

---

# Real-Time Optimisation for Online Learning in Auctions

---

Lorenzo Croissant<sup>1</sup> Marc Abeille<sup>1</sup> Clément Calauzènes<sup>1</sup>

## Abstract

In display advertising, a small group of sellers and bidders face each other in up to  $10^{12}$  auctions a day. In this context, revenue maximisation via *monopoly price* learning is a high-value problem for sellers. By nature, these auctions are *online* and produce a very high frequency stream of data. This results in a computational strain that requires algorithms be *real-time*. Unfortunately, existing methods inherited from the batch setting suffer  $\mathcal{O}(\sqrt{t})$  time/memory complexity at each update, prohibiting their use. In this paper, we provide the first algorithm for online learning of monopoly prices in online auctions whose update is constant in time and memory.

## Introduction

Over the last two decades, online display advertising has become a key monetisation stream for many businesses. The market for the trading of these ads is controlled by a very small ( $< 10$ ) number of large intermediaries who buy and sell at auction, which means that a seller-buyer pair might trade together in  $10^{10}$  to  $10^{12}$  auctions per day. Repeated auctions on this scale raise the stakes of revenue maximisation, while making computational efficiency a key consideration. In his 1981 seminal work on revenue maximisation, Myerson described *the* revenue-maximising auction when buyers' bid distributions are known. In the context of online display ads these distributions are private, but the large volume of data collected by sellers on buyers opens the way to *learning* revenue maximising auctions.

**Optimal vs. tractable.** The learning problem associated with the Myerson auction has infinite pseudo-dimension (Morgenstern & Roughgarden, 2015), making it unlearnable (Pollard, 1984). 2<sup>nd</sup>-price auctions with personalised reserve prices (*i.e.* different for each bidder) stand as the

commonly accepted compromise between optimality and tractability. They provide a 2-approximation (Roughgarden & Wang, 2016) to the revenue of the Myerson auction while securing finite pseudo-dimension.

**Monopoly prices.** 2<sup>nd</sup>-price auctions with personalised reserves can be either *eager* or *lazy*. In the *eager* format, the item goes to the highest bidder *amongst* those who cleared their reserve prices and goes unsold if none of them did. In the *lazy* format, the item goes to the highest bidder *if* he cleared his reserve price and goes unsold if he did not. While an optimised *eager* version leads to better revenue than an optimised *lazy* version, solving the eager auction's associated Empirical Risk Minimisation (ERM) problem is NP-hard (Paes Leme et al., 2016) and even APX-hard (Roughgarden & Wang, 2016). In contrast, solving the ERM for the *lazy* version can be done in polynomial time (Roughgarden & Wang, 2016): it amounts to computing a bidder-specific quantity called the *monopoly price*. Not only is the *monopoly price* the optimal reserve in the *lazy* 2<sup>nd</sup>-price auction, but it is also a provably good reserve in the *eager* one (Roughgarden & Wang, 2016), and the optimal reserve in posted-price (Paes Leme et al., 2016). This makes learning monopoly prices for revenue maximisation an important and popular research direction.

**Online learning.** Finding the monopoly price in a repeated 2<sup>nd</sup>-price auction is a natural sequential decision problem based on the incoming bids. All three aforementioned settings relating to the monopoly price have been studied: posted-price (Amin et al., 2014; Blum et al., 2004), eager (Cesa-Bianchi et al., 2014; Roughgarden & Wang, 2016; Kleinberg & Leighton, 2003), and lazy which we study (Blum & Hartline, 2005; Blum et al., 2004; Mohri & Medina, 2016; Rudolph et al., 2016; Bubeck et al., 2017; Shen et al., 2019). Each setting also corresponds to a different observability structure. The offline problems are well understood, but no online method offers the  $\mathcal{O}(1)$  *efficiency* crucial for real-world settings. We focus, therefore, on the key problem of learning monopoly prices, online and efficiently, in the stationary and non-stationary cases.

**Structure of the paper.** We propose a real-time first-order algorithm which makes online learning of monopoly prices computationally feasible, when interacting with sta-

---

<sup>1</sup>Criteo AI Lab, Paris, France. Correspondence to: Lorenzo Croissant <ld.croissant@criteo.com>.

tionary and non-stationary buyers. In Sec. 1, we detail the setting and problem we consider. We review, in Sec. 2, the existing approaches and stress the challenges of the problem including overcoming computational complexities. Our approach, based on convolution and the  $\mathcal{O}(1)$  Online Gradient Ascent algorithm, is described in Sec. 3. We study performance for stationary bidders in Sec. 4 with  $1/\sqrt{t}$  convergence rate to the monopoly price, and for non-stationary bidders in Sec. 5 with  $\mathcal{O}(\sqrt{T})$  dynamic regret.

## 1. Setting

A key property of a personalised reserve price in a lazy second price auction is that it can be optimised separately for each bidder (Paes Leme et al., 2016). For a bidder with bid cdf  $F$ , the optimal reserve price is the monopoly price  $r^*$ , i.e. the maximiser of the *monopoly revenue* defined as

$$\Pi^F(r) = r(1 - F(r)). \quad (1)$$

Thus, without loss of generality, we study each bidder separately in the following repeated game: the seller sets a reserve price  $r$  and simultaneously the buyer submits a bid  $b \in [0, \bar{b}]$  drawn from his private distribution  $F$ , whose pdf is  $f$ . The seller then observes  $b$  which determines the instantaneous revenue

$$p(r, b) = r\mathbb{1}_{r \leq b} \quad \text{with} \quad \mathbb{E}_F[p(r, b)] = \Pi^F(r). \quad (2)$$

In this work, we consider two settings, depending whether the bid distribution is stationary or not.

**Stationary setting.** Stationarity here means  $F$  is fixed for the whole game. We thus have a stream of i.i.d. bids from  $F$ , where the seller aims to maximise her long term revenue

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T p(r, b_t) = \Pi^F(r).$$

Or, equivalently, tries to construct a sequence of reserve prices  $\{r_t\}_{t \geq 1}$  given the information available so far encoded in the filtration  $\mathcal{F}_t = \sigma(r_1, b_1, \dots, b_{t-1})$  such that  $\Pi^F(r_t) \rightarrow \Pi^F(r^*)$  as fast as possible.

