of a click.
- For a CPA offer, p_i is the predicted probability of an action/conversion.

We refer to p_i as the *probability of action*, using the term *action* in a general sense, to refer to a display for a CPM offer, a click for a CPC offer, and an action/conversion for a CPA offer.

The (predicted) expected payout for offer i is

$$e_i \equiv p_i b_i.$$

Let

$$w = \underset{i}{\operatorname{argmax}}\ e_i.$$

Then offer w wins the auction. (In case of a tie, select w uniformly at random from indices of tied expected payouts.) Let

$$s = \underset{i \neq w}{\operatorname{argmax}}\ e_i.$$

Call offer s the second-place offer. The second price is $p_s b_s$. The charge for the winning offer is

$$a_w = \min\left(\frac{p_s b_s}{p_w} + \varepsilon, b_w\right),$$

where ε is the minimum bid increment, usually \$0.01. If offer w is a CPM offer, then the advertiser is charged a_w. For a CPC offer, the advertiser is charged a_w if the ad is clicked. For a CPA offer, the advertiser is charged a_w if showing the ad leads to a conversion. For our analysis and results, we will ignore ε, treating the second price charge for the winning offer as

$$\frac{p_s b_s}{p_w}.$$

Let p_w^* be the actual probability of action for which p_w is an estimate. Then the expected revenue from the auction is

$$r = a_w p_w^*.$$

The actual revenue is

$$\hat{r} = \begin{cases} a_w & \text{if the impression elicits an action} \\ 0 & \text{otherwise} \end{cases}.$$

In most cases, the auction provider and the publisher get a fixed share of the revenue. So both have an interest in increasing revenue.

While increasing revenue is important, it is also important to award impressions to the offer with the highest value. Define

$$w^* = \underset{i}{\operatorname{argmax}}\ b_i p_i^*$$

to be the index of the *ideal winner* - the offer with the highest value. Similarly, define

$$s^* = \underset{i \neq w^*}{\operatorname{argmax}} b_i p_i^*$$

to be the index of the ideal runner-up. Call an auction *fair* if $w = w^*$. The ideal expected revenue, r^*, is the expected revenue if actual probabilities are used in the auction instead of estimates:

$$r^* = \frac{p_{s^*}^* b_{s^*}}{p_{w^*}^*} p_{w^*}^* = p_{s^*}^* b_{s^*}.$$

2.2 Sequence of Auctions

Now we define a sequence of auctions with learning for auction winners. The same set of offers participate in all auctions in the sequence. For each offer i in each auction, let n_i be the number of impressions previously awarded to offer i, and let k_i be the number of actions resulting from those impressions. The estimated probability of action for each offer is the observed frequency of action:

$$p_i = \frac{k_i}{n_i}.$$

Assume each offer is awarded a number, n_{i0}, of "learning" impressions before the sequence of auctions begins, in order to observe the number of resulting actions, k_{i0}, to form an initial estimate of the probability of action for the offer. (For a more sophisticated approach to forming initial estimates, see Richardson, Dominowska & Ragno (2007) and Regelson & Fain (2006).)

The sequence of auctions proceeds as follows:

- Use the auction described in the previous subsection to ultip determine the index w of the winning offer and the index s of the runner-up.
- Award the impression to offer w, and increment n_w.
- If an action occurs as a result of the impression, increment k_w and charge the winner the estimated value of the second price, as specified in the auction definition in the previous subsection.

The auction process described here is idealized in the sense that it assumes instant feedback on user actions and incorporation of that data into estimates of values for winners. In practice, estimates are usually updated based on batches of data collected over time periods on the order of an hour or a day.

In a sequence of auctions, append a subscript to indicate a value after the auction indicated by the subscript. For example, use n_{it} to denote the number of impressions awarded to offer i before auction t in a sequence. Define the fairness of the first t auctions in a sequence to be the fraction of auctions with $w = w^*$:

$$F_t = \frac{\sum_{j=1}^{t} I\left(w_j = w^*\right)}{t},$$

where $I(\cdot)$ is the indicator function - one for a true argument and zero otherwise.

