Repeated Bidding with Dynamic Value

Benjamin Heymann, Alexandre Gilotte and Rémi Chan-Renous

Journée des rencontres ENSAE-ENSAI, 12-13 septembre 2023

Criteo AI Lab and Fairplay

Dynamic Value: motivation and challenge

Display Advertising

Figure 1: Display advertising allows the monetizing of publisher content on the internet.

Textbook solution to the bidding problem[1]

Bid = value

[1] In second price auctions, it is optimal to bid the value

Textbook solution to the bidding problem[2]

$$\text{Bid} = \overbrace{\underbrace{\alpha}_{\text{constant factor}} \times \underbrace{\Pr(\textit{Conversion}|\textit{Display})}_{\text{ctr}}}^{\text{value for the display opportunity}}$$

[2] In second price auctions, it is optimal to bid the valuation of the display

Dynamic value

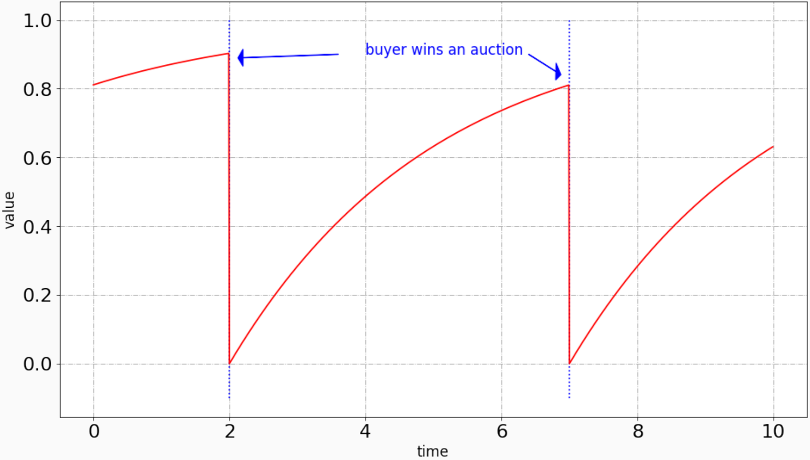

In this context, the formula is not true anymore

```
          ???
(*_*)'
|   |
```

Bid = ~~value for the display opportunity~~

A possible approach

The optimal bid satisfies

$$b^\star = \arg\max_{b} \mathbb{E}\left[D \cdot (V - Cost) | Bid = b\right]$$

with

$$\underbrace{V}_{\text{display valuation}} = \underbrace{\alpha}_{\text{constant factor}} \cdot \Delta S - \Delta FCost$$

$$D = 1 \quad \text{if we win the auction}$$

$$\Delta S = \Pr(conversion | D = 1) - \Pr(conversion | D = 0)$$

$$\Delta FCost = \text{"How much the display changes future cost"}$$

Martin Bompaire, Alexandre Gilotte, and BH. “Causal Models for Real Time Bidding with Repeated User Interactions"

The previous optimality condition is hard to solve in practice

Observe that $\Delta FCost$ and ΔS are

1. functions of the optimal bid
2. counterfactual quantities

$\rightarrow$ Solving an analytic example would be a step in the right direction

Contributions

- We analyze the case $\mathit{value} = k(\tau)$, and provide an algorithm to compute the optimal bidding strategy.
- We observe that empirically, there are constant shading factors that perform very well.

Repeated Bidding with Dynamic Value

Dynamic value

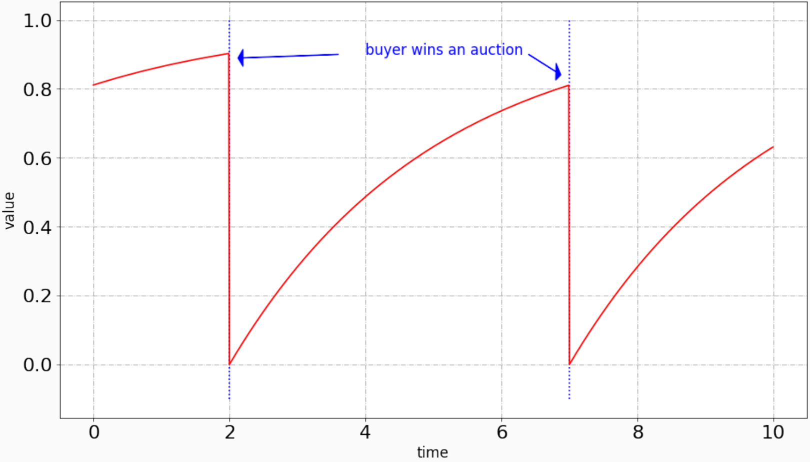

Model: (μ, k, q)

- τ: **age** of the last won auction
- μ: **intensity** of the auction arrivals
- $k(\tau)$: **value** of the item for the bidder (non-decreasing and bounded)
- $q(b)$: **win rate** probability that the buyer wins with a bid equal to b

---

- $p(b)$ average payment of the user when bidding b (second price auction)
- γ: discount rate

$\rightarrow$ Markov Chain with continuous time state and action.

