

Cost Per Action Constrained Auctions

Benjamin Heymann

ABSTRACT

A standard result from auction theory is that bidding truthfully in a second price auction is a weakly dominant strategy. The result, however, does not apply in the presence of Cost Per Action (CPA) constraints. Such constraints exist, for instance, in digital advertising, as some buyer may try to maximize the total number of clicks while keeping the empirical Cost Per Click (CPC) below a threshold. More generally the CPA constraint implies that the buyer has a maximal average cost per unit of value in mind.

We discuss how such constraints change some traditional results from auction theory.

Following the usual textbook narrative on auction theory, we focus specifically on the symmetric setting. We formalize the notion of CPA constrained auctions and derive a Nash equilibrium for second price auctions. We then extend this result to combinations of first and second price auctions. Further, we expose a revenue equivalence property and show that the seller's revenue-maximizing reserve price is zero.

In practice, CPA-constrained buyers may target an empirical CPA on a given time horizon, as the auction is repeated many times. Thus, his bidding behavior depends on past realization. We show that the resulting buyer dynamic optimization problem can be formalized with stochastic control tools and solved numerically with available solvers.

CCS CONCEPTS

• **Theory of computation** → **Computational pricing and auctions**; *Stochastic control and optimization*; • **Applied computing** → Economics.

KEYWORDS

Ad Auctions, Bidding Strategy, Equilibrium, HJB.

ACM Reference Format:

Benjamin Heymann. 2019. Cost Per Action Constrained Auctions. In *The 14th Workshop on the Economics of Networks, Systems and Computation (NetEcon'19)*, June 28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3338506.3340269>

1 INTRODUCTION

What should you bid in a second price ad auction for a display with a known click-through rate (CTR), for a given cost-per-click (CPC)? The most probable (and possibly incorrect) answer is "CPC times

CTR". However, the right answer is "What do you mean by cost-per-click?". Indeed, if the CPC is the maximal amount the buyer is ready to pay for each click, the first answer is correct. But the story is different when CPC means the maximal average cost per click.

In this article, we challenge the current modeling approach for ad auctions. We argue that when some advertisers' business constraints apply, the expected outcome of the auction may depart from the traditional literature.

Advertising is a major source of revenue for Internet publishers, and as such, is financing a large part of the Internet. About 200 billion USD spent in 2017 on digital advertising [17]. Banners for display advertising are usually bought through a high frequency one unit auction mechanism called RTB (Real-Time Bidding) by or on behalf of advertisers.

When a user reaches a publisher page, it triggers a call to an RTB platform. The RTB platform then calls potential buyers, who have a few milliseconds to answer with a bid request. This results in an allocation of the display to a bidder, in exchange for a payment. The allocation and the payment are defined by the rules of the auction, which depends on the RTB platforms. The winner of the auction can show a banner to the internet user. Yet, for many end buyers, what really matters is not the display itself, but what will result from it: clicks, conversions, etc... This is why empirical CPA measures such as average cost per click (CPC) or average cost per order (CPO) are of paramount importance.

The literature on display advertising auctions has been growing over the last decade, pushed by the emergence of Internet giants whose business models rely on digital advertising. One track of research takes the seller's point of view and focuses on how to build "good" auctions, and relies on mechanism design theory [3; 13], while the dual track brings the buyer perspective under scrutiny, and focuses on the design of bidding strategies [1; 6; 10; 12; 24]. The reader can refer to [19] and [4] for an introduction to auction theory.

Usually, in performance markets, the advertiser has a target Cost Per Action (CPA) and/or a budget in mind. We will focus the analysis on the CPA constraint. We will assume that the budget constraint is absent or not binding. This topic of budget constraints has already been addressed in the literature [1; 3; 12]. The properties of the market equilibrium are discussed in [8; 9].

The CPA is measured in term of the average quantity of money spent per action, the action is a click, a visit, a conversion... For instance, if (a) we are facing a competition uniformly distributed between 0 and 1, (b) the CTR of every display is 0.5, (c) the auction is second price and (d) we are ready to pay 1 USD on average for a click, then if we bid 0.5, we will win half of the time. For every display won, we get on average half a click, and we pay $2 \int_0^{0.5} t dt = 1/4$, thus our expected CPC is $1/(4 * 0.5) = 0.5 < 1$. We should increase the bid to raise the empirical CPC.

One reason for using CPA constraints instead of budgets constraints in performance markets is that if the campaign performs well, there is no reason to shut it down at the middle of the month

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

NetEcon'19, June 28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6837-7/19/06...\$15.00

<https://doi.org/10.1145/3338506.3340269>

because of a strict budget limit. Those constraints, implemented by algorithms, can be modified by comparing the performances of different channels.

Auctions with ROI constraints are a hot topic [2; 14; 23]. In particular Golrezaei et. al [14] provide some empirical evidence that some buyers are ROI constrained and propose an optimal auction design and Wilkens et. al [22] explain the success of generalized second-price auctions (GSP) in the search market by the fact that buyers are value maximizers with ROI constraints.

Our work departs from them because we use a different notion of ROI: in the definition of CPA, we do not subtract the cost to the revenue in the numerator. The reason to do so is that end buyers reason in « cost per click » or « cost per order ». The numerator is not expressed in a unit of money and does not include the buyer payment to the seller. Also, our focus is more buyer than seller oriented.

In this work, we focus on the buyers perspective, but we also show on simple market instances that the CPA constraints may impact the seller design, or give rise to undesirable competitive behaviors.