Similarly, define R_t to be the average expected revenue:

$$R_t = \frac{\sum_{j=1}^{t} r_j}{t},$$

define $\hat{R}_t$ to be the average actual revenue:

$$\hat{R}_t = \frac{\sum_{j=1}^{t} \hat{r}_j}{t},$$

and define R_t^* to be the (average) ideal revenue:

$$R_t^* = r^*.$$

3 Second Price is Monotone Decreasing

In the process outlined in the previous section, the estimate of the winner's value fluctuates as the winner wins auctions. The estimate of the runner-up's value remains steady until the estimate of the winner's value goes below the estimate of the runner-up's value. When this happens, the runner-up becomes the new winner, and the estimated second price decreases, because it is the estimate of the old winner's value. As a result, the estimated second price is non-increasing. In essence, the estimated second price is the minimum of a kind of random walk - a relay walk in which winners hand off the baton to runners-up when their estimated values cross.

Figure 1 shows an example of these dynamics. The figure shows estimated offer values for four offers over the course of a few hundred auctions. All auctions have the same \$1 bids, the same actual probabilities of action $p_i^* = 0.05$, and the same number of initial learning impressions $n_{i0} = 1000$. They have different numbers of initial actions, with $k_{10} = 52, k_{20} = 51, k_{30} = 50$, and $k_{40} = 49$, resulting in slightly different initial estimates for probability of action and value.

At the beginning, offer 1 wins each auction. As it wins more auctions and fails to elicit actions, its estimated value decreases until it goes below the estimated value for offer 2. Then offer 2 wins and fails to elicit an action. As a result, offer 1 becomes the winner again. Offers 1 and 2 alternate winning the auction and failing to elicit actions, driving down the first and second price until they fall below the third price. Then offers 1, 2, and 3 alternate being the winner until one of the offers elicits an action, causing its estimated value to increase. The estimated value of this winner fluctuates for almost 100 auctions, alternating between small decreases when it fails to elicit actions and large increases when it elicits actions. Meanwhile, the second price remains steady, until the fluctuation of the winner's estimated value brings it below the estimates of the runner-up offers again.

Note that while the first price fluctuates, the second price either remains steady or decreases throughout the sequence of auctions. Also note that when more offers are alternating at winning the auction, the second price falls more slowly. (Observe that the slope of the first price line decreases more slowly when more offers are involved.)

Figure 1: Example of Multiple Bidder Dynamics.