Continuation function

For a bidding function $b : \tau \longrightarrow b(\tau) \in \mathbb{R}^+$, the expectation of the bidder's future payoff when the state is τ, is

$$V_b(\tau) \stackrel{def}{=} \mathbb{E} \sum_{i=1}^{\infty} e^{-\gamma T_i} \underbrace{(k(\tau(T_i)) - C_i) \mathbf{1}\{b(\tau(T_i)) > C_i\}}_{\text{auction } i \text{ payoff}}.$$

where $T_1, T_2 \dots T_n \dots$, are times of the next auctions, $C_1 ... C_n ...$ the competition at these times.

$$\underbrace{V^\star(\tau)}_{\text{Bellman value}} \stackrel{def}{=} \sup_{b \in \mathcal{B}} V_b(\tau).$$

Dynamic programming

Lemma

We have the relation

$$V_t^\star = \int_0^{+\infty} \mu e^{-(\mu+\gamma)t} \Big(\pi(k_t + V_0^\star - V_t^\star) + V_t^\star\Big) \mathrm{d}t,$$

where

$$U(v, b) = q(b) \cdot v - p(b) \quad (= \textit{static payoff})$$
$$\pi(v) = U(v, v) \quad (= \textit{static optimal payoff})$$

Moreover,

$$b^\star(\tau) = \max\left(0; \quad \underbrace{k(\tau) + V^\star(0) - V^\star(\tau)}_{\textit{incremental gain from winning the auction}}\right) .$$

ODE Reformulation

Lemma
Set $\Phi(t, v, \lambda) = \gamma v - \mu\pi(k_t + \lambda - v)$. *The value function* $V^\star$ *is the solution of the ordinary differential equation*

$$\begin{cases} \dot{Y}_t = \Phi(t, Y_t, y_0) \\ Y_0 = y_0 \end{cases} \qquad (\mathcal{F}_{y_0})$$

for some $y_0 \in \mathbb{R}_+$.

It should be noted that parameter y_0 is not given.

Sufficient condition for monotony of $b^\star$

Theorem

If k is concave, then $b^\star$ is increasing with τ, and strictly increasing on any interval where k strictly increases.

Counter-example

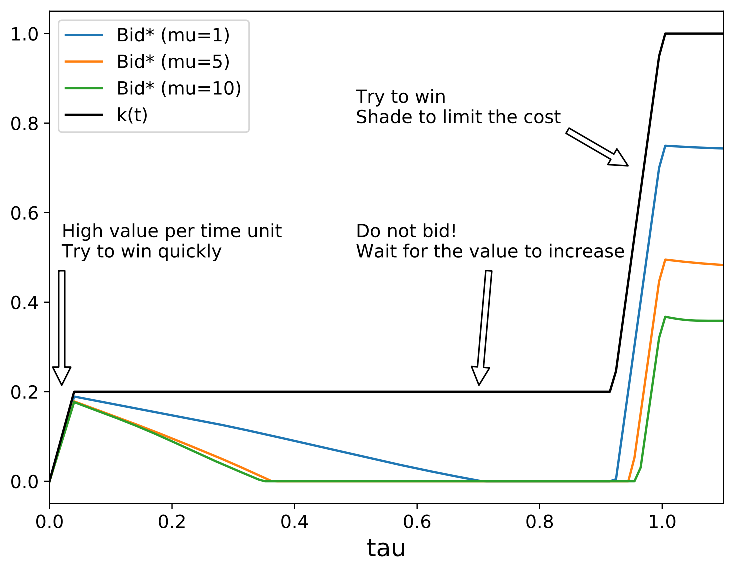

Main result

By the Cauchy-Lipschitz Theorem, the solution of the ordinary differential equation

$$\begin{cases} \dot{Y}_t = \Phi(t, Y_t, \lambda) \\ Y_0 = v_0 \end{cases} \qquad (\mathcal{F}_{y_0,\lambda})$$

admits a unique maximal solution $Z^{y_0,\lambda} : t \to Z^{y_0,\lambda}(t)$ for any $y_0 > 0$ and $\lambda > 0$. We set $Z^v(t) = Z^{v,v}(t)$.