This work brings several contributions to the table. First, we introduce the buyer's CPA constrained optimization problem and exhibit a solution, then we find the symmetric equilibrium in this setting and compare its properties with standard results from the literature. In particular, no reserve price should be used. Last, since in practice, the CPA is computed over a time window of repeated auctions, we move the discussion to the dynamic settings for which we explain why one can expect adaptive bid multipliers to provide solutions close to the optimal for the buyer.

2 MODELING ASSUMPTIONS AND NOTATIONS

We start by introducing the CPA-constrained bidding problem.

A seller is selling one item (a display opportunity in our context) through an auction. The item brings a value v_i to bidder $i \in 1..n$. These values are independently distributed among bidders. For example, in the context of the ad auction, v can be thought of as the expected number of clicks (or sales) the bidder would get after winning the auction. Thus, this value is not expressed in money but in unit of action. We will see later how these values, expressed in actions relates to bids, expressed in dollars. Let F_i be the cumulative distribution and f_i the density distribution of v_i . For simplicity, we assume the support of these density distributions to be compact intervals.

When we take the viewpoint of bidder i , we denote by b^{-i} the greatest bid of the other bidders (the price to beat). We denote by g^i and G^i its density and cumulative distributions. Unless otherwise stated, we also assume the auction to follow a second price rule, with $n > 1$ bidders competing for the item.

Bidders with CPA constraints compare:

- the expected value they get in the auction
- the expected payment they incur.

We will refer to the CPA with the letter T because it is our "targeted cost per something". We will denote by b^i the bid of a given bidder i of interest.

The bidder wants to maximize his expected value, subject to an *ex-ante* constraint in expectancy representing the targeted CPA. The constraint is *ex-ante* because, in practice, the same bidding strategy is going to be applied repeatedly (or even simultaneously) on several similar auctions. If for a random variable X , we denote by $[X] \in \{0, 1\}$ the binary random variable that takes the value 1 when $X \geq 0$, and 0 otherwise, then the expected value earned by the buyer bidding b^i is $E_{v, b^{-i}}[b^{-i} < b^i(v)]v$. Here $E_{v, b^{-i}}$ is there expectation operator on the distribution of v and b^{-i} . Similarly the bidder expected cost is $E_{v, b^{-i}}[b^{-i} < b^i(v)]b^{-i}$.

We can now express the bidder optimization problem. The bidder is looking for a bid function $b^i : [0, 1] \rightarrow \mathbb{R}^+$ solution of

$$\max_{b^i} E_{v, b^{-i}}[b^{-i} < b^i(v)]v, \quad (1)$$

subject to $E_{v, b^{-i}}[b^{-i} < b^i(v)]b^{-i} \leq T^i E_{v, b^{-i}}[b^{-i} < b^i(v)]v$ (CPA). Observe that if we remove the constraint, the buyer would buy all the opportunities no matter the cost. We pinpoint that problem (1) is not equivalent to a budget constraint problem: we would have had $E_{v, b^{-i}}[b^{-i} < b^i(v)]b^{-i} \leq \text{budget}$ instead of the (CPA) constraint. Nor is it equivalent to the maximization of the yield $E_{v, b^{-i}}(T^i v - b^{-i})[b^{-i} < b^i(v)]$, as illustrated in the introduction.

We model the interactions among buyers with a game. Observe that this game has constraints on the strategies profile. For a given set of competing bidding strategies, a best reply of the bidder is a solution of (1). A (constrained) Nash equilibrium is a strategies profile with components that are best replies against the others .

In the following, the superscript i is often omitted to help readability.

3 BIDDING BEHAVIOR AND SYMMETRIC EQUILIBRIUM

The main result of this section is the derivation of a symmetric Nash equilibrium for a CPA constrained second price auction. We then generalize this result to linear combinations of first price and second price auctions.

3.1 Second Price

Going back to the motivational example of the introduction, if a buyer bids Tv , then he would pay *at most* T per unit of action. Since he can increase his bid until he pays T per unit of action *on average*, it is intuitive that he should bid higher than Tv to maximize the criterion. The next theorem formalizes this idea.

THEOREM 3.1 (BEST REPLY). *For any bounded distribution of the price to beat b^{-i} , there exists $\lambda^* \in \mathbb{R}^+$ such that a solution of (1) writes:*

$$b^i(v, \lambda^*) = (1/\lambda^* + T^i)v. \quad (2)$$

LEMMA 3.2 (CPA MONOTONY). *If the bidder bids proportionally to the value i.e. if there exists α such that $b(v) = \alpha v$, then the CPA is non-decreasing in α .*

PROOF. Fix $u \in [0, 1]$, $\alpha > 0$, then the expected cost knowing that the value is v and the auction is won is $\int_0^{\alpha v} tg(t)dt$, where g is the probability density of the price to beat. This quantity is increasing in α . Meanwhile the value is v . We get the result by integration on v . \square

If v_1, v_2, \dots, v_n are n independent draws from F , we denote by $v_i^{(n)}$ the average of the i^{th} greater draw (the i th order statistic).

THEOREM 3.3 (EQUILIBRIUM BID). *The unique symmetric constrained Nash equilibrium strategy is*

$$b^{i*}(v) = \frac{T}{\gamma} v, \quad (3)$$

where

$$\gamma(F, n) = n \left(\frac{v_1^{(n-1)}}{v_1^{(n)}} - \frac{n-1}{n} \right). \quad (4)$$

Discussion. Theorem 3.1 is the answer to the introductory question. Despite its simplicity, it shows something that is probably overlooked: it may happen in a second price auction that the bidder's optimal bid depends on the competition. Another useful insight is that the bid is linear w.r.t. the value, which implies that simple bid multipliers could be used optimally. Observe that the result holds also for non-symmetric settings.