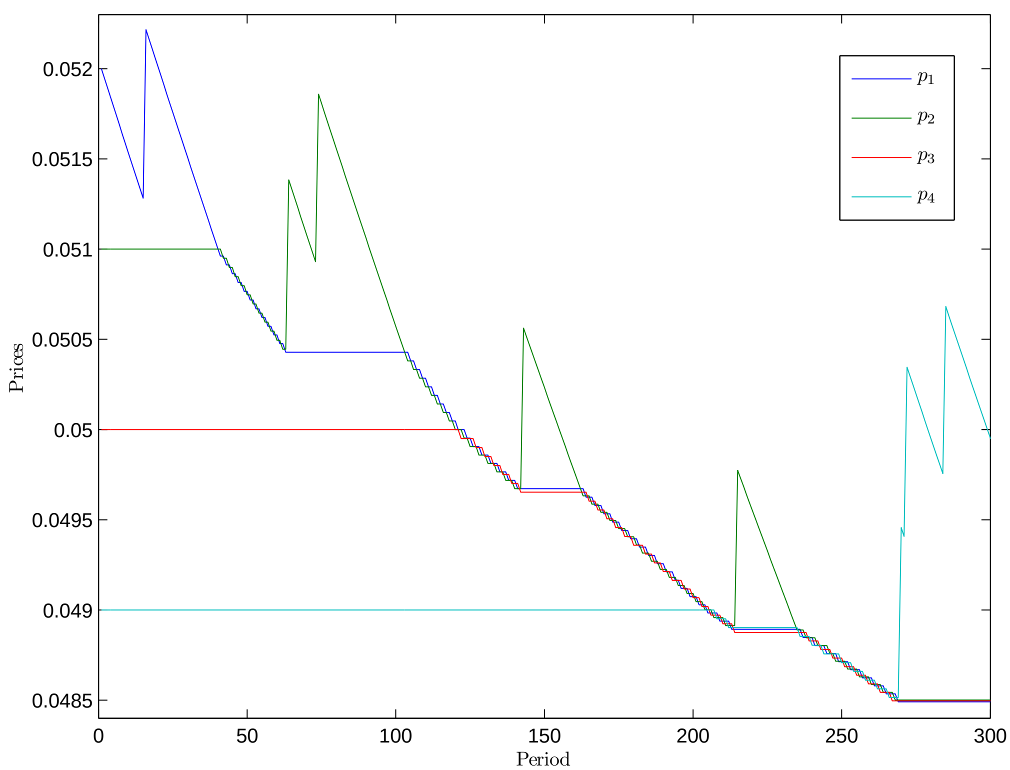

Table 1 shows the effect of this behavior on the second price. Each row shows the average second price at auctions $1000, 5000, \ldots, 100{,}000$ over 1000 simulated sequences of auctions. For each sequence, $p_1^* = p_2^* = 0.05$, $n_{10} = n_{20} = 1000$, k_{10} and k_{20} are shown in the first two columns, and $b_1 = b_2 = \$1$. Second price values are shown in cents. The ideal expected revenue, r^*, is 5¢.

Table 1: Average Second Price

		t					
k$_1$	**k$_2$**	**1000**	**5000**	**10000**	**20000**	**50000**	**100000**
45	45	4.3649	4.351	4.3506	4.3506	4.3506	4.3506
47	47	4.5479	4.5182	4.5165	4.5163	4.5163	4.5163
49	49	4.711	4.657	4.6489	4.6468	4.6459	4.6459
50	50	4.7991	4.7208	4.7064	4.7015	4.6988	4.6984
51	51	4.8823	4.7856	4.7671	4.7562	4.7504	4.7482
53	53	5.0426	4.9027	4.8627	4.838	4.8198	4.8138
55	55	5.1897	4.9877	4.9297	4.89	4.8655	4.8554
49	51	4.7811	4.7204	4.7092	4.7051	4.7033	4.703
47	53	4.6783	4.6491	4.6451	4.6443	4.6441	4.6441
45	55	4.4983	4.4918	4.4913	4.4913	4.4913	4.4913

Note that

- The estimated second value is always decreasing across each row, over the sequences of auctions. Even when both probabilities are underestimated initially, with $k_{10} = k_{20} = 45$, the estimated second value decreases over time.
- Even when both initial probabilities p_{10} and p_{20} are overestimated, with $k_{10} = k_{20} = 55$, the second value becomes underestimated on average at 5000 auctions and beyond.

Table 2 shows average actual revenues $\hat{R}_t$ in cents for the sequences used to generate Table 1. Note that

- Revenues are below the ideal expected revenue of 5¢. Revenues are generally lower when offers have lower initial value estimates.
- Unlike the estimated second value, revenue is not always decreasing. Revenue depends on p_w as well as the second price, with lower p_w generating more revenue.
- Average revenues converge over the sequences of auctions.

Tables 3 and 4 show second prices and revenues for different numbers of bidders, from 2 to 10. Each offer has $p_i^* = 0.05$ and $b_i = \$1$. For each offer, $n_{i0} = 1000$ and $k_{i0} = 50$, so $p_{i0} = p^*$. Each row shows second prices averaged over 1000 sequences of auctions, at the times t specified at the top of each column.

Table 2: Average Revenue

		t					
$\mathbf{k_0^1}$	$\mathbf{k_0^2}$	**1000**	**5000**	**10000**	**20000**	**50000**	**100000**
45	45	4.65161	4.47483	4.42491	4.39632	4.37398	4.36347
47	47	4.72505	4.56547	4.5423	4.52674	4.5239	4.52118
49	49	4.72959	4.65453	4.63366	4.62712	4.63063	4.63341
50	50	4.74069	4.67109	4.66794	4.67352	4.68087	4.68833
51	51	4.79218	4.70061	4.70444	4.7127	4.71859	4.72788
53	53	4.82072	4.76634	4.76163	4.76843	4.77718	4.78868
55	55	4.81343	4.79793	4.79866	4.80011	4.82012	4.82667
49	51	4.71141	4.67069	4.66643	4.67215	4.68152	4.68809
47	53	4.52953	4.56119	4.57848	4.59613	4.61762	4.62667
45	55	4.2305	4.35852	4.39195	4.42362	4.45828	4.47226

Note that

- Average second prices are monotonically decreasing over sequences of auctions for all numbers of offers, but the decrease is not as fast when there are more offers. This is consistent with the observation in Figure 1 that more offers cause the second price to decrease more slowly.
- Similarly, average revenues are higher when there are more offers. however, average revenues are consistently below the ideal revenue of 5¢, by about 0.2 to 0.3¢.