**Non-stationary setting.** In real-world applications, bid distributions may change over time based on the current context. For example, near Christmas the overall value of advertising might go up since customers spend more readily, and thus bids might increase. The bidder could also refactor his bidding policy for reasons entirely independent of the seller. We relax the stationarity assumption by allowing bids to be drawn according to a sequence of distributions  $\{F_t\}_{t \geq 1}$  that varies over time. As a result, the monopoly prices  $\{r_t^*\}_{t \geq 1}$  and optimal monopoly revenues

$\{\Pi^{F_t}(r_t^*)\}_{t \geq 1}$  fluctuate and convergence is no longer defined. Instead, we evaluate the performance of an adaptive sequence of reserve prices by its expected dynamic regret

$$R(T) = \mathbb{E} \left( \sum_{t=1}^T \Pi^{F_t}(r_t^*) - \Pi^{F_t}(r_t) \right), \quad (3)$$

and our objective is to track the monopoly price as fast as possible to minimise the dynamic regret.

## 2. Related Work, Challenges and Contributions

### 2.1. Related Work

Lazy 2<sup>nd</sup>-price auctions have been studied both in batch (Mohri & Medina, 2016; Shen et al., 2019; Rudolph et al., 2016; Paes Leme et al., 2016) and online (Blum & Hartline, 2005; Blum et al., 2004; Bubeck et al., 2017) settings. All existing approaches aim to optimise, at least up to a precision of  $1/\sqrt{t}$ , the ERM objective

$$\Pi^{\hat{F}_t}(r) = r(1 - \hat{F}_t(r)) = \frac{1}{t} \sum_{i=1}^t r\mathbb{1}_{r \leq b_i}. \quad (4)$$

However, regardless of how well-behaved  $\Pi^F$  is,  $\Pi^{\hat{F}_t}$  is very poorly behaved for optimisation: it is non-smooth, non-quasi-concave, discontinuous, and is increasing everywhere (see Fig.1, center, dashed). Thus direct optimisation with first order methods is not applicable. Some attempts have been made in the batch setting to optimise surrogate objectives, but ended up with an irreducible bias (Rudolph et al., 2016) or with hyper-parameters whose tuning is as hard as the initial problem (Shen et al., 2019). The classical approach relies on sorting the bids  $\{b_i\}_{i=1}^t$  to be able to enumerate  $\Pi^{\hat{F}_t}(\cdot)$  linearly over  $\{b_i\}_{i=1}^t$  (Mohri & Medina, 2016; Paes Leme et al., 2016). A popular improvement in terms of complexity, especially used in online approaches (Blum et al., 2004; Blum & Hartline, 2005) consists in applying the same principle on a regular grid of resolution  $1/\sqrt{t}$ , which in the end provides an update with complexity  $\mathcal{O}(\sqrt{t})$  and a memory requirement of  $\mathcal{O}(\sqrt{t})$ .

This idea of discretising the bid space was widely adopted in partially observable settings – e.g. online *eager* or online posted-price auctions – as it reduces the problem to a multi-armed bandit with its well-studied algorithms (Kleinberg & Leighton, 2003; Cesa-Bianchi et al., 2014; Roughgarden & Wang, 2016) at the price of still suffering the same update and memory complexities of  $\mathcal{O}(\sqrt{t})$ .

Numerous approaches with adversarial bandits also followed this discretisation approach to adapt Exp3/Exp4 (Kleinberg & Leighton, 2003; Cohen et al., 2016; Bubeck

et al., 2017) to all the settings. As a bonus, it also allows to handle the case of non-stationary bidders. However, the work of Amin et al. (2014) stresses that bidders cannot behave in an arbitrary way, as they optimise their own objective that is not incompatible with the seller’s<sup>1</sup>. Hence, the non-stationarity mostly comes from the item’s value changing over time. This suggest adapting a regular stochastic algorithms (ERM, UCB...), e.g. using sliding windows (Garivier & Moulines, 2011; Lattimore & Szepesvári, 2018).

Non-smooth or non-differentiable objectives such as  $\Pi^{\hat{F}}$  have been studied in both stochastic and 0-order optimisation. In both, convolution smoothing has been employed to circumvent these problems. In Duchi et al. (2012) stochastic gradient with decreasing convolution smoothing is studied for the convex case. Unfortunately, very few distributions yield a concave  $\Pi^F$ . In 0-order optimisation, the only feedback received for an input is the value of the objective at that input. In this setting, Flaxman et al. (2005) perturb their inputs to estimate a convolved gradient. In contrast, we obtain a closed form and do not need to perturb inputs.

## 2.2. Challenges

The directing challenge of our line of work is to devise an online learning algorithm for monopoly prices with minimal cost, to handle very large real-world data streams. With  $10^{10}$  daily interactions in one seller-bidder pair, it is acceptable to forfeit some convergence speed in exchange for feasibility of the algorithm. It is not possible to accept update complexity or memory requirement scaling with  $t$ . Our objective is thus to find a method that **1)** converges to  $r^*$  in the stationary setting or has a low regret  $R(T)$  in the non-stationary one, **2)** has  $\mathcal{O}(1)$  memory footprint, and **3)** computes the next reserve  $r_{t+1}$  with  $\mathcal{O}(1)$  computations.

**Computational Complexity.** Unfortunately, none of the previously proposed methods fit these requirements. On one hand, all methods based on solving ERM by sorting (Cesa-Bianchi et al., 2014; Roughgarden & Wang, 2016) need to keep all past bids in memory ( $\mathcal{O}(t)$  dependence) and their update steps require at best  $\mathcal{O}(\sqrt{t})$  computations. On the other hand, adversarial methods such as Exp3 or Exp4 (Cohen et al., 2016; Bubeck et al., 2017) are designed for finite action spaces and thus need to discretise  $[0, \bar{b}]$  into  $\sqrt{t}$  intervals (to keep their regret guarantees), also leading to a complexity of  $\mathcal{O}(\sqrt{t})$  from sampling to compute  $r_{t+1}$ .

**Gradient bias.** First order methods (e.g. Online Gradient Ascent a.k.a. OGA) are standard tools in online learning and enjoy  $\mathcal{O}(1)$  update and memory. This makes them

<sup>1</sup>An auction is not a zero-sum game: if the item goes unsold, neither player receives payoff.

great candidates for our problem. OGA requires 3 ingredients to converge: an objective whose gradients always point towards the optimum<sup>2</sup>, and a gradient estimator which has bounded variance and is unbiased. Unfortunately, discontinuity of  $p$  makes  $\nabla p$  a *biased* estimator of  $\nabla \Pi^F$ . A natural approach is to construct a surrogate for  $p$  which has unbiased gradients and preserves the other two conditions.