The proof relies on the strong incentive compatibility of the second price auction in standard setting reinterpreted on the Lagrangian of the optimization problem. The parameter λ^* should be interpreted as the Lagrangian multiplier associated with the CPA constraint.

Informally speaking, the harder the constraint, the bigger λ , the smaller the bid. To get an intuitive interpretation, let us look at the extreme values. If λ is set to $+\infty$, one bids only for displays that do not consume the constraint, while when λ goes to zero, the bid diverges (It is an intuition, since in practice for a sufficiently high bid, the bidder win all the auctions), which can be interpreted as maximizing the buyer's total value without any spend related constraints: the solution in such scenario would be to buy all available displays.

Lemma 3.2 is useful to understand what is happening. Observe that since the bid is linear w.r.t. the value and the objective increasing in the bid multiplier, the optimal bid multiplier is the maximal admissible one.

Theorem 3.3 is one of our main results. Observe that the **competition factor** γ is only a function of the number of opponents and F . Since it is smaller than one (see Appendix) the bid is proportional to Tv with a factor greater than one. Compared with the standard setting, the seller's expected revenue is multiplied by T/γ . The asymptotic value of γ , as well as other natural questions concerning the equilibrium will be studied in Section 4. Note that Theorem 3.3 is a necessary and sufficient condition for a symmetric Nash equilibrium, therefore it is unique, but we cannot claim that we have identified the only Nash Equilibrium (see Section 4).

3.2 Generalization

We now generalize the argument to auctions which are convex combinations of first price and second price. We motivate this extension by the existence in the industry of auctions which mix together first and second pricing rules, such as soft-floor.

If we denote by $S(b, b^-)$ the payment rule of the auction when the buyer bids b and the best competing bid is b^- , then we can make the following observations: (1) for any $s \geq 0$, $S(sb, sb^-) = sS(b, b^-)$, (2) $S(b, b) = b$, (3) $S(b, b^-)$ is the sum of a linear function of b and a linear function of b^- .

Consider a standard auction with payment rule S and symmetric buyers with i.i.d. values distributed according to F (no CPA constraints). Let \hat{b} be a symmetric equilibrium strategy in such a setting.

THEOREM 3.4. *The bid $b^{i*}(v) = \frac{T}{\gamma} \hat{b}(v)$ is a symmetric equilibrium strategy, of the CPA constrained auction with payment rule S .*

Example. : Consider a first price auction with 2 buyer having uniformly distributed values in $[0, 1]$. It is known [4] that an equilibrium bid for a standard auction is $v \rightarrow 0.5v$. By (16), $\gamma = 0.5$. Thus an equilibrium bid for the CPA constrained auction is $v \rightarrow Tv$. Indeed: in first price, we ensure the saturation of the CPA constraint by bidding Tv .

4 PROPERTIES OF THE EQUILIBRIUM

Observe that the expected payment of a buyer is equal to his expected value times T . Since the total welfare is fixed, the expected payment of every buyer is the same no matter the auction. We thus recover a standard result [21] from auction theory, adapted to CPA constrained auctions.

THEOREM 4.1 (REVENUE EQUIVALENCE). *All the auctions described in Section 3.2 bring the same expected revenue to the seller.*

PROOF. Using the proof of Theorem 3.3, we see that constraint in (1) is binding at the symmetric equilibrium. Therefore the payment of a buyer to the seller is equal to $T\mathbb{E}v[v > v^-] = T\mathbb{E} \max_{i \in \{1..n\}}(v_i)/n$. \square

On the other hand, the following observation is quite unusual.

THEOREM 4.2 (OPTIMAL RESERVE PRICE). *With CPA constraints, the optimal reserve price in a symmetric setting is zero.*

Explanation: a reserve price would decrease the social welfare, and thus decrease the expected payment because of the CPA constraint.

PROOF. Same idea as before: In the presence of a reserve price, the constraint in (1) tells us that the buyer expected payment is smaller than $T\mathbb{E}v[v > v^-i][v > r]$, which is smaller than $T\mathbb{E}v[v > v^-i]$ for $r > 0$. Using the proof of Theorem 3.3 we see that the constraint should be binding at the equilibrium, and thus this last quantity corresponds to the expected payment when $r = 0$. Thus the reserve price should be set to zero. \square

We display in Figure 1 the results of a simulation that illustrates Theorem 4.2. Observe that the proofs of the two last results rely on the symmetry of the setting (same CPA target T , same value distribution F). We argue that this should not be seen as a limitation of our results, because (a) Revenue Equivalence in the classical setting also requires symmetry, (b) Theorem 4.2 provides a striking counterexample to the common belief that the reserve price increases the revenue of the seller. Furthermore we believe (because the proofs mostly rely on the binding of the CPA constraints) that Theorem 4.2 to be extendable to non-symmetric setting with non i.i.d. valuations and different CPA constraints, as long as one can guarantee that the CPA constraint is binding.