4 Asymmetry in Correcting the Second Price

Suppose there are two offers, A and B. Suppose offer A has higher actual and estimated value than offer B. Consider two cases: the case where the value of B is overestimated, and the case where the value of B is underestimated. Fluctuations in the estimate of the value of A will either bring that estimate down to the overestimated value of B and not to the underestimated value of B, or they bring that estimate down to the overestimated value of B sooner than to the underestimated value of B. So in the case of an overestimate of the second price, the second price gets adjusted downward

Table 3: Average Second Value

	t					
Number of Bidders	**1000**	**5000**	**10000**	**20000**	**50000**	**100000**
2	4.8042	4.7331	4.7201	4.7136	4.7102	4.7094
3	4.8537	4.7813	4.7674	4.7596	4.755	4.7542
4	4.8864	4.8206	4.8061	4.7986	4.7935	4.7922
5	4.9052	4.8454	4.8256	4.8145	4.8078	4.8066
6	4.9214	4.8649	4.8479	4.8368	4.8304	4.8287
7	4.9293	4.8751	4.8561	4.8451	4.8382	4.836
8	4.9351	4.8867	4.8676	4.8552	4.847	4.8451
9	4.942	4.8953	4.8781	4.8653	4.8578	4.855
10	4.9478	4.9055	4.8886	4.8769	4.8674	4.865

either with more likelihood or more quickly than the second price gets adjusted if it is underestimated. So, in general, waiting for the first price to cross the second price before learning more about the second price causes underestimated second prices to be less likely to be corrected or to take longer to be corrected. As a result, publishers and the auctioneer lose revenue on average, even if the estimate of the second price is unbiased.

Figures 2(a) and 2(b) illustrate this dynamic. In both cases, $b_1 = b_2 = \$1$, $p_1^* = 0.05$, $p_2^* = 0.04$, and $p_1 = p_1*$ initially. The initial numbers of learning impressions are $n_1 = n_2 = 1000$. Figure 2(a) shows a sequence of auctions with $p_{20} = 0.045$, so the second price is initially overestimated. Figure 2(b) shows a sequence with $p_{20} = 0.0035$, so the second price is initially under-estimated. Both figures use the same random seed. Compare the two figures. In Figure 2(a), where the second price is initially overestimated, fluctuations in the estimate of the first price cause the first price to cross the second price. As a result, the second offer wins some impressions, adding information about p_2. On average, that information causes the estimate p_2 to approach p_2^*, reducing p_2. In Figure 2(b), where the second price is initially underestimated, the estimate of p_1 does not fluctuate enough to cause the estimated first price to cross the estimated second price. As a result, the underestimate of the second price persists.

Tables 5 and 6 show statistics over multiple trails similar to the single trials shown in Figures 2(a) and 2(b). For Tables 5 and 6, $b_1 = b_2 = \$1$, $p_1^* = 0.05$, $p_2^* = 0.04$, and initially $p_1 = p_1^*$. The initial number of learning impressions is $n_1 = n_2 = 1000$. For each initial value p_{20} in $\{0.025, 0.028, ..., 0.055\}$ and each number of auctions in