**Surrogate consistency.** Optimising a surrogate objective inherently creates a *bias*, which has to be reduced over-time. To do so without breaking the convergence of OGA, we must conduct a careful finite time analysis of the algorithm, which is an *analytical* challenge. We must re-analyse classical results (e.g. Bach & Moulines (2011); Duchi et al. (2012)) for varying objectives: the challenge is to design a bias reduction procedure, and then integrate it into these proofs to show we preserve consistency.

**Non-stationarity.** Resolving the above challenges is sufficient to achieve efficient convergence in the stationary setting. However, it is not sufficient in order to track non-stationary bid distributions. Taking a constant surrogate and learning rate, it is possible to adapt the stationary solution to the non-stationary case and keep its computational efficiency. The challenge is to devise this adaptation, and then to derive (sub-linear) regret for it.

## 2.3. Contributions

We propose a smoothing method for creating surrogates in pseudo-concave problems with biased gradients. We use it to create a first-order real-time optimisation algorithm which reduces the surrogate’s bias during optimisation. We prove convergence and give rates in the stationary setting and dynamic regret bounds for tracking. In more detail:

**Smooth surrogates for first order methods.** We first translate standard auction theory assumptions (e.g. increasing virtual value) into properties of generalised concavity of the monopoly revenue (Prop. 1). Next, we introduce our smoothing method and show (in Prop. 2) that it preserves the properties from Prop. 1 while offering arbitrary smoothness, which fixes the biased gradient problem. Finally, we provide controls, via the choice of the kernel, on the bias and variance of the gradient estimates of our surrogate, which is now ready for OGA (Prop. 3).

**Consistent algorithm for stationary bidders.** We construct an algorithm (V-CONV-OGA) which performs gradient ascent while simultaneously decreasing the strength of the smoothing over time, reducing the bias to zero. As a result our algorithm almost surely converges to the

<sup>2</sup>Pseudo-concave or variationally coherent.

monopoly price (Thm. 1) while enjoying computational efficiency. Further, under a minimum curvature assumption, we provide the rate of convergence and optimal tuning parameters (Thm. 2 and Cor. 1). At the cost of a slight degradation in convergence speed (from  $\mathcal{O}(1/t)$  to  $\mathcal{O}(1/\sqrt{t})$ ), our algorithm has update and memory complexity of  $\mathcal{O}(1)$  which is vital for real-world applications. Results are summarised in Table 1.

	Update	Memory	Convergence
ERM	$\mathcal{O}(t)$	$\mathcal{O}(t)$	$\mathcal{O}(1/t)$
Discrete ERM	$\mathcal{O}(\sqrt{t})$	$\mathcal{O}(\sqrt{t})$	$\mathcal{O}(1/t)$
V-CONV-OGA	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1/\sqrt{t})$

Table 1. Comparison of our method (V-CONV-OGA) against solving ERM at each step and ERM discretised on a grid of resolution  $1/\sqrt{t}$ , in terms of complexity and convergence – i.e.  $\|\Pi^F(r_t) - \Pi^F(r^*)\|$ .

**Tracking for non-stationary bidders.** Contrary to the stationary setting, when tracking we do not decrease the strength of the smoothing over time. When the bias created is smaller than the noise, our algorithm can still achieve sub-linear dynamic regret when tracking changing bid distributions. For reasonably varying distribution, we show a regret bound of  $\mathcal{O}(\sqrt{T})$  (see Thm. 3 and Cor. 2).

### 3. Smooth Surrogate for First Order Methods

Our objective, to reiterate, is to design an online optimisation procedure to learn or track the optimal reserve price whose updates require  $\mathcal{O}(1)$  computational and memory cost. To this end, we focus on first order method and consider vanilla Online Gradient Ascent. Unfortunately, the specific problem of learning a monopoly price doesn’t provide a way to compute unbiased gradient estimates for  $\Pi^F$  from bid samples. We therefore want to design a surrogate that makes  $p$  sufficiently smooth so that differentiation and integration commute. This is a well known property of convolutional smoothing, suggesting its use. In addition, we must ensure our surrogate preserves the optimisation properties that  $\Pi^F$  already has. These must thus be studied first, before smoothing to obtain a surrogate.

#### 3.1. Properties of the Monopoly Revenue

The standard assumptions of auction theory are made to guarantee that the monopoly price exists – and that the optimisation problem is well-posed. This is generally stated as “the monopoly revenue is *quasi-concave*”. We refine this characterisation by translating the assumptions we make into specific concavity properties of the monopoly revenue in Prop 1.

**(A1).**  $F \in \mathcal{C}^2([0, \bar{b}])$  and  $f = F' > 0$  on  $(0, \bar{b})$ .

**(A2).**  $F$  is strictly regular on its domain of definition i.e. the virtual value  $\psi(b) = b - \frac{1-F(b)}{f(b)}$  is increasing.

**(A3).**  $F$  has strongly increasing hazard rate on its domain i.e. the hazard rate  $\lambda(b) = \frac{f(b)}{1-F(b)}$  satisfies:

$$\forall 0 \leq b_1 \leq b_2 \leq \bar{b}, \lambda(b_2) - \lambda(b_1) \geq \mu(b_2 - b_1).$$

(A1) is made for ease of exposition. (A2) is a standard auction theory assumption (see Krishna (2009) for a review), and implies a pseudo-concave revenue, as shown by Prop 1. (A2) is satisfied by common distributions, exhaustively listed in (Ewerhart, 2013) and by real-world data – see e.g. (Ostrovsky & Schwarz, 2011). (A3) strengthens (A2) by requiring a minimum curvature around the maximum.

**Proposition 1.** Let  $F$  satisfy (A1) and  $\Pi^F$  be its associated monopoly revenue. Then,  $\Pi^F \in \mathcal{C}^2([0, \bar{b}])$ ,  $\Pi^F > 0$  on  $(0, \bar{b})$  and:

- under (A2),  $\Pi^F$  is strictly pseudo-concave,
- under (A3),  $\Pi^F$  is  $\mu$ -strongly log-concave, i.e.  $\log \Pi^F$  is  $\mu$ -strongly concave on  $(0, \bar{b})$ .