This completely departs from the idea that reserve prices should be used to increase the seller's revenue. Observe that in practice

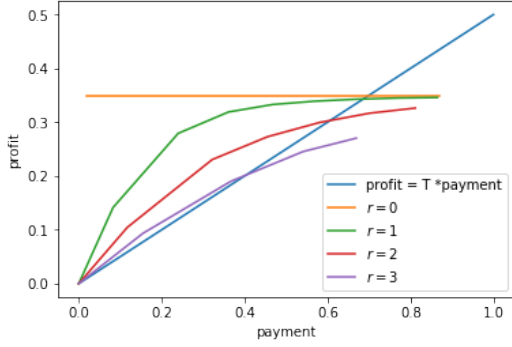


Figure 1: Potential outcomes (buyers' profits and payments) for different values of the bid multiplier in a symmetric setting. Each curve corresponds to a different value of the reserve price. The equilibrium points are the intersections of those curves with the line $\text{profit} = T \cdot \text{payment}$. We took $T = 0.4$, uniform distribution of values, and two bidders for the numerical simulation.

in the case of display advertising, the buyer may take some time to react to a change of empirical CPA. Consequently, measuring the seller's long term expected revenue uplift is a technical and business challenge. Moreover, the reserve price, by decreasing the welfare, may on the long run trigger an increase of the target CPA, since buyers have to do an arbitrage between volume and CPA.

We pinpoint that in our model, there is only one seller. If there were several sellers, the conclusion on the reserve price would be different. To share the intuition, we can envision a situation with one big and one very small sellers. If the small seller set a reserve price, the short term effect will be an increase of its revenue. But it will not be completely compensated by the long term effect of a bid level decrease, since the overall CPA of the buyer will not be that much affected. More generally, with several sellers, we would get a game theoretical situation where the reserve price of one seller affects the other sellers' revenues.

Last but not least, we do not claim that reserve prices should be set to zero for CPA constrained bidders in all situations. A trivial counterexample to such claim is embodied by a setting with only one buyer and no competition.

THEOREM 4.3 (CONVERGENCE). *For any number of bidders n , $\gamma(F, n) \leq 1$. Moreover if the value v is bounded, then*

$$\lim_{n \rightarrow \infty} \gamma(F, n) = 1 \quad (5)$$

PROOF. Let us denote by $Y^{(k)}$ the maximum of k random variables drawn with the distribution F . One just need to observe that

$$\gamma(F, n) = 1 - \frac{\mathbb{E}(v - Y^{n-1} | v > Y^{n-1})}{v_1^{(n)}}. \quad (6)$$

□

This can be interpreted in the following way: when there is a great number of bidders, the competition is such that the payment tends to become first price.

REMARK 1. *There may be a nonsymmetric equilibrium.*

The basic idea is that one bidder can bid higher than necessary to force another bidder to leave the auction because his (CPA) constraint cannot be satisfied.

PROOF. We concentrate on exhibiting a counter-example. Take $n = 2$, $T = 1$, F is the uniform distribution over $[2, 3]$, and denote by $(\alpha_i)_{i=1,2}$ the bid multipliers of the two buyers. Set $\alpha_2 = 6$. Then if $\alpha_1 \leq 4$, buyer 1 does not win any auction, but the (CPA) constraint is satisfied.

On the other hand, if $\alpha_1 > 4$, then the cost vs value ratio is bigger than 4, and the (CPA) constraint is violated. We can check that if $\alpha_1 = 0$ the (CPA) constraint is satisfied for 2. Therefore $(\alpha_1, \alpha_2) = (0, 6)$ is an asymmetric Nash equilibrium. □

In addition to this, observe that the buyers may be tempted to bid other strategies than the linear best reply. For example consider this example: Take $n = 2$, an exponential distribution with parameter 1, $T = 1$. One buyer may increase its profit by bidding with an affine function. Compare $(b_1(v_1), b_2(v_2)) = (2v_1 + 1, 2v_2)$ with $(b_1(v_1), b_2(v_2)) = (3v_1, 3v_2)$.

On a simulation with 10^8 auctions, we get in the first case an empirical CPA of 0.92, and a revenue of 0.84 for bidder 1 (resp. 1.0 and 0.61 for bidder 2), while in the second case, we get an empirical CPA of 1, and a revenue of 0.75 for bidder 1 (resp. 1.0 and 0.75 for bidder 2). Those simple, informal examples indicate that the bidders may be tempted to bid aggressively to weaken the other bidders CPA.

5 DYNAMIC BIDDER PROBLEM

In practice, the buyer behaviors may differ from the static case: (1) the buyer can adapt his bidding strategy to the past events, (2) the linear constraint (CPA) does not reflect the buyer risk aversion, (3) the benefit the buyer gets from a won auction is stochastic (for example, if he is only interested in clicks or conversions), (4) from a business perspective, there is a trade-off between CPA and volumes in the buyer mind. This trade-off can be expressed in different ways.

We propose in this section a continuous time optimal control framework to express and study the buyer's problem in a dynamic fashion. The approach combines the advantages of a powerful and mathematically clean expressiveness with theoretical insights and numerical tools.

We refer the reader to [18] for a reference on stochastic control, and to [11] for a presentation of the numerical methods involved. The User Guide [5] provides a hands-on presentation of the topic.

We only model one individual buyer. Yet note that the optimal control formulation is the first building block for the study of the full market dynamics (with mean field games [7; 15] or differential games [20]). Our main discovery in the stochastic case is that (a) the optimal bid is not linear in the value anymore, but (b) we can still reasonably propose a linear bid (at the cost of an approximation to be discussed thereafter).

5.1 Deterministic Dynamic Formulation

We use the deterministic case to introduce some notations, and recover with optimal control tools some of the results we already

derived with the static model. Indeed, when we neglect the stochastic aspect of the reward, the problem is very similar to the static one.

We consider a continuous time model. The buyer receives a continuous flow of requests. The motivation to propose a continuous time model is that it makes the theoretical analysis easier [16; 18].

