

A heuristic for estimating Nash equilibria in first-price auctions with correlated values

Benjamin Heymann¹, Panayotis Mertikopoulos^{1,1}

^a*Criteo AI Lab*

^b*Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, 38000, Grenoble, France*

Abstract

Our paper concerns the computation of Nash equilibria of first-price auctions with correlated values. While there exist several equilibrium computation methods for auctions with *independent* values, the correlation of the bidders' values introduces significant complications that render existing methods unsatisfactory in practice. Our contribution is a step towards filling this gap: inspired by the seminal fictitious play process of Brown and Robinson, we present a learning heuristic – that we call *fictitious bidding* (FB) – for estimating Bayes-Nash equilibria of first-price auctions with correlated values, and we assess the performance of this heuristic on several relevant examples.

1. Introduction

In single-item sealed bid first-price auctions (sometimes referred to as *pay-as-bid*), several potential buyers bid simultaneously for an item to be bought; the highest bidder then gets the item and pays their bid to the seller. A common approach to model a sealed bid first-price auction is to frame it as a game. More precisely, since the buyers' valuations are not necessarily known by the other participants, one usually supposes that each of those valuations are random variables sampled from a probability distribution. In this narrative, each buyer observes their own valuation – a *private* signal – before submitting their bid, so that the resulting situation belongs to the class of *Bayesian* games. A good share of the literature on auctions makes the additional structuring assumptions that the valuations are independently distributed.

Given the probability distributions of the buyers' values, an important research question is to determine the outcome of the auction. While several solution concepts exist in the game-theoretic literature, the notion of Bayes-Nash equilibrium is of primary theoretical and practical importance for the study of auctions, and the numerical estimation of Bayes-Nash equilibrium has been a vigorously researched question for quite some time, with several breakthroughs along the way. However, a major challenge that arises in practice is that (a) the above methods invariably rely on a first-order characterization of the solution and (b) they require the bidders' values to be independently distributed, which is a very strong limitation for real applications.

Our contributions

Our paper seeks to overcome this limitation via a learning approach whereby the bidders' interactions are stylized as a repeated game, and the auction's Bayes-Nash equilibrium emerges as the players' empirical bid distribution profile over time. This heuristic – both simple and natural – specifies the bidders' behavior and simulates the dynamics of the bidders' strategy profiles as the auction is repeated, and shares some features with the seminal fictitious play procedure introduced by Brown and Robinson as a solution method for zero-sum games. This is why we name it fictitious bidding (FB).

Given the immense difficulty of calculating Bayes-Nash equilibria in first-price auctions – a problem which is PP-complete, and may even be PSPACE-complete [?] – any theoretical guarantees of our method would necessarily concern smaller, more tractable classes of first-price auctions. To avoid narrowing down our focus in this way, we forego a theoretical analysis of the proposed heuristic, and we undertake instead a numerical evaluation campaign in a series of diverse, relevant examples. Also, given that our heuristic does not rely on the structure of the first-price auction, it is also tested successfully on other auction rules. In all the examples we tested, the method provides an equilibrium outcome in a few seconds or minutes, depending on the example, a feature which we consider particularly appealing for practical applications.

To the best of our knowledge, the proposed heuristic is the first approach in the literature that remains agnostic to any correlations or dependencies between the users' value distributions. Moreover, the adopted learning viewpoint offers the distinct advantage of being oblivious to the actual auction mechanism (and any first-order characterization of the auction's equilibrium), so the proposed heuristic can also be applied directly to other auction settings – such as mixtures of first and second-price auctions or other combinations thereof.

Related work

The literature on auctions is immense [?], and is impossible to survey in a short paper; we therefore discuss below only the most relevant works we are aware of concerning the numerical computation of Bayes-Nash equilibria in first-price auctions.