#### 3.2. A Method Based on Smoothing

Prop. 1 ensures that the first condition for the convergence of OGA is met under standard assumptions (A2) or (A3). The main difficulty in the way of using OGA for revenue optimisation – it must be stressed – lies in the undesirable shape of the instantaneous revenue  $p$ . Indeed,  $p$  is non-smooth (discontinuous even) and cannot be used to construct an unbiased estimate of  $\nabla \Pi^F(r)$ , which is necessary for first order-methods.

Mohri & Medina (2016) suggests replacing  $p(\cdot, b)$  by a continuous upper bound. This surrogate can be used for OGA, but it has potentially large areas of zero-gradient, which means it doesn’t learn from all samples. We give a general surrogate construction (based on convolutional smoothing) which **1**) can approximate the original monopoly revenue to arbitrary accuracy, **2**) preserves the concavity properties of  $\Pi^F$ , **3**) offers the desired level of smoothness, and **4**) exhibits no areas of zero gradient.

Formally, given a kernel  $k$  (a metaparameter), we use convolution smoothing to create surrogates for  $p$  and  $\Pi^F$ :

$$p_k(r, b) = (p(\cdot, b) \star k)(r), \quad \Pi_k^F(r) = (\Pi^F \star k)(r). \quad (5)$$

This smoothing guarantees that  $\nabla p_k(\cdot, b)$  is an unbiased estimate of  $\nabla \Pi_k^F$ . On Fig. 1, we illustrate the effect of this smoothing on  $p$ ,  $\Pi^F$ , and  $\Pi_k^F$ . We introduce a set  $\mathcal{K}$  of admissible kernels which contains all strictly positive, strictly log-concave,  $\mathcal{C}^1(\mathbb{R})$ ,  $\mathcal{L}^1(\mathbb{R})$ , normalised (i.e.  $\int_{\mathbb{R}} k(x)dx = 1$ ) functions.  $\mathcal{K}$  contains a large family of




Figure 1. The effect of smoothing the monopoly revenue of a bidder with  $F$  a Kumaraswamy<sup>4</sup>(1,0.4) distribution with a Gaussian kernel. Left: smoothing of  $p(r, b)$  for  $b = 0.7$ . Center: smoothing of the empirical revenue  $\Pi_t^F(r)$  for some random values of  $b_t$ . Right: smoothing of the expected revenue ( $\Pi_k^F$  vs  $\Pi^F$ ). Note the difference in the optima (*i.e.* the surrogate bias).

kernels, including standard smoothing ones such as Gaussians and mollifiers. Prop. 2 shows that convolution with elements of  $\mathcal{K}$  preserves pseudo- and log-concavity.

**Proposition 2.** *Let  $F$  satisfy (A1) and  $\Pi^F$  be its associated monopoly revenue. Let  $k \in \mathcal{K}$ , then:*

- $\Pi_k^F$  and  $p_k$  are  $\mathcal{C}^1(\mathbb{R})$ ,
- $\Pi_k^F(r) = \mathbb{E}_F(p_k(r, b))$  and  $\nabla \Pi_k^F(r) = \mathbb{E}_F(\nabla p_k(r, b))$ ,
- under (A2),  $\Pi_k^F$  is strictly pseudo-concave on  $\mathbb{R}$ ,
- under (A3),  $\Pi_k^F$  is strictly log-concave on  $\mathbb{R}$ .

*Proof.* See App. B.2.  $\square$

---

#### Algorithm 1 CONV-OGA

---

**input:**  $r_0, \{\gamma_t\}_{t \in \mathbb{N}}, k \in \mathcal{K}, \mathcal{C} \subset [0, \bar{b}]$   
**for**  $t = 1$  **to**  $+\infty$  **do**  
 | observe  $b_t$   $r_t \leftarrow \text{proj}_{\mathcal{C}}(r_{t-1} + \gamma_t \nabla p_k(r_{t-1}, b_t))$

---

Prop. 2 guarantees that the surrogate satisfies the pseudo-concavity and unbiased gradient conditions of OGA. Applying OGA to the surrogate  $\Pi_k^F$  gives Alg. 1. Note that as a property of convolution,  $\nabla p_k(\cdot, b) = p(\cdot, b) \star \nabla k$ , which is a simple (generally closed-form) computation.

Prop. 3 will show that the bounded variance condition of OGA also holds. Since OGA’s three conditions are satisfied, we can guarantee convergence to the maximum  $r_k^*$  of  $\Pi_k^F$  (see *e.g.* (Bottou, 1998)). However, in general  $r_k^*$  is not the monopoly price  $r^*$  and the surrogate is biased. Prop. 3 also gives a control on this bias in terms of the  $L^1$  distance between the cdf  $K$  of  $k$  and the cdf  $\mathbb{1}_{\mathbb{R}^+}$  of the Dirac mass  $\delta_0$ , which is the only kernel to guarantee  $r_{\delta_0}^* = r^*$ .

<sup>4</sup>This distribution satisfies our concavity assumptions and can display highly eccentric behaviour for easy visualisation of the impact of the surrogate.

**Proposition 3.** *Let  $F$  satisfy (A1),  $k \in \mathcal{K}$ . Let  $r^*$  and  $r_k^*$  be the monopoly prices associated with  $\Pi^F$  and  $\Pi_k^F$ . Then, the bias  $B_k = |\Pi^F(r^*) - \Pi^F(r_k^*)|$  and the instantaneous convolved gradient second moment  $V_k = \max_{r \geq 0} \mathbb{E}_{b \sim F}(|\nabla p_k(r, b)|^2)$  are upper bounded by*

- $B_k \leq 2 \|\nabla \Pi^F\|_{\infty} \|K - \mathbb{1}_{\mathbb{R}^+}\|_1$ ,
- $V_k \leq 1 + \bar{b}(1 + \|\nabla \Pi^F\|_{\infty}) \|k\|_{\infty}$ .

If one chooses a family of kernels, these bounds can be expressed in terms of its parameters. For instance, when  $k$  is zero-mean Gaussian with variance  $\sigma^2$ , one easily recovers:

$$\|K - \mathbb{1}_{\mathbb{R}^+}\|_1 = \sigma \sqrt{2/\pi}, \quad \|k\|_{\infty} = (\sqrt{2\pi}\sigma)^{-1}. \quad (6)$$

CONV-OGA converges only to  $r_k^*$ . To remedy this, we would like to decrease  $B_k$  over time by letting  $k \rightarrow \delta_0$ . However, since  $\|k\|_{\infty}$  diverges as  $k \rightarrow \delta_0$ , we will have to tread carefully in our analysis which occupies the next section.