Lemma

Suppose q continuous. The value $V_0^\star$ is the unique v for which $\lim_{t\to+\infty} Z^v(t)$ is finite.

We can solve numerically using a dichotomy

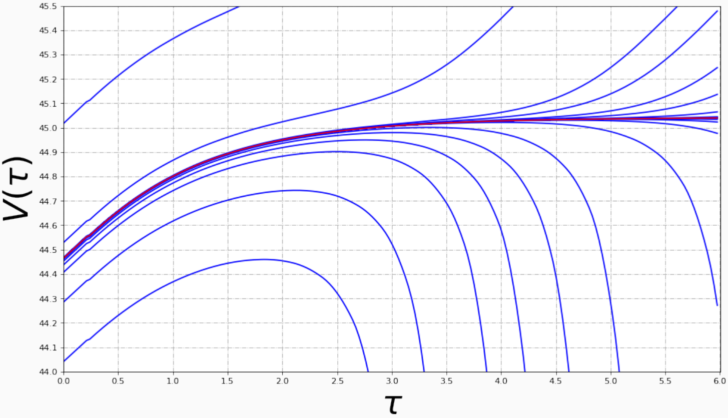

Figure 2: An example of Algorithm run. In red is the output of the algorithm, in blue the iterates.

What about shading policies? (1/2)

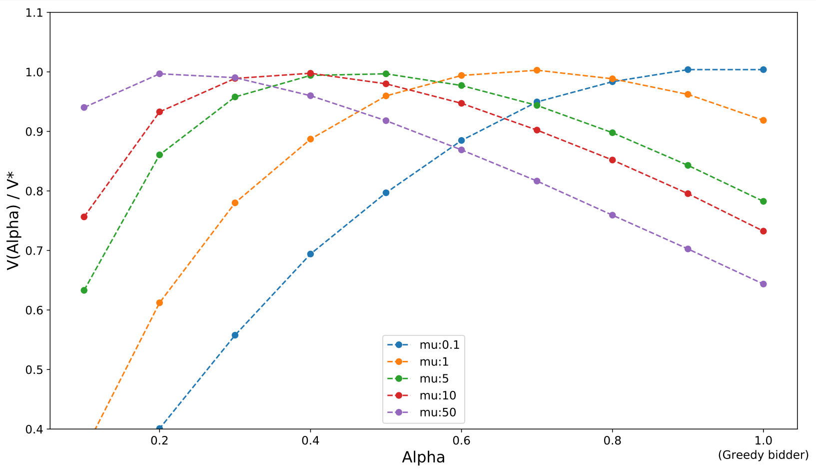

Figure 3: Ratio $V_\alpha / V^\star$ as a function of α with $k_\tau = 1 - e^{-t}$ and $\mu = 5$

What about shading policies? (2/2)

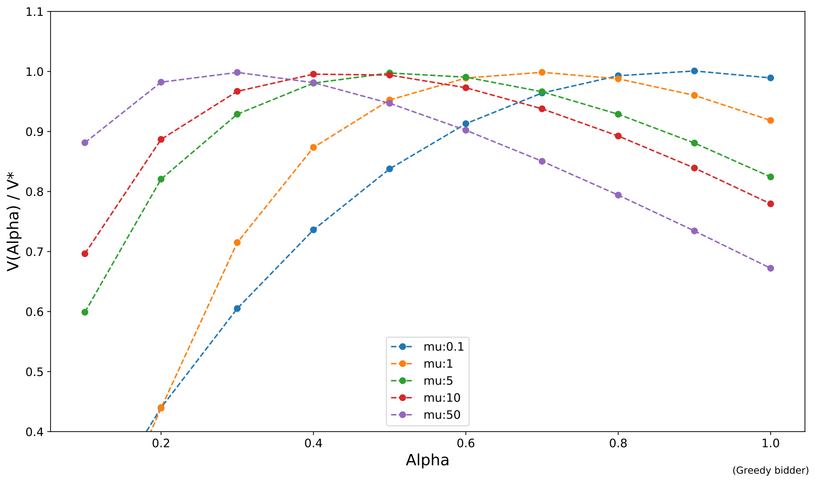

Figure 4: Ratio $V_\alpha / V^\star$ as a function of α with $\mu = 5$ and $k(t) = 1 - 1/1+t$

Conclusion

___ "In this end, it is not *that* bad, I just need to tweak α"
(*_*)
| |

$$\text{Bid} = \underbrace{\alpha^{\#}}_{\text{tweaked factor}} \times \Pr(Conversion|Display)$$

Next

1. Non-asymptotic guaranties for the shading policies
2. More general dynamics
3. Online learning of the parameters