In his seminal paper [?], Vickrey derived a characterization of Nash equilibria when the bidders' values are independent and the bidders otherwise identical. Later, Plum [?] showed how to compute the Nash equilibrium of 2-bidder auctions for some special cases when the price is a combination of first and second price. A first general numerical method for computing the Nash equilibrium of a first-price auction with independent values appears in [?]. Theoretical analyses of the equilibrium structure are provided in particular in [? ? ?]. In order to study bidding rings, Bajari introduces several heuristics to compute the Nash equilibrium of first-price auctions [?]. Several other computation methods for first-price auctions with independent values have been proposed since then [? ? ? ? ?], and more recently, [?], that addresses the cases of discrete value distributions. These methods rely on a first-order characterization of the best

reply of the players to produce a system of ordinary differential equations that are then solved using various methods. One of the main difficulty lies in the numerical instability of the solution to the system.

Research on first-price auctions received renewed impetus in 2019 when Google switched its display advertising market place to first-price [? ?]. This led to a surge of interest in new topics such as computational complexity [?], numerical approximation [?] or more recently, the use of neural networks to compute the auction’s equilibrium [?].

We also mention that there is a very active track of research that, in the of wake of [?], aims at maximizing the seller’s revenue with pragmatism [? ? ?].

Our paper takes a complementary, online learning approach to the above: the bidders’ interactions are modeled as a repeated game that unfolds in discrete time, with each agent playing an approximate ”best response” to the observed empirical distribution of bids. This heuristic closely resembles – and was inspired by – the fictitious play process as pioneered concurrently by Brown [?] and Robinson [?]. This is one of the most widely studied procedures for learning in games, and it involves each player playing a best response to her beliefs about her opponents, given here by the the empirical frequency of past play.

The convergence of these beliefs to Nash equilibrium was first established in two-player zero-sum games by Robinson [?]. Subsequently, the method has been shown to converge in 2×2 games [?], general N -player potential games [?], symmetric games with an interior evolutionarily stable strategy [?], as well as certain classes of supermodular games [? ? ? ?]. Variants of fictitious play involving a certain degree of explicit exploration / randomness have also been considered in the literature: the most widely studied of these processes is that of *stochastic* – or *perturbed* – fictitious play, which was introduced by Fudenberg & Kreps [?], and which was shown by Hofbauer & Sandholm [?] to converge to an approximate Nash equilibrium – a quantal response equilibrium to be exact – in the same classes of games as fictitious play.¹

At the same time, the literature on the convergence of (stochastic) fictitious play should not be interpreted as suggestion that these processes converge to equilibrium in *all* games: notable examples include Jordan’s three-player matching pennies variant [?], as well as the counterexamples of Shapley [?] and Gaunersdorfer and Hofbauer [?]. In a similar vein, we should stress that we do not make any claims of global convergence to Bayes-Nash equilibrium in *all* first-price auctions. However, the series of numerical examples presented is sufficiently wide in scope and breadth to provide reasonable optimism for the use of the proposed heuristic in practice.

¹Stochastic fictitious play is related – *but not in any way equivalent* – to the class of no-regret learning policies known as ”follow the regularized leader” [?]; for a detailed discussion, we refer the reader to [? ?], and references therein.

2. Formulation of the game, and agent-based reformulation

The choice of mathematical formalism to model the bidders' behavior (their *strategies*) plays a crucial role in the analysis. For instance, when the support of the bidders' values is an interval, it is well known that, under mild assumptions, there exists a unique pure Nash equilibrium [? ? ?], so randomization is not needed for the analysis. By contrast, if the bidders' values belong to a discrete set, bid randomization is required to ensure the existence of a Nash equilibrium [?]. In both cases — discrete or continuous support — most of the literature captures the bidders' strategies via their cumulative bid distributions. This is natural because, in both cases, the bid cannot decrease when the value increases.

However, this formalism is not well-adapted to our algorithmic proposal. For this reason, after providing a quick refresher on the standard first-price auction model, we present an alternative, agent-based formulation which is much better suited to the discussion of our learning heuristic.

2.1. The standard approach to model an auction as a game

A standard way to model one-item auctions is as follows: The auction is framed as a Bayesian game with players (the auction participants) belonging to a finite set P . Each player $p \in P$ has access to a (random) private signal $v_p \in \mathbb{R}^+$ that corresponds to their value for the item being auctioned, and places a *bid* $b_p = \beta_p(v_p)$. The auction is then cleared² and the payoff of the p -th player in a first-price auction is given by $(v_p - b_p) \prod_{j \in P; j \neq p} \{b_p > b_j\}$.