## 4. Convergence with a Stationary Bidder

To decrease  $B_k$  over time, we introducing a decaying kernel sequence  $\{k_t\}_{t \in \mathbb{N}}$  into CONV-OGA, giving V-CONV-OGA (Alg. 2). This section will demonstrate its consistency and convergence by controlling the trade-off between bias  $B_k$  and variance  $V_k$ , as  $B_k$  is reduced to zero over time. This trade-off decomposes the total error as:

$$\Pi^F(r^*) - \Pi^F(r_t) = \underbrace{\Pi^F(r^*) - \Pi^F(r_k^*)}_{(\text{surrogate bias})} + \underbrace{\Pi^F(r_k^*) - \Pi^F(r_t)}_{(\text{estimation})}.$$

This stresses that the kernel should converge to  $\delta_0$  *fast enough* to cancel the bias  $B_k$ , yet *slowly enough* to control  $V_k$  and preserve the convergence speed of OGA.

### 4.1. General Convergence Result

Thm. 1 provides sufficient conditions on the schedules of  $k_t$  and  $\gamma_t$  that guarantees V-CONV-OGA converges *a.s.* to

**Algorithm 2** V-CONV-OGA

**input:**  $r_0, \{\gamma_t\}_{t \in \mathbb{N}}, \{k_t\}_{t \in \mathbb{N}} \in \mathcal{K}^{\mathbb{N}}, \mathcal{C} \subset [0, \bar{b}]$   
**for**  $t = 1$  **to**  $+\infty$  **do**  
 | observe  $b_t$   $r_t \leftarrow \text{proj}_{\mathcal{C}}(r_{t-1} + \gamma_t \nabla p_{k_t}(r_{t-1}, b_t))$

$r^*$ . It is derived by adapting stochastic optimisation results (see e.g. Bottou (1998)) to the changing objective  $\Pi_{k_t}^F$ .

**Theorem 1.** Let  $F$  satisfy (A1) and (A2) and  $\{k_t\}_{t \in \mathbb{N}} \in \mathcal{K}^{\mathbb{N}}$ . Then, by running V-CONV-OGA with  $\mathcal{C} = [0, \bar{b}]$ , we have

$$r_t \xrightarrow{\text{a.s.}} r^*$$

as long as  $\sum_{t=1}^{+\infty} \gamma_t = +\infty$ ,  $\sum_{t=1}^{+\infty} \gamma_t \|K_t - \mathbb{1}_{\mathbb{R}^+}\|_1 < +\infty$  and  $\sum_{t=1}^{+\infty} \gamma_t^2 \|k_t\|_{\infty} < +\infty$ .

*Proof Sketch.* The proof relies on decomposing the error  $\|r_t - r^*\|^2$  into three terms related respectively to the pseudo-concavity of  $\Pi^F$ , the bias  $B_{k_t}$ , and the instantaneous gradient second moment  $V_{k_t}$ . Then, following Bottou (1998), we use a quasi-martingale argument to ensure the convergence of the stochastic error process. The full proof is available in App. C.1.  $\square$

If a constant kernel sequence were to be used in V-CONV-OGA, we would recover the usual stochastic approximation conditions on the step size  $\gamma_t$ , namely that  $\sum_{t=1}^{\infty} \gamma_t = +\infty$  and  $\sum_{t=1}^{\infty} \gamma_t^2 < +\infty$ . This suggests setting  $\gamma_t \propto 1/t$ . For such a choice of step-size, Thm. 1 asserts convergence if  $\sum_{t=1}^{\infty} \gamma_t \|K_t - \mathbb{1}_{\mathbb{R}^+}\|_1 < +\infty$ , which is guaranteed by  $k_t \rightarrow \delta_0$ . This means  $\|k\|_{\infty} \rightarrow \infty$  as  $t \rightarrow \infty$ , but  $\sum_{t=1}^{+\infty} \gamma_t^2 \|k_t\|_{\infty} < +\infty$  tells us explicitly how slow our decay must be in terms of the family of kernels. For example in the case of a Gaussian kernel, (6) implies that a suitable choice of kernel variance is  $\sigma_t \propto t^{-\alpha}$  for  $\alpha \in (0, 1)$ .

## 4.2. Finite-time Convergence Rates

While Thm. 1 provides sufficient conditions on the kernel sequence  $\{k_t\}_{t \in \mathbb{N}}$  for V-CONV-OGA to be consistent, it does not characterise the rate of the convergence, and thus cannot be leveraged to *optimise* the step size  $\gamma_t$  and the decay rate of the kernel.

To obtain finite time guarantees on the rate of convergence, we must impose stronger conditions on the monopoly revenue  $\Pi^F$ . Recall that under (A2),  $\Pi^F$  is strictly pseudo-concave. It is well known that such functions can have large areas of arbitrarily small gradient. Since they can make first order methods arbitrarily slow, no meaningful rate can be obtained for them. Strengthening the assumption to (A3), i.e. excluding vanishing gradients by ensuring  $\Pi^F$  is  $\mu$ -strongly log-concave (see Prop. 1), will give a rate in Thm. 2 under the further technical assumption (A4).

**(A4).** The seller is given a compact subset  $\mathcal{C} \subseteq [0, \bar{b}]$  and a constant  $c > 0$  such that  $r^* \in \mathcal{C}$  and for all  $r \in \mathcal{C}$ ,  $\Pi^F(r) \geq c$ .

(A4) ensures that the seller can lower bound revenue on a compact subset of  $[0, \bar{b}]$ . It should be understood as prior knowledge of the seller based on the format of the auction and the type of item sold.  $\mathcal{C}$  exists for any  $c < \Pi^F(r^*)$ , so this hypothesis is not restrictive relative to (A3).