We denote by \mathcal{R} the instantaneous revenue and by C the instantaneous cost:

$$\mathcal{R}(b) = \mathbb{E}v[b^- < b(v)], \quad \text{and} \quad (7)$$

$$C(b) = \mathbb{E}b^-[b^- < b(v)], \quad (8)$$

If we denote by τ the time horizon, the bidder is now maximizing with respect to b

$$\int_0^\tau \mathcal{R}(b_t) dt, \quad (9)$$

subject to $\dot{X}_t = T\mathcal{R}(b_t) - C(b_t)$, $X_\tau \geq 0$ and $X_0 = x_0$.

The state X_t represents the (CPA) constraint is should be positive at $t = \tau$ for the (CPA) constraint to be satisfied. We use the Pontryagin Maximum Principle to argue that a necessary condition for an optimal control is that there exists a multiplier $p(t)$ so that $b(t)$ maximizes $H = p(T\mathcal{R} - C) + \mathcal{R}$. Moreover $\dot{p} = -\frac{\partial H}{\partial x} = 0$. We thus recover the main result of the static case: the bid is linear in v . Moreover in the absence of stochastic components, the multiplier is constant. Observe that in practice, one could use the solution of problem (9) to implement a Model Predictive Control (MPC), using x_0 as the current state and thus get an online bidding engine.

5.2 Stochastic Formulation

The number of displays is much higher than the number of clicks/sales, therefore we neglect the randomness of the competition/price over the randomness that comes from the *actions*. Because of the number of displays involved, we argue that by the law of large number, the uncertainty on the action outcome can be apprehended a white noise. We thus add an additional, stochastic term in the revenue: $T\sigma_i dW_t$, where W_t is a Brownian motion. We get that an optimal bid b should maximize (see [18] for a reference)

$$\mathbb{E}_{b^-, v}[b > b^-]((Tv - b^-)p + v + T^2\sigma^2 M), \quad (10)$$

for some p and M .

Since the realization of the action follows a binomial law, $\sigma \propto v(1 - v)$. Assuming $v \ll 1$, we can approximate it as v . Therefore every infinitesimal term of the Hamiltonian becomes $((Tv - b^-)p + v + vMT^2)[b > b^-]$ reproducing the previous reasoning we get

$$b^* = v(T + \frac{MT^2 + 1}{p}). \quad (11)$$

Conclusion: Once again, the bid factor approach is justified! Observe that this conclusion comes from an approximation, (therefore the bid is not strictly speaking optimal for the initial problem), but by doing so, we have reduced the control set from $\mathbb{R}^+ \rightarrow \mathbb{R}^+ \rightarrow \mathbb{R}^+$. This reduction gives access to a well-understood approach to solve this kind of problem: the continuous time equivalent of Dynamic Programming.

Observe that our problem definition is incomplete: we need to take into account the constraint on the final state. We restate

this constraint with a penalty $K(X)$. Using the dynamic programming principle, the buyer problem is equivalent to the following Hamilton-Jacobi-Bellman equation:

$$V_t + \sup_\alpha \left((T\mathcal{R}(\alpha) - C(\alpha))V_x + \frac{T^2}{2}\mathcal{R}(\alpha)V_{xx} + \mathcal{R}(\alpha) \right) = 0 \quad (12)$$

$$\text{and } V_\tau = K, \quad (13)$$

with

$$\mathcal{R}(\alpha) = \mathcal{R}(v \rightarrow \alpha v), \quad \text{and} \quad (14)$$

$$C(\alpha) = C(v \rightarrow \alpha v) \quad (15)$$

5.3 Numerical Resolution

Our aim here is to illustrate that HJB approaches provide powerful numerical and theoretical tools to model the buyer's preferences. A quantitative comparison of the performance of the stochastic method is out of the scope of this article.

We solve an instance of the problem with the numerical optimal stochastic control solver BocopHJB [5]. BocopHJB computes the Bellman function by solving the non-linear partial differential equation (HJB). Observe that the state is one dimensional, which makes this kind of solver very efficient.

We take $G(b^-) = (b^-)^\alpha$ on $[0, 1]$ and v uniformly distributed on $[0, 1]$. We get that $\mathcal{R}(\alpha) = \frac{\alpha^\alpha}{\alpha+2}$ if $\alpha < 1$ and $\mathcal{R}(\alpha) = \frac{\alpha^\alpha}{(a+2)(\alpha^{a+2})} + \frac{1}{2}(1 - \frac{1}{\alpha^2})$ otherwise, similarly, $C(\alpha) = \frac{a\alpha^{a+1}}{(a+1)(a+2)}$ for $\alpha \leq 1$ and $C(\alpha) = \frac{a}{(a+1)(a+2)\alpha}((a+2)\alpha - a - 1)$ for $\alpha \geq 1$ (see Figure 2). We take $T = 0.8$ and a linear penalty for negative X_τ . The result of a simulated trajectory is displayed in Figure 3(a).

We see that on this specific instance of the stochastic path, the control is decreased at the end of the time horizon to adapt to the sudden spike of CPA. This is the kind of behavior we can expect from a stochastic controller. By contrast, a constant multiplier may result in the kind of trajectories displayed on Figure 3 (b).


6 CONCLUSION

We have formalized the existence of CPA constraints in the buyers bidding problem. These constraints are of paramount importance in performance markets, yet they tend to be neglected in the current literature. We have seen how standard results translate into this context, sometimes with surprising insights. We also envisioned the dynamic, stochastic setting by combining a first-order approximation with the Hamilton-Jacobi-Bellman approach. This allowed us in particular to derive a numerical method.