Figure 2: Asymmetric Correction for the Second Price

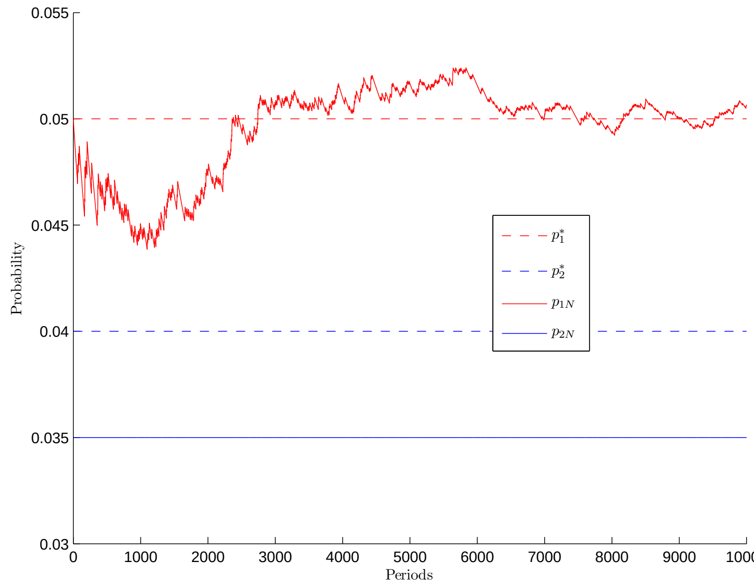

(a) Initially Overestimated Second Price

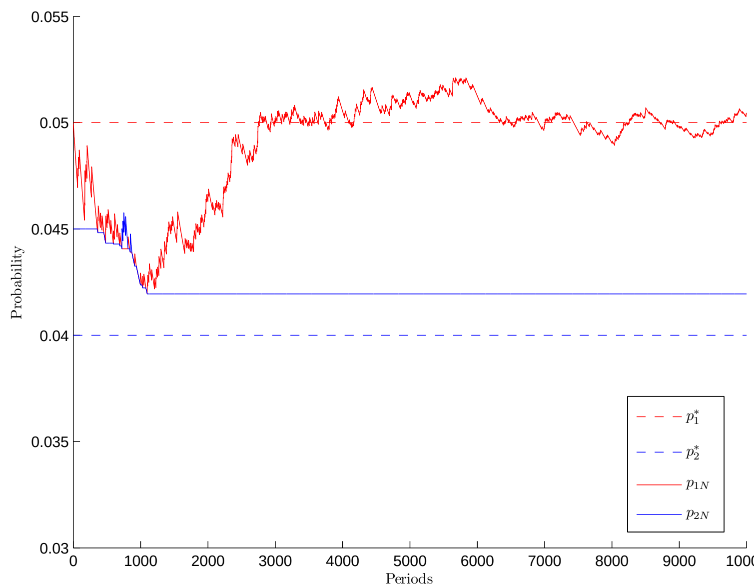

(b) Initially Underestimated Second Price

Table 4: Average Revenue

	t					
Number of Bidders	**1000**	**5000**	**10000**	**20000**	**50000**	**100000**
2	4.78163	4.68476	4.67291	4.67919	4.68987	4.69409
3	4.75881	4.68565	4.68694	4.69694	4.71571	4.73184
4	4.76421	4.70046	4.69823	4.71146	4.74276	4.76149
5	4.77055	4.69552	4.69752	4.71802	4.74836	4.77038
6	4.79688	4.69299	4.70335	4.72928	4.76123	4.78544
7	4.75987	4.68539	4.70073	4.72479	4.76492	4.78716
8	4.7644	4.69224	4.69738	4.72434	4.76579	4.79373
9	4.7417	4.6919	4.70212	4.72583	4.76783	4.79869
10	4.75285	4.69978	4.70741	4.72946	4.77278	4.8055

the sequence in $\{100, 500, 1000, 5000, 10{,}000\}$, there are results averaged over 10,000 sequences of simulated auctions. Table 5 shows the revenue loss

$$\frac{R_t^* - \hat{R}_t}{R_t^*},$$

which is the fraction of ideal revenue lost by using estimated values rather than actual values. Table 6 shows the fraction of impressions in the sequence that are awarded to offer 2, the offer with the lower eCPM.

In Tables 5 and 6, $p_{20} = p_2^*$ when $k_2 = 40$. For $k_2 > 40$, $p_{20} > p_2^*$, and the second price is initially overestimated. For $k_2 < 40$, $p_{20} < p_2^*$, and the second price is initially underestimated.

In Table 5, observe that for sequences of 100 auctions, the average revenue loss grows as p_{20} decreases from p_2^*, because as p_{20} shrinks, the second price, and hence charge to the winner, shrinks as well. As p_{20} increases from p_2^*, there is initially a growing revenue gain, because the growing second price increases the charge to the winner. Once p_{20} goes above p_{10}, the second offer becomes the auction winner. As p_{20} rises from there, the expected revenue decreases, because the charge to the second offer is inversely proportional to its estimated probability of action. Figure 3 shows the relationship between p_2 and expected revenue. In Table 5, note that as the number of auctions per sequence increases, the revenue gap becomes negative for all p_{20} values above p_2^*, because p_2 transitions through the higher-revenue zone of values in 0.05 to 0.04 as it descends from the initial overestimate.