**Theorem 2.** Let  $F$  satisfy (A1) and (A3), let  $\mathcal{C}$  and  $c$  as in (A4) be given, and let  $\{k_t\}_{t \in \mathbb{N}} \in \mathcal{K}^{\mathbb{N}}$  such that there are  $\nu_1, \nu_{\infty}, \alpha_1, \alpha_{\infty}$  with

$$\|K_t - \mathbb{1}_{\mathbb{R}^+}\|_1 \leq \nu_1 t^{-\alpha_1} \text{ and } \|k_t\|_{\infty} \leq \nu_{\infty} t^{\alpha_{\infty}}.$$

Then, by running V-CONV-OGA on  $\mathcal{C}$  for  $\gamma_t = \nu t^{-\alpha}$  with  $\nu \leq (2c\mu)^{-1}$ , we have for all  $t \geq 2$ ,

$$\begin{aligned} \alpha = 1 : \quad & \mathbb{E}(\|r_t - r^*\|^2) = \tilde{O}(t^{-\alpha_1} + t^{\alpha_{\infty}-1} + t^{-2\mu c\nu}) \\ \alpha \in (0, 1) : \quad & \mathbb{E}(\|r_t - r^*\|^2) = \tilde{O}(t^{-\alpha_1} + t^{\alpha_{\infty}-\alpha}) \\ & + \tilde{O}\left((t^{1-\alpha-\alpha_1} + t^{1+\alpha_{\infty}-2\alpha})e^{-\mu c\nu t^{1-\alpha}}\right) \end{aligned}$$

where  $\tilde{O}$  potentially hides a logarithmic term depending on the values of  $\alpha, \alpha_1$ , and  $\alpha_{\infty}$ .

*Proof Sketch.* The extended statement of Thm. 2 with explicit constants and its proof are detailed in App. C.2. The proof builds on Bach & Moulines (2011, Thm.2), derived for log-concave functions, and is adapted to our varying kernel approach and its changing objective. In contrast with the proof of Thm. 1, where we only leveraged pseudo-concavity, we show here that (A3), together with (A4), guarantees more refined control on the curvature of  $\Pi^F$  around  $r^*$ :

$$\forall r \in \mathcal{C}, \quad (r - r^*) \nabla \Pi^F(r) \leq -\mu c \|r^* - r\|^2. \quad (7)$$

This way we can better control the stochastic process  $\{\|r_t - r^*\|^2\}_{t \in \mathbb{N}}$ : like for Thm. 1, we decompose the error into three terms related to concavity (Eq. 7), bias  $B_{k_t}$ , and instantaneous gradient smoothness  $V_{k_t}$ . The error is then bounded in expectation by manipulating finite series.  $\square$

Thm. 2 show two distinct regimes for both choices of  $\alpha$ : transient ( $t^{-2\mu c\nu}$  and  $e^{-\mu c\nu t^{1-\alpha}}$  resp.) and stationary ( $t^{-\alpha_1} + t^{\alpha_{\infty}-\alpha}$  and  $t^{-\alpha_1} + t^{\alpha_{\infty}-1}$  resp.). On Fig. 2 (top), the transient phase is visible up to  $2 \times 10^3$  steps. Since the transient regime's rate depends only on  $\nu = \gamma_0, c$  known from (A4), and  $\mu$  known from (A3), we can set  $\nu$  to make the *stationary* regime the driver of the rate.

To optimise the stationary regime we face a bias-variance trade-off. Like Thm. 1, Thm. 2 requires that  $k \rightarrow \delta_0$  (via  $\alpha_1 > 0$ ) while imposing a bound on the growth speed

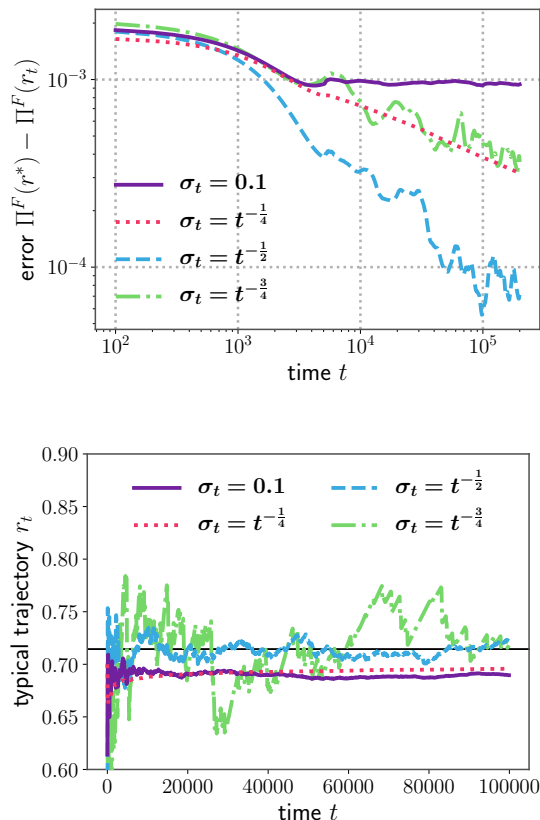


Figure 2. Stationary case. Numerical behaviour of V-CONV-OGA for different  $\sigma_t$  on i.i.d samples from a Kumaraswamy (1, 0.4). Top: averaged convergence speeds of instant regret (log-log scale). Bottom: representative reserve price trajectories.

of  $V_k$  (via  $\alpha_\infty < \alpha$ ). This time, however, we have exact rates which we can use to determine optimal parameters for the trade-off, taking into account the antagonistic effects of  $\alpha_1$  and  $\alpha_\infty$ . From Thm. 2, we recover that the optimal learning rate is  $\gamma_t \propto 1/t$ . To tune the kernels it is sensible to fix a parametric family and tune its parameter(s). For zero-mean Gaussian kernels, we have Cor. 1.

**Corollary 1.** *If we fix  $\gamma_t \propto 1/t$ , and let  $\{k_t\}_{t \in \mathbb{N}}$  be Gaussian  $(0, 1/t)$  kernels in Thm. 2 we have for all  $t \geq 2$  that:*

$$\mathbb{E}(\|r_t - r^*\|^2) = \tilde{O}(t^{-1/2}).$$

*This rate is optimal up to logarithmic factors.*

Fig. 2 demonstrates this optimality: the  $\sigma_t = 1/\sqrt{t}$  (blue) curve is the optimal rate on the top pane, and attains the rate of Cor. 1. The bottom pane illustrates the bias variance trade-off at hand in Thm. 2. If the kernel decays slower than  $1/\sqrt{t}$  (red), the learning rate shrinks much faster and convergence is very slow but very smooth. If  $\sigma_t$  decreases too fast (green) the variance becomes overwhelming and

noise swallows the performance.