Formally, a one-item auction is a tuple $(P, (I_p)_{p \in P}, F)$ such that

1. P is a finite set;
2. Each I_p is a compact subset of \mathbb{R}^+ ;
3. F is a distribution on $\prod_{p \in P} I_p$.

For such an auction, a strategy for player $p \in P$ is a measurable map $\beta_p: I_p \rightarrow \mathbb{R}^+$. We denote by Σ_p the set of strategies of player p so the expected payoff of player p given a strategy profile $(\beta_p)_{p \in P}$ is

$$\pi_p(\beta_p, \beta_{-p}) = \mathbb{E}_F [(v_p - \beta_p(v_p)) \prod_{j \neq p} \{\beta_p(v_p) > \beta_j(v_j)\}], \quad (1)$$

where we have used the standard game-theoretic shorthand $-p = P \setminus \{p\}$. A (pure strategy) Nash Equilibrium (NE) is then defined as a strategy profile β_P such that

$$\pi_p(\beta_p, \beta_{-p}) \geq \pi_p(\beta'_p, \beta_{-p}) \quad \forall \beta_p \in \Sigma_p, \quad \forall p \in P. \quad (2)$$

In the remaining of this paper, we will refer to this formulation as the *player-based* formulation.

²In case of equality of the highest bids, a tie breaking rule is needed.

As we mentioned above, when the value sets $(I_p)_{p \in P}$ are intervals, there exists under mild assumptions a (unique) Nash equilibrium in pure strategies, so there is no need to introduce mixed strategies. However, such mixed strategies are required in a discrete setting where the existence of a Nash equilibrium is only guaranteed in *mixed* strategies. We will therefore introduce mixed strategies only for the discrete setting because, as observed by Aumann [?], the definition of mixed strategies when pure strategies are continuous mappings between two continuous sets introduces significant technical difficulties that are beyond the scope of this work.³ We hence call a mixed strategy a probability distribution on Σ_p , and denote by $\Delta\Sigma_p$ the set of mixed strategies of player p .

2.2. The agent-based auction model

We propose a variant to the *player-based* auction model introduced in Section ?? . An agent-based auction model is a tuple $(\mathbb{A}, v_{\mathbb{A}}, \mu)$ such that

1. \mathbb{A} is a finite set of agents;
2. $v_{\mathbb{A}}$ is a vector of values (assigning a single value per agent).
3. μ is a probability distribution on $2^{\mathbb{A}}$.

The strategy set Γ is defined to be the set of probability distributions on \mathbb{R}^+ , so a strategy profile $(\gamma_a)_{a \in \mathbb{A}} \in \Gamma^{\mathbb{A}}$, denoted $\gamma_{\mathbb{A}}$ induces a probability distribution on $(\mathbb{R}^+)^{\mathbb{A}}$. An ϵ -Nash equilibrium in this setting is a strategy profile $\gamma_{\mathbb{A}}^*$ such that for all $a \in \mathbb{A}$ and $\gamma_{\mathbb{A}} \in \Gamma^{\mathbb{A}}$

$$\pi_a(\gamma_a^*, \gamma_{-a}^*) \geq \pi_a(\gamma_a, \gamma_{-a}^*) - \epsilon, \quad (3)$$

where, for a first-price auction,

$$\pi_a(\gamma_a, \gamma_{-a}) = \int (a - b_a) \prod_{a' \in S \setminus a} \{b_a > b_{a'}\} d\mu(S|a \in S) d\gamma_{\mathbb{A}}(b_{\mathbb{A}}). \quad (4)$$

2.3. Equivalence of the two formulations

The relation between the two formulation might not be a priori obvious. In view of this, we show below that the agent-based formulation is equivalent to the player-based one when the underlying set of values is discrete.

Lemma 1. *Let (P, I_P, F) be a player-based one-item auction formulation. Suppose that I_p is discrete for all $p \in P$. Then there exist*

1. An agent-based auction model $(\mathbb{A}, v_{\mathbb{A}}, \mu)$
2. A partition $(\mathbb{A}_p)_{p \in P}$ of \mathbb{A}
3. A family of application $(\phi_p)_{p \in P}$, where ϕ_p is a bijection from the player's mixed strategy set $\Delta\Sigma_p$ to its agents' strategy set $\Gamma^{\mathbb{A}_p}$.