Table 5: Revenue Loss (Percentage)

	t				
k_2	**100**	**500**	**1000**	**5000**	**10000**
25	37.6	37.4	37.3	37.2	37.3
28	29.5	29.6	29.7	29.8	29.8
31	22.2	22.5	22.3	22.2	22.3
34	15.1	14.6	14.6	14.7	14.8
37	7.2	7.3	7.2	7.2	7.3
40	−0.4	−0.2	−0.3	−0.3	−0.3
43	−7.6	−7.7	−7.7	−7.5	−7.4
46	−14.7	−13.9	−13.2	−12.7	−12.7
49	−17.1	−14	−13.5	−13.9	−14.2
52	0.6	−6.5	−8.8	−12.4	−13.5
55	7.5	1.2	−3.2	−10.6	−12.5

Table 6: Percentage of Impressions Awarded to Offer 2

	t				
$\mathbf{k_2}$	**100**	**500**	**1000**	**5000**	**10000**
25	0	0	0	0	0
28	0	0	0	0	0
31	0	0	0	0	0
34	0	0	0	0	0
37	0	0	$5e-05$	$1e-05$	$5e-06$
40	0	0.00404	0.0406	0.106	0.0987
43	0	0.291	0.758	0.862	0.636
46	0.0685	4.34	6.1	4.8	3.13
49	17.6	24.9	23.5	12.7	7.86
52	89.6	60.2	47.9	21.9	13.3
55	100	83.5	67.2	30.7	18.7

Figure 3: Relationship Between Second Price and Revenue.

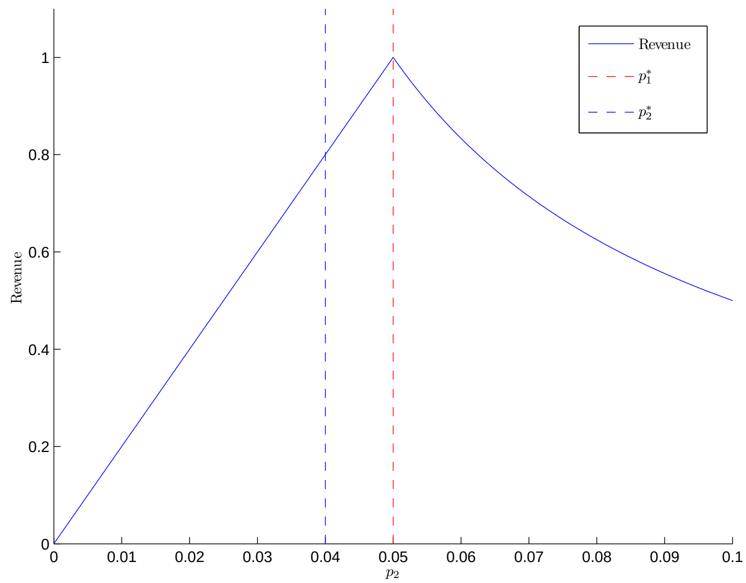

From Table 6, note that:

- For the higher values of p_{20}, the fraction of impressions awarded to offer 2 decreases with the number of auctions in the sequence, as the system learns over time that offer 2 has less value than offer 1.
- For the lower values of p_{20}, the fraction of impressions awarded to offer 2 is 0% or close to it, indicating that fluctuations in the estimate p_1 never or only rarely take p_1 across the lower values of p_{20}. As a result, the system does not correct the underestimates of p_2.

The impact on revenue from these asymmetric effects are shown in Table 5. For sequences of 10,000 auctions, compare the revenue gain from each initial $k_2 = 40 - \delta$ to each corresponding initial $k_2 = 40 + \delta$. The corresponding initial values are approximately equally likely under Bernoulli sampling. But the revenue loss from each initial $k_2 = 40 - \delta$ more than offsets the gain from the corresponding initial $k_2 = 40 + \delta$. The net result is a revenue loss to publishers and the auctioneer.