This work raises questions that deserve a more extended study. In particular, can we generalize the first three theorems to other constraints, auction rules and settings? Can the existence of the aggressive behaviors pointed out at the end of Section 5 pose a problem to the market? More generally, the questions that were raised for budget constrained auctions and pacing equilibrium could be posed for CPA constrained auctions.

ACKNOWLEDGMENTS

I would like to thank Jérémie Mary, Nicolas Chrysanthos and the anonymous referees for their valuable feedback



(a) $a=1$

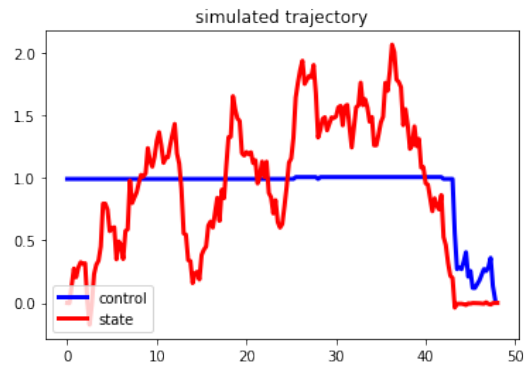


(b) $a=5$

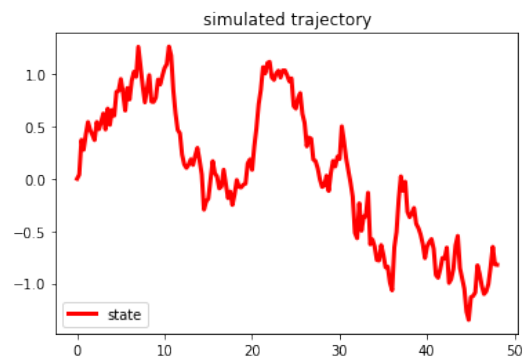
Figure 2: \mathcal{R} and C and empirical T for two values of a .

REFERENCES

- [1] Agarwal, D., Ghosh, S., Wei, K., You, S.: Budget pacing for targeted online advertisements at linkedin. In: Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 1613–1619. ACM (2014)
- [2] Auerbach, J., Galenson, J., Sundararajan, M.: An empirical analysis of return on investment maximization in sponsored search auctions. In: Proceedings of the 2nd International Workshop on Data Mining and Audience Intelligence for Advertising - ADKDD '08. pp. 1–9. ACM Press, New York, New York, USA (2009).
- [3] Balseiro, S.R., Besbes, O., Weintraub, G.Y.: Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* **61**(4), 864–884 (2015)
- [4] Baranwal, G., Kumar, D., Raza, Z., Vidyarthi, D.P.: Auction theory. Academic press (2018).
- [5] Bonnans, F., Giorgi, D., Heymann, B., Martinon, P., Tissot, O.: BocopHJB 1.0.1 – User Guide (2015).
- [6] Cai, H., Ren, K., Zhang, W., Malialis, K., Wang, J., Yu, Y., Guo, D.: Real-Time Bidding by Reinforcement Learning in Display Advertising. In: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining. pp. 661–670. ACM (2017).
- [7] Carmona, R., Delarue, F.: Probabilistic Analysis of Mean-Field Games. *SIAM Journal on Control and Optimization* **51**(4), 2705–2734 (1 2012).
- [8] Conitzer, V., Kroer, C., Panigrahi, D., ..., O.S.a.p.a., 2018, u.: Pacing Equilibrium in First-Price Auction Markets. arxiv.org
- [9] Conitzer, V., Kroer, C., Sodomka, E., Stier-Moses, N.E.: Multiplicative Pacing Equilibria in Auction Markets. arxiv.org (2017),
- [10] Diemert, E., Meynet, J., Galland, P., Lefortier, D.: Attribution Modeling Increases Efficiency of Bidding in Display Advertising. In: Proceedings of the ADKDD'17. p. 2. ACM (7 2017),
- [11] Falcone, M., Ferretti, R.: Semi-Lagrangian Approximation Schemes for Linear and Hamilton–Jacobi Equations (2013).



(a) stochastic controller



(b) deterministic controller

Figure 3: Simulation result.

- [12] Fernandez-Tapia, J., Guéant, O., Lasry, J.M.: Optimal real-time bidding strategies. *Applied Mathematics Research eXpress* **2017**(1), 142–183 (2016)
- [13] Golrezaei, N., Lin, M., Mirrokni, V., Nazerzadeh, H.: Boosted Second-price Auctions for Heterogeneous Bidders. *SSRN Electronic Journal* (8 2017). <https://doi.org/10.2139/ssrn.3016465>
- [14] Golrezaei, N., Lobel, I., Paes Leme, R.: Auction Design for ROI-Constrained Buyers. *Ssrn* pp. 1–35 (2018)
- [15] Guéant, O., Lasry, J.M., Lions, P.L.: Mean field games and applications. In: *Lecture Notes in Mathematics*, vol. 2003, pp. 205–266 (2011).
- [16] Heymann, B., Bonnans, J.F., Martinon, P., Silva, F.J., Lanas, F., Jiménez-Estévez, G.: Continuous optimal control approaches to microgrid energy management. *Energy Systems* **9**(1), 59–77 (2 2018).
- [17] Kafka, P., Molla, R.: 2017 was the year digital ad spending finally beat TV - Recode (2017),
- [18] Pham, H.: *Continuous-time Stochastic Control and Optimization with Financial Applications*, vol. 61 (2009).
- [19] Roughgarden, T.: *Twenty lectures on algorithmic game theory*. Cambridge University Press (2016)
- [20] Vajda, S., Isaacs, R.: *Differential Games. A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*, vol. 51. Dover Publications (2007).
- [21] Vickrey, W.: Auctions and bidding games. *Recent advances in game theory* **29**, 15–27 (1962)
- [22] Wilkens, C.A., Cavallo, R., Niazadeh, R.: GSP - The Cinderella of Mechanism Design pp. 1–14 (2017)
- [23] Wilkens, C.A., Cavallo, R., Niazadeh, R., Taggart, S.: Mechanism Design for Value Maximizers (7 2016)
- [24] Zhang, Y., Wei, Y., Ren, J.: Multi-touch attribution in online advertising with survival theory. In: 2014 IEEE International Conference on Data Mining (ICDM). pp. 687–696. IEEE (2014)