³The problem is that one cannot write in this case $\Delta\Sigma_p$.

such that for any player strategy profile $\beta_P \in \prod_{p \in P} \Sigma_p$, we have

$$\pi_p(\beta_p) = \sum_{a \in \mathbb{A}_p} \pi_a((\phi_p(\beta_p))_{p \in P}) \mu(S \ni a) \quad (5)$$

Proof. Set $\mathbb{A}_p = [|I_p|]$, $v_{\mathbb{A}_p} = I_p$ for all $p \in P$, and $\forall S \in 2^{\mathbb{A}}$, $\mu(S) = 0$ if $S \notin \times_{p \in P} \mathbb{A}_p$ and $\mu(S) = F(v_S)$ otherwise. Then ϕ_p is obtained by desintegrating the players' mixed strategies over the set of agents. \square

3. Fictitious bidding: definition, implementation, and results

3.1. Definition of the heuristic

As we discussed in the introduced, the proposed heuristic for calculating the Nash equilibrium of an auction is inspired by the process of fictitious play for learning in finite games [? ?]. The key novelty of our approach — which leads to an easily implementable heuristic — is that the process unfolds in the *agent-based* representation of the game, not its normal form (or a discretized version thereof).

To describe the heuristic formally, let \mathbb{B} be a discrete grid of bid values, let $\eta_k > 0$ be a "learning rate" sequence, and let $(\gamma_a^0(b))_{a \in [N_a], b \in \mathbb{B}}$ be an initial strategy profile. Then the proposed heuristic proceeds by taking at each stage $k = 0, 1, \dots$, a best reply to a generalized exponential moving average of the players' empirical bid distributions up to stage k . More precisely, we have the following sequence of events:

1. Strategies are initialized with $(\gamma_a^0(b))_{a \in [N_a], b \in \mathbb{B}}$
2. At each time step k , every agent $a \in \mathbb{A}$ performs the bid update

$$b_a^{(k)} \in \arg \max_{b \in \mathbb{B}} \pi_a(\delta_b, \gamma_{-a}^{(k)}) \quad (6)$$

3. Subsequently, mixed strategies are updated as

$$\gamma_a^{(k+1)} = (1 - \eta_k) \gamma_a^{(k)} + \eta_k \delta_{b_a^{(k)}} \quad (7)$$

4. Repeat until a termination criterion is triggered

At a high level, this process is conceptually similar to fictitious play as, at each stage, the algorithm outputs a best response to "some" average of past play. This analogy would be precise if the averaging assigned equal weight to all past states, which in turn corresponds to the choice $\eta_k = 1/k$ for the method's learning rate. However, in many cases, it is more advantageous to assign exponentially more weight to recent observations relative to past ones, as this would effectively discount the irrelevant impact of the method's (arbitrary) initialization; this corresponds to the choice $\eta_k = \text{const.}$, in which case the averaging undertaken is a straightforward exponential moving average.

3.2. Implementation and results

In terms of implementation, the fictitious bidding heuristic can be coded in less than 100 lines in Julia [?] and is available online.⁴ All our experiments were run with this implementation on a personal laptop. For the case of independent probability distributions we can use the simplifying observation that $\gamma_{-a}^{(k)}$ is the same for every agent of the same player, which in turn alleviates the computational and memory burden. In the plots we provide, we report the cumulative distribution of the agent’s bid as well as their payoff, which allows for checking visually that the support of the bids is approximately in the maximand of the payoff.

Example 1. We take $|\mathbb{A}| = 4$, $v_{\mathbb{A}} = (0, 0, 1., 1.)$, and μ take the value 0.25 on the following 4 scenarios: (a_1, a_2) , (a_3, a_4) , (a_1, a_4) , (a_2, a_3) . This corresponds to a symmetric setting with two bidders whose values are sampled in $\{0, 1\}$ uniformly and independently. This example is interesting because it posses an analytic solution. Indeed, suppose that in the equilibrium, no agent bids above its value, then the bid of agent a_1 and agent a_2 should be zero. Since the setting is symmetric, let us denote by G the cumulative of the bid of agent a_3 and a_4 . Observe that $G(b) = 1/(1-b) - 1$ for $b \in [0, 1/2]$ corresponds to a Nash equilibrium because (1) the expected payoff of agent a_3 (resp. a_4) is constant, equal to 0.5, for $b \in [0, 1/2]$ and strictly smaller than 0.5 for any $b > 1/2$. We took $\mathbb{B} = [0, 1/400, 2/400, \dots 1]$, $\eta_k = 0.01/k$ and performed 10^5 iterations. The strategy estimate we obtain is 8×10^{-5} -Nash equilibrium. The result is displayed in Figure ??.