5 Learning the Second Price

Suppose that with probability $z \in [0, 1]$, rather than awarding the advertisement space to the bidder with the highest estimated value (offer w), the auction awards the impression to the bidder with the second highest estimated value (offer s). Also, suppose

Figure 4: Revenue Impact as a Function of Probability of Learning the Second Value

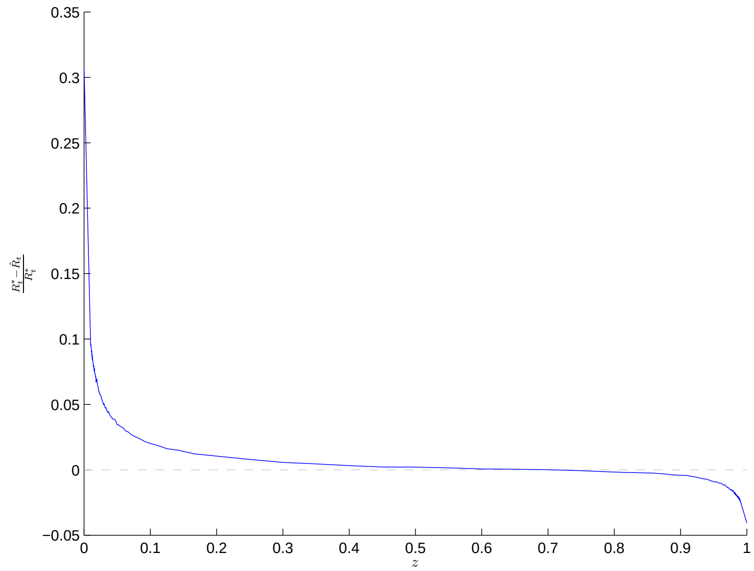

we charge the second price for these impressions, making the expected revenue:

$$\frac{b_s p_s}{p_s} p_s^* = b_s p_s^*$$

The goal of awarding some impressions to the runner-up is to learn the second price. However, if too many impressions are awarded to the runner-up, fairness decreases. As more impressions are awarded to the runner-up, the incentive to state full value in the bid in order to win the auction decreases. So the optimal rate of learning, z, mediates a tradeoff between short-term revenue and fairness/long-term revenue.

Figures 4 and 5 show the effects on revenue and fairness of awarding a fraction of impressions to the runner-up. Figure 4 shows the gap between ideal and actual revenue:

$$\frac{R^*_{100,000} - \hat{R}_{100,000}}{R^*_{100,000}}$$

averaged over 10,000 sequences of auctions for each value of the probability z of awarding the impression to the runner-up. Figure 5 shows the fairness $F_{100,000}$, averaged over the same 10,000 sequences of 100,000 auctions for each z value. Each sequence of simulated auctions has two offers, with $b_1 = b_2 = \$1$, $n_{10} = n_{20} = 100$ initial learning impressions per offer, $p_1^* = 0.05$, $p_2^* = 0.045$, each k_{10} generated i.i.d. by a binomial distribution over n_{10} samples with probability p_1^*, and each k_{20} is generated i.i.d. by a binomial distribution over n_{20} samples with probability p_2^*.

From Figure 4, note that:

