CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 32: April 19

Lecturer: Éva Tardos *Scribe: Aravind Gollakota*

Continuing with our recent material on obtaining price of anarchy bounds from data, this lecture and the next is based on some work by Hoy et al. [HNS]. Our setting is generalized second-price auctions, and in particular the Bayesian case.

Last time, we plotted expected number of clicks $x(b)$[1] against the size of the bid b, and (on empirical basis) assumed the curve to be smooth and increasing. We had a similar plot for total price paid $P(b)$ vs. b. Then we looked at utility $u_i = v_i x(b) - P(b)$ and a Nash equilibrium assumption to obtain a rough range of possible values for v_i. But one problem was that this range was too fuzzy for us to plug the value back into our POA expressions and obtain useful bounds on the POA.

Today we look at a different approach where we try to bound POA directly without ever trying to estimate the value v_i. Let us also change our focus from total price paid $P(b)$, to price-per-click (ppc), $pc(b)$, because it is the latter that is much more clearly in the data. Indeed, assuming again that $\gamma_1 b_i \geq \ldots \gamma_n b_n$, recall that the ppc was just the price you needed to pay to keep your position in the sorted order, $pc(b_i) = \frac{b_{i+1}\gamma_{i+1}}{\gamma_i}$ — and each of the terms here is in the data. Meanwhile the total price paid is this times the expected number of clicks, $P(b) = pc(b)x(b)$, and $x(b)$ is not as easy to get from the data.

We propose that the following is a nice new perspective: plot a curve with $x_i(b)$ on the x-axis and $pc_i(b)$ on the y-axis. This is of course an implicit function of b. At low bids, we will see low (in fact zero) $x_i(b)$ and low $pc_i(b)$, and the curve then slopes up as at high bids both $x_i(b)$ and $pc_i(b)$ will be high. We refer to this function as τ_i, so that a point on the curve is $(x_i(b), pc_i(b)) = (\xi, \tau_i(\xi))$ (for some bid b), and formally we have $\tau_i(\xi) = pc_i(x_i^{-1}(\xi))$. We're going to assume that τ is continuous since it is the composition of functions we already assumed to be continuous. Intuitively, τ_i tells us how much we need to pay per click when we bid high enough to have our expected number of clicks be ξ; the bid itself is of secondary importance.

Figure 32.1: $pc_i(b)$ vs. $x_i(b)$ (subscript everything by i)

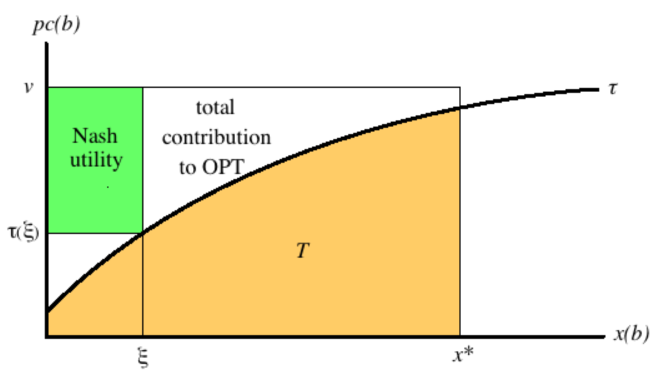

[1]In the past we have sometimes treated $x(b)$ as the probability of a click as well, which is of course proportional.

The idea now is to take advantage of this new picture using our geometric tools. Let $(\xi, \tau(\xi))$ be the point of Bayesian Nash equilibrium. Let v_i be our value per click, an upper bound on how much we're ever willing to pay per click. Our utility at ξ is $u_i = (v_i - \tau(\xi))\xi$, and this can be seen pictorially as a box on the plot (see figure). Let x_i^* be the point where OPT is achieved. (For convenience we plotted $x_i^* > \xi$ on the figure.) Then $v_i x_i^*$ is i's contribution to the OPT social welfare, in that it is i's utility, $(v_i - \tau(x_i^*))x_i^*$, plus contribution to revenue, $\tau(x_i^*))x_i^*$. This can also be represented by a (much bigger) box.