The novel analysis of V-CONV-OGA showed its a.s. convergence under (A2), and that with a bit of curvature (A3) and the technical (A4) we could fully characterise its convergence rates. We could thus derive optimal learning rates and place conditions on optimal kernel decay rates. We made the optimal decay rate explicit for Gaussian kernels. This concludes the primary discussion on V-CONV-OGA, and we now move to the non-stationary setting.

## 5. Tracking a Non-stationary Bidder

In practical applications of online auctions, such as display advertising, bidders might change their bid distribution over time. These changes often result from non-stationarity in the private information of bidders. It is therefore beneficial to be able to effectively adapt one’s reserve price over time to *track* changing bid distributions  $\{F_t\}_{t \in \mathbb{N}}$ . We use the dynamic regret  $R(T)$  to measure the quality of an algorithm’s tracking.

The difficulty in the non-stationary setting is to trade-off adaptability (how fast a switch is detected) vs. accuracy (how accurate one is between switches). Convergent algorithms like ERM or V-CONV-OGA will have high accuracy in the first phase, but then suffer as they try to adapt to changes later on, when their learning rate is very small. Windowed methods are more adaptable but still carry with them a lag, directly dependent on their window size. First-order methods like CONV-OGA (with constant learning rate  $\gamma$ ) are much more adaptable, but their convergence rate ( $\mathcal{O}(1/\sqrt{t})$ ) hurts their accuracy. Nevertheless, we show that CONV-OGA is effective, with  $\mathcal{O}(\sqrt{T})$  regret.

The dynamic regret  $R(T)$  cannot be meaningfully controlled for arbitrary sequences  $\{F_t\}_{t \in \mathbb{N}}$ . As such it is customary to assume (A5) that  $\{F_t\}_{t \in \mathbb{N}}$  contains at most  $\tau - 1$  switches up to a horizon  $T$  (see e.g. (Garivier & Moulines, 2011; Lattimore & Szepesvári, 2018)). This corresponds to approximating a slowly changing sequence of  $F_t$  (e.g. Lipschitz) by a piece-wise constant sequence.

(A5). *Given some horizon  $T$ , there exists  $\tau \leq T$  such that  $\sum_{t=1}^{T-1} \mathbb{1}_{F_t \neq F_{t+1}} \leq \tau - 1$ .*

Under (A5), the game (up to  $T$ ) decomposes into  $\tau$  phases. The first step towards controlling the regret is to bound the tracking performance in each phase. We do this in Thm. 3, which shows an incompressible asymptotic error (the bias of our surrogate plus the variance) and a transient phase with exponential decay.

**Theorem 3.** *Let  $F$  satisfy (A1) and (A3) with parameter  $\mu$ , let  $\mathcal{C}$  and  $c$  be as in (A4) and  $k \in \mathcal{K}$ . Then, by running CONV-OGA on  $\mathcal{C}$  with a constant stepsize  $\gamma > 0$ , for any*




Figure 3. Non-stationary case. Tracking by CONV-OGA of three Kumaraswamy distributions (with parameters (1, 4), (1, 0.4), and (1, 1) resp.) with different Gaussian kernels and learning rates.

$t \geq 1$  we have

$$\mathbb{E}(\|r_t - r_k^*\|^2) \leq (\bar{b}^2 + C(\gamma, k)(t-1))e^{-\frac{\mu c \gamma}{2}t} + \frac{2C(\gamma, k)}{\mu c}$$

where  $C(\gamma, k) = \mathcal{O}(\gamma \|K - \mathbb{1}_{\mathbb{R}^+}\|_1 + \gamma^2 \|k\|_\infty)$ .

Thus, immediately after a switch there will be a transient regime of order  $te^{-\frac{\mu c \gamma}{2}t}$  (high adaptability), but afterwards  $r_t$  will oscillate in a band of size  $\frac{2C(\gamma, k)}{\mu c}$  around  $r_k^*$  (low accuracy). We can then use Thm. 3 to derive a sub-linear regret bound given  $T, \tau$  (Cor. 2).

**Corollary 2.** Let  $\{F_t\}_{t \geq 1}$  satisfy (A1), (A3), (A4), (A5) and  $k \in \mathcal{K}$ . Then, there exists  $\Xi(k, \gamma)$  and  $\Omega(k, \gamma)$  such that CONV-OGA has a non-stationary regret of

$$R(T) \leq \Xi(k, \gamma)T + \Omega(k, \gamma)\tau.$$

Further, if the horizon  $T$  is known in advance, running CONV-OGA with  $\gamma = T^{-\frac{1}{2}}$  and  $k$  a kernel with  $\|K - \mathbb{1}_{\mathbb{R}^+}\|_1 \leq T^{-\frac{1}{2}}$  and  $\|k\|_\infty \leq T^{\frac{1}{2}}$ , then  $R(T) = \mathcal{O}(\sqrt{T})$ .

*Proof.* See App. D.  $\square$

Figure 3 illustrates the behaviour of CONV-OGA in a non-stationary environment. In agreement with Thm. 3 and Cor. 2,  $\gamma$  controls the length of the transient regime due to the  $e^{-\frac{\mu c \gamma}{2}t}$  term. Increasing  $\gamma$  shortens it but increases the width of the band of the asymptotic regime as  $C(\gamma, k)$  increases with  $\gamma$  (blue vs. green curves). For a fixed  $\gamma$ , the stationary regime in terms of  $k$  exhibits a bias-variance trade-off:  $\|K - \mathbb{1}_{\mathbb{R}^+}\|_1$  corresponds to the bias and  $\|k\|_\infty$  to the variance (see Prop. 3). In the case of a Gaussian kernel, increasing  $\sigma$  reduces variance but increases bias (green vs. red curve).

CONV-OGA with a constant learning rate is an efficient real-time algorithm for tracking monopoly prices of non-stationary bidders. It incurs  $\mathcal{O}(\sqrt{T})$  regret given the horizon and  $\tau$ , by tuning  $\gamma$  and  $k$ , while maintaining the computational efficiency of online methods.

## 6. Discussion

In this paper we introduced V-CONV-OGA, the first *real-time* ( $\mathcal{O}(1)$  update-time and memory) method for monopoly price learning. We first gave some theoretical results bridging auction theory and optimisation. We then showed how to fix the biased gradient problem with smooth surrogates, giving CONV-OGA. Next, we let the smoothing decrease over time in V-CONV-OGA, for whom we showed convergence of  $\mathcal{O}(1/\sqrt{t})$ . Finally, we adapted CONV-OGA to perform tracking of non-stationary bid distributions with  $\mathcal{O}(\sqrt{T})$  dynamic regret.