A EXAMPLES

A.1 Power Law Case: derivation of the competition factor

We take $F(v) = v^a$ over $[0, 1]$ with $a > 0$. We get that $v^{(n)_1} = \frac{an}{an+1}$.

Therefore we can write $\frac{v_1^{(n-1)}}{v_1^{(n)}} = \frac{n-1}{n} \frac{an+1}{a(n-1)+1}$, which leads to

$$\gamma(a, n) = (n-1) \left(\frac{an+1}{a(n-1)+1} - 1 \right). \quad (16)$$

We observe that γ is increasing in n and a . It converges to 1 as n (or a) goes to infinity. A plot of γ is displayed in Figure 4(a).

Remark: If we consider the 2 bidders case, and denote by $\alpha > 0$ and $\beta > 0$ their respective bid multipliers, then the payment and expected value of the first bidder are respectively equal to

$$\pi = \beta \left(\frac{\alpha}{\beta} \right)^{a+2} \frac{1}{(a+2)(2a+3)}, \text{ and } V = \frac{\left(\frac{\alpha}{\beta} \right)^{a+1}}{(a+1)(2a+3)}. \quad (17)$$

We see that the CPA of the first bidder does not depend on the second bidder strategy. Therefore, we have in this case something similar to a strategy dominance.

A.2 Numerical estimation of the competition factor

The trick is to remark that γ is equal to the ratio of cost and value when $T = 1$. We derive the following algorithm to estimate γ for different F and n :

- (1) Choose n , F and a number of samples N
- (2) $Cost = 0$; $Value = 0$;
- (3) Generate $v \sim F$ and $v^- \sim F^n$
- (4) If $v > v^-$ then: $Cost+ = v^-$ and $Value+ = v$
- (5) Repeat from 3. $(N-1)$ times
- (6) $\gamma \approx Cost/Value$

The results are displayed in Figure 4. For the log-normal distribution, the results are similar to those of the power distribution: monotony and convergence to 1. It seems that the bigger the variance, the bigger γ . Without providing a proof, we see that we can get an insightful intuition from Equation 6 of the next section.

The exponential distribution provides more exotic results: γ seems to be only a function of n (i.e. independent of the scale parameter). This can be easily proved analytically using the formula $\mathbb{E}(\max(v_1, v_2)) = \frac{1}{\lambda_1} + \frac{1}{\lambda_2} - \frac{1}{\lambda_1 + \lambda_2}$ and a recursion on $v_1^{(n)}$. It is however still increasing in n , and seems to converge to 1 as n goes to infinity.

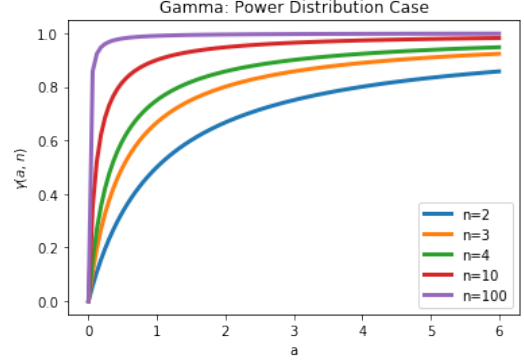
B PROOF OF THEOREM 3.1

Set $b(\alpha, v) = \alpha v$ for $\alpha > 0$. If $b(\alpha, v)$ is admissible for any α , then we can set $\lambda = 0$ and we are done (edge case). Else, there exists α^+ such that $\gamma(\alpha^+) > 0$, where $\gamma(\alpha) := \mathbb{E}_v[b^{-i} < b(\alpha, v)](b^{-i} - T^i v)$

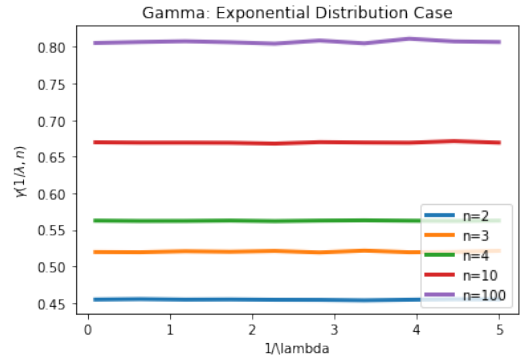
Moreover:

- (1) since $\alpha \rightarrow \int_{b^{-i}/\alpha}^1 (b^{-i} - T^i v) f^i(v) dv$ is continuous for any $b^{-i} > 0$, so is γ ,
- (2) for α smaller than T^i , $\gamma(\alpha) \leq 0$.

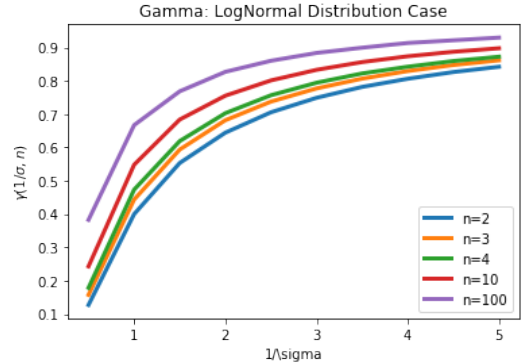
So there exists α^* satisfying $\gamma(\alpha^*) = 0$.