We'll see next class that there is good reason to think of the yellow part of the picture as revenue, and now the (white) parts of the big box that are contributing neither to the utility nor to the revenue are the ones that are to blame for high POA. We claim that the area under the curve, $T_i = \int_0^{x_i^*} \tau_i(\xi)\, d\xi$, is a sensible parameter to look at, and one that we can get cleanly from the data. We will revisit and justify this claim more properly next time (as x_i^* is actually not in the data). For now our hope is to show that if the total price paid is a substantial fraction of T_i, formally $P_i(b) \geq \frac{1}{\mu} T_i$, then the POA ($\frac{\text{Nash welfare}}{\text{OPT welfare}}$) is bounded above, specifically by $\frac{1-e^{\mu}}{\mu}$. We'd like μ to be as small as possible, and we hope to get it from the data.

A priori no such bound on T_i needs to hold for general τ_i and ξ, since we can make the ratio as terrible as we like. We need to use the Nash condition: that $u_i \geq (v_i - \tau_i(\xi))\xi$ for all ξ (or else the player would change the bid to get expected win up to ξ). This means $\tau_i(\xi) \geq v_i - u_i/\xi$ for all ξ, and in particular for $u_i/v_i \leq \xi \leq x^*$ (where below u_i/v_i, we use the nonnegative bound on τ_i instead). Integrating, we have

$$\begin{aligned}
\int_0^{x_i^*} \tau_i(\xi)\, d\xi &\geq \int_{u_i/v_i}^{x_i^*} (v_i - u_i/\xi)\, d\xi \\
&= v_i x_i^* - u_i - u_i \int_{u_i/v_i}^{x_i^*} \frac{1}{\xi}\, d\xi \\
&= v_i x_i^* - u_i - u_i(\ln x_i^* - \ln \frac{u_i}{v_i}) \\
&= v_i x_i^* - u_i + u_i \ln \frac{u_i}{v_i x_i^*}.
\end{aligned}$$

We claim that this is all we need to get our POA bound.

Theorem 32.1 *Let b be a Nash bid. If $P_i(b) \geq \frac{1}{\mu} T_i$, then the POA is at most $\frac{1-e^{\mu}}{\mu}$.*

Proof: The POA here is $\frac{\text{Nash welfare}}{\text{OPT welfare}} = \frac{u_i + P_i(b)}{v_i x_i^*}$. Using our assumption that $P_i(b) \geq \frac{1}{\mu} T_i$, what we want to show is that

$$\begin{aligned}
u_i + P_i(b) &\geq u_i + \frac{1}{\mu} T_i \\
&\geq u_i + \frac{1}{\mu}(v_i x_i^* - u_i + u_i \ln \frac{u_i}{v_i x_i^*}) \\
&\overset{?}{\geq} \frac{1-e^{\mu}}{\mu} v_i x_i^*.
\end{aligned}$$

For convenience let $y = u_i / v_i x_i^*$. Dividing our inequality throughout by $v_i x_i^*$, we see that what we need is that

$$y + \frac{1}{\mu}(1 - y + y \ln y) \geq \frac{1-e^{\mu}}{\mu}.$$

We prove this by showing that even the minimum of the LHS is at least the RHS. Solving for the point of minimum by differentiating the LHS wrt y, we get $1 + \frac{1}{\mu}(-1 + \ln y + 1) = 0$, or $y = e^{-\mu}$. Substituting $y_{\min} = e^{-\mu}$ into our inequality, we see quickly that it does indeed hold, and we are done. ■

At this point the idea is to look at empirical data to estimate μ, and ideally see that it is small. This will give us a good empirical POA bound.

Next time, we'll repurpose this analysis with $\mu = 1$ to obtain the optimal POA bound of $1 - 1/e$ in the first price auction setting. Another point to revisit is our heavy usage of our Nash inequality, $\tau_i(\xi) \geq v_i - u_i/\xi$. This may be far from tight at many ξ, leading to a weak POA bound. Inherently, this framework can never prove a better POA bound than $1 - 1/e$, whereas in reality the situation could be a bit better.

References

[HNS] Darrell Hoy, Denis Nekipelov and Vasilis Syrgkanis, "Robust Data-Driven Efficiency Guarantees in Auctions," arXiv:1505.00437.