Figure 5: Effect of Learning Second Value on Fairness in Auctions

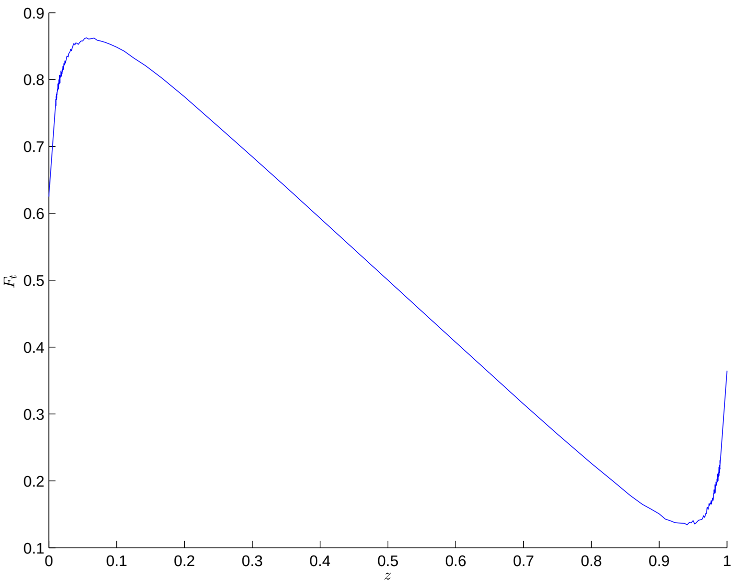

- When there is no learning of the second price ($z = 0$), there is about a 30% gap between ideal and actual revenue.
- The revenue gap declines steeply as z is increased to about 0.10. Devoting about 10% of impressions to learning the second price decreases the revenue gap from about 30% to less than 5%.
- Awarding a significant majority of impressions to the runner-up actually increases revenue above ideal revenue. To see why, consider setting $z = 1$. In this case, all impressions are awarded to the offer with the second-highest estimated value. When this offer has the second-highest actual value, the expected revenue is $b_s p_s^*$, which is equal to the ideal revenue r^*. However, when the runner-up based on estimates is really the ideal winner, the expected revenue is $b_w p_w^*$, which is greater than r^*. If p_{20} is sufficiently higher than p_{10} and p_1^*, then offer 1, as the runner-up, is awarded all impressions and p_1 converges to p_1^* without affecting p_2. As a result, the ideal winner (offer 1) is consistently awarded the impressions and charged the ideal first price, making revenue above r^*.

From Figure 5, note that:

- Increasing the fraction, z, of impressions awarded to the runner-up initially increases fairness, as learning the second price increases the probability of learning when the runner-up based on estimates is actually the ideal winner.
- Increasing the fraction of impressions awarded to the runner-up eventually decreases fairness, as the sheer volume of impressions awarded to the runner-up

ensures that neither offer gets a significant majority of the impressions. (With $z = 1/2$, fairness is about $1/2$.)

- Increasing z close to one increases fairness, because awarding so many impressions to the runner-up causes the ideal winner to gain more impressions in the case discussed in the notes for Figure 4. We certainly do not recommend setting z close to one to increase fairness - even at $z = 1$, the fairness is less than $1/2$.

Taken together, Figures 4 and 5 show that learning the second price can improve both revenue and fairness.

6 Discussion

This paper shows that failing to learn the second price harms revenue and fairness, presents a simple method to learn the second price, and shows that it can improve revenue and fairness. Directions for future work include:

- Apply the insights about the dynamics of the second price from this paper to more sophisticated methods to award impressions to runner-up offer than the simple random method used in this paper. More sophisticated methods are developed in Gonen & Pavlov (2007), Pandey & Olston (2007), and Wortman, Vorobeychik, Li & Langford (2007).
- Examine the simple method of giving the second price an initial boost in order to correct for the asymmetry in learning the second price.
- Explore optimally combining learning the second price with squashing (Varian 2007), which is designed to adjust the emphasis in offer value between the bid and the probability of action.
- Analyze the dynamics of awarding impressions to third place and subsequent offers as well as to the winner and runner-up. For example, learning the third price increases the second price, and hence revenue, when the third price crosses the second price. However, learning the third price cannot decrease the second price.
- Examine the effect of different auction mechanisms on bidder strategies. For example, if the second price is learned with probability $z > 0$, then the auction mechanism may no longer be incentive compatible, and the bidders may bid below their true values, causing a decrease in revenue.

References

Edelman, B., M. Ostrovsky & 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.

Feng, J., H. Bhargava & D.M. Pennock. 2004. "Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms.".

Gonen, R. & E. Pavlov. 2007. "An incentive-compatible multi-armed bandit mechanism." *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing* pp. 362–363.

Krishna, V. 2002. *Auction Theory.* Academic Press San Diego.

Lahaie, S. 2006. "An Analysis of Alternative Slot Auction Designs for Sponsored Search." *Proceedings of the 7th ACM conference on Electronic commerce* pp. 218–227.

Lahaie, S. & D.M. Pennock. 2007. "Revenue analysis of a family of ranking rules for keyword auctions." *Proceedings of the 8th ACM conference on Electronic commerce* pp. 50–56.

Milgrom, PR. 2004. *Putting Auction Theory to Work.* Cambridge University Press.

Pandey, S. & C. Olston. 2007. "Handling Advertisements of Unknown Quality in Search Advertising." *Advances in Neural Information Processing* 19:1065.

Regelson, M. & D. Fain. 2006. "Predicting click-through rate using keyword clusters." *Proceedings of the Second Workshop on Sponsored Search Auctions* .

Richardson, M., E. Dominowska & R. Ragno. 2007. "Predicting clicks: estimating the click-through rate for new ads." *Proceedings of the 16th international conference on World Wide Web* pp. 521–530.

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

Wortman, J., Y. Vorobeychik, L. Li & J. Langford. 2007. "Maintaining Equilibria During Exploration in Sponsored Search Auctions." *Lecture Notes in Computer Science* 4858:119.