**Towards optimal rates.** In the context of high-frequency auctions, computational efficiency trumps numerical precision, so we traded  $\mathcal{O}(\sqrt{t})$  complexity and  $\mathcal{O}(1/t)$  speed for  $\mathcal{O}(1)$  complexity and  $\mathcal{O}(1/\sqrt{t})$  speed. Whether or not it is possible to reach the optimal rate with a real-time algorithm remains an open question. We conjecture this to be impossible in general, but we know it is possible in some instances. If  $F$  is a symmetric distribution, then CONV-OGA with a constant *symmetric* kernel has no bias and  $\mathcal{O}(1/t)$  convergence. Adapting the choice kernel to some *a priori* knowledge on  $F$  is a possible direction to match the optimal rate.

**Extension to stationary bandit.** The second question concerns the extension to partially observable settings, such as online *eager* auctions, when the seller does not observe bids under the reserve. Obviously, extensions using a reduction to multi-armed bandits (UCB, Exp3, Exp4, etc.) via a discretisation of the bid space cannot be *real-time*: the discretisation creates a need for  $\mathcal{O}(\sqrt{t})$  in memory and the same for the update. Yet, it is possible to obtain a straightforward extension of V-CONV-OGA in this setting, by plugging it into an Explore-The-Commit (ETC) (Perchet & Rigollet, 2013) algorithm: V-CONV-OGA learns an estimate of the monopoly price during the exploration period, which is then used during the exploitation period. As for other algorithms, by using a doubling trick to handle an unknown horizon, ETC+V-CONV-OGA exhibits a sub-linear regret. Unfortunately, like in the *lazy* auction setting, the regret is not optimal and the question of whether a *real-time* algorithm can match this optimal regret is still open.



**Extension to non-stationary bandit.** The question of the partially observable setting also applies to non-stationary bidders. In this case, extending CONV-OGA with ETC is no longer straight-forward, as the switching times are unknown. Thus, it is not obvious when to re-trigger an exploration phase of ETC to adapt to the change of the bidder’s distribution. A potential way to tackle this problem could be to use randomised resets for the algorithm (Allesiardo et al., 2017) or change-point detection algorithms to trigger exploration (Hartland et al., 2006).

## References

- Allesiardo, R., Féraud, R., and Maillard, O.-A. The non-stationary stochastic multi-armed bandit problem. *International Journal of Data Science and Analytics*, 3(4): 267–283, Jun 2017.
- Amin, K., Rostamizadeh, A., and Syed, U. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pp. 622–630, 2014.
- Bach, F. and Moulines, E. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, 2011.
- Blum, A. and Hartline, J. D. Near-optimal online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1156–1163. Society for Industrial and Applied Mathematics, 2005.
- Blum, A., Kumar, V., Rudra, A., and Wu, F. Online learning in online auctions. *Theoretical Computer Science*, 324(2-3):137–146, 2004.
- Bottou, L. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998.
- Bubeck, S., Devanur, N. R., Huang, Z., and Niazadeh, R. Online auctions and multi-scale online learning. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 497–514. ACM, 2017.
- Cesa-Bianchi, N., Gentile, C., and Mansour, Y. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2014.
- Cohen, M., Lobel, I., and Leme, R. P. Feature-based dynamic pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016.
- Duchi, J. C., Bartlett, P. L., and Wainwright, M. J. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- Ewerhart, C. Regular type distributions in mechanism design and  $\rho$ -concavity. *Economic Theory*, 53(3):591–603, 2013.
- Flaxman, A. D., Kalai, A. T., and McMahan, H. B. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 05*, pp. 385394, USA, 2005. Society for Industrial and Applied Mathematics. ISBN 0898715857.
- Garivier, A. and Moulines, E. On upper-confidence bound policies for switching bandit problems. In *International Conference on Algorithmic Learning Theory*, pp. 174–188. Springer, 2011.
- Hartland, C., Gelly, S., Baskiotis, N., Teytaud, O., and Sebag, M. Multi-armed Bandit, Dynamic Environments and Meta-Bandits. working paper or preprint, November 2006. URL <https://hal.archives-ouvertes.fr/hal-00113668>.
- Ibragimov, I. A. On the composition of unimodal distributions. *Theory of Probability & Its Applications*, 1(2): 255–260, 1956.
- Kleinberg, R. and Leighton, T. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pp. 594–605. IEEE, 2003.
- Krishna, V. *Auction theory*. Academic press, 2009.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms. preprint*, 2018.
- Mohri, M. and Medina, A. M. Learning algorithms for second-price auctions with reserve. *The Journal of Machine Learning Research*, 17(1):2632–2656, 2016.
- Morgenstern, J. H. and Roughgarden, T. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems*, pp. 136–144, 2015.
- Myerson, R. B. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- Ostrovsky, M. and Schwarz, M. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC ’11*, pp. 59–60, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0261-6.
- Paes Leme, R., Pal, M., and Vassilvitskii, S. A field guide to personalized reserve prices. In *Proceedings of the 25th International Conference on World Wide Web, WWW*

'16, pp. 1093–1102, Republic and Canton of Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.

Perchet, V. and Rigollet, P. The multi-armed bandit problem with covariates. *The Annals of Statistics*, 41(2):693–721, 04 2013.

Pollard, D. *Convergence of stochastic processes*. Springer Science & Business Media, 1984.

Roughgarden, T. and Wang, J. R. Minimizing Regret with Multiple Reserves. In *Proceedings of the 2016 ACM Conference on Economics and Computation - EC '16*, volume 9, pp. 601–616, 2016.

Rudolph, M. R., Ellis, J. G., and Blei, D. M. Objective variables for probabilistic revenue maximization in second-price auctions with reserve. In *Proceedings of the 25th International Conference on World Wide Web*, pp. 1113–1122. International World Wide Web Conferences Steering Committee, 2016.

Saumard, A. and Wellner, J. A. Log-concavity and strong log-concavity: A review. *Statistics Surveys*, 8:45–114, 2014.

Shen, W., Lahaie, S., and Paes Leme, R. Learning to clear the market. In *International Conference on Machine Learning*, pp. 5710–5718, 2019.