(a) Power Distribution



(b) Exponential Distribution



(c) LogNormal Distribution

Figure 4: The scaling factor as a function of the distribution parameter, for different numbers of bidders.

If we set λ^* so that $\frac{1}{\lambda^*} + T^i = \alpha^*$, then by strategic dominance of the truthful strategy in a second price auction, $b(\alpha^*, v)$ solves the unconstrained optimization problem of maximizing

$$\mathbb{E}_{b^-, v}[b^- < b](v - \lambda^*(b^- - T^i v)). \quad (18)$$

We can also assert that $b(\alpha^*, \cdot)$ is also an optimal solution for the constrained optimization problem (1). Why? First observe that $b(\alpha^*, \cdot)$ is admissible since $\gamma(\alpha^*) = 0$, second, let \hat{b} be an optimal

solution to (1), then, if we write $k(b) = \mathbb{E}_{b^-, v}[b^- < b]v$, we have that

$$k(b(\alpha^*)) - \lambda^* \gamma(\alpha^*) \geq k(\hat{b}) - \lambda^* \mathbb{E}_v[b^{-i} < \hat{b}](b^{-i} - T^i v) \quad (19)$$

thus since $\gamma(\alpha^*) = 0$ and $\mathbb{E}_v[b^{-i} < \hat{b}](b^{-i} - T^i v) \leq 0$

$$k(b(\alpha^*)) \geq k(\hat{b}). \quad (20)$$

C PROOF OF THEOREM 3.3

To derive the first order condition, we follow the same path as in the proof of Theorem 3.4. We thus only discuss the uniqueness. We start with a small technical consideration: observe that the bids can be modified on a zero measure set without impacting the profit. So when we say unique, we, in fact, mean essentially unique.

We denote by b a symmetric equilibrium bidding strategy. It induces a distribution G of the price to beat. It is clear that we cannot be in the edge case envisioned in the previous proof. So take α^* as introduced in the previous proof. Observe that to be a best reply b has to maximize (18). Assume there exists v_0 such that $\alpha v \neq b(v)$ on a neighborhood V of v_0 . Since $g(b(v)) > 0$ for $v \in V$, αv is strictly dominating $b(v)$ in the optimization of (18), which is a contradiction.

D PROOF OF THEOREM 3.4

First observe that, the Lagrangian of the bidder optimization problems writes:

$$\lambda \mathbb{E}_{v, b^-} ((1/\lambda + T)v - S(b(v), b^-)) [b > b^-],$$

thus we can look for a point-wise maximizer. Denote by (\hat{b}, \hat{b}^-) the equilibrium bid of the same auction without any CPA constraints. Observe that if the price to beat is distributed like $(1/\lambda + T)\hat{b}^-(v^-)$ (for a given $\lambda > 0$), then:

$$\begin{aligned} b^*(v) &= \operatorname{argmax}_b \mathbb{E}_{v^-} (1/\lambda + T)v - S(b, b^-(v^-)) \\ &= (1/\lambda + T) \operatorname{argmax}_b \mathbb{E}_{v^-} (1/\lambda + T)v - S((1/\lambda + T)b, (1/\lambda + T)\hat{b}^-(v^-)) \\ &= (1/\lambda + T) \operatorname{argmax}_b \mathbb{E}_{v^-} v - S(b, \hat{b}^-(v^-)) \\ &= (1/\lambda + T)\hat{b}(v). \end{aligned}$$

Therefore, for a given λ , if the competition uses the proposed linear bid, one should also bid linearly to maximize the Lagrangian. We denote by G the price to beat distribution and by κ the degree of "second priceness": $S(b, b^-) = ((1 - \kappa)b + \kappa b^-)[b > b^-]$. Observe that the first term of the integrand in the Lagrangian definition is equal to $\mathbb{E}_v b_\lambda G(b_\lambda)$. So, let us deal with the second term. Since

$$\begin{aligned} \mathbb{E}_{v, v^-} S(b_\lambda(v), b^-(v^-)) [v > v^-] &= \mathbb{E}_v \int_0^{b_\lambda(v)} S(b_\lambda(v), t) g(t) dt \\ &= \mathbb{E}_v S(b_\lambda(v), b_\lambda(v)) G(b) - \int_0^{b_\lambda(v)} \kappa G(t) dt \\ &= \mathbb{E}_v b_\lambda G(b_\lambda) - \int_0^{b_\lambda(v)} \kappa G(t) dt, \end{aligned}$$

we get that

$$L = \lambda \kappa \mathbb{E}_v \int_0^{b_\lambda(v)} G(t) dt.$$

The first order condition on λ writes

$$0 = \mathbb{E}_{v, b^-} ((1/\lambda + T)v - S(b(v), b^-)) [b > b^-] - \kappa/\lambda \mathbb{E}_v G(b_\lambda(v)) v dv.$$

Then using successively the homogeneity of S , a simplification by κ , and the de facto definition of γ as $\lambda T/(1 + \lambda T)$, we get

$$\gamma = \frac{\mathbb{E} v^- [v > v^-]}{\mathbb{E} v [v > v^-]},$$

which we simplify using the following formula on the order statistics (see [4]): $\mathbb{E} Y_2^{(n)} = \mathbb{E} n Y_1^{(n-1)} - (n-1) Y_1^{(n)}$.