Figure 1: The bid cumulative of agent 3 (of value 1) in Example ??

Example 2. Next we take $|\mathbb{A}| = 3$ and $\mathbb{A} = (1/3, 2/3, 3/3)$ with the two equiprobable scenarios (a_1, a_2) and (a_2, a_3) . We took $\mathbb{B} = [0, 1/600, 2/600, \dots 1]$,

⁴we will release the code with the definitive version of the paper




Figure 2: Results for Example ??




Figure 3: Example ??

$\eta_k = 0.01/k$ and performed 10^6 iterations. The strategy profile we obtain is a (1.5×10^{-4}) -Nash equilibrium. Results are displayed in Figure ??.

Example 3. Let $v_{\mathbb{A}} = (1/4, 2/4, 2/4, 4/4)$ with the four equiprobable scenarii $(a_1, a_2), (a_2, a_3), (a_1, a_3), (a_1, a_2, a_3, a_4)$. We took $\mathbb{B} = [0, 1/400, 2/400, \dots, 1]$, $\eta_k = 0.01/k$ and performed 10^5 iterations. The strategy estimate we obtain is a (2.5×10^{-3}) -Nash equilibrium. The results are displayed in Figure ?? .

Example 4 (From Wang et al.). We test the heuristic on Example 8 from [?], which was used to test the stability of continuous method on discrete setting. It is notable that we do better than the result reported by [?] for the continuous methods. We use the library developed in [?] to compare our heuristic with their solution. The setting is as follow. There are 3 players, each represented by 3 agents, so there are 9 agents in total. The potential values of all players belong to $\{0.1, 0.2, 0.25\}$, and are sampled independently. What differs between the players is the relative probability of those three values:

- For the first player those probabilities are 1/4, 1/4, and 1/2
- For the second and third player, they are 0.05, 0.45, and 0.5.

The results are reported in Figure ??, and the strategy profile we obtain is a (4×10^{-5}) -Nash equilibrium of the auction.

Example 5 (Another comparison with Wang et al.). We use the library developed in [?] to compare fictitious bidding with another example. The setting




Figure 4: Result for Example ??




Figure 5: Result for Example ??

is as follow. There are 2 players, 2 agents for player 1, 3 for player 2. . The potential values of player 1 are 0.1 and 0.25 with respective probability 0.25 and 0.75, and the potential values of player 2 are 0.1, 0.2 and 0.25 with respective probability 0.05, 0.45 and 0.5. The values are sampled independently. The results are reported in Figure ?? and the strategy profile we obtain is a (9×10^{-4}) -Nash equilibrium.

Example 6. In order to test the robustness of the fictitious bidding, we generated random instances of auction games. Each instance contains 10 agents whose values are samples uniformly between 0 and 1. We then generate 20 scenarios by taking ten 20 times at random a pair in $\{1 \dots 10\}$. We took $\mathbb{B} = [0, 1/100, 2/100, \dots, 1]$, $\eta_k = 0.01/k$ and performed 10^6 iterations. The strategy profile we obtain is ϵ -Nash equilibrium, with ϵ respectively equals to 0.00048, 0.00083, 0.00139, 0.00058, 0.00402, 0.01374, 0.00501, 0.00305, 0.00326 and 0.00472.

As a side note, we observed during our experiments that the fictitious bidding heuristic seems to work well for other auctions, such as mixture of first and second-price (The item is still allocated to the highest bidder, but the winner now pays the average of the first and second highest bids).

4. Conclusion

We have introduced a fictitious-play like heuristic to compute the Bayes-Nash equilibrium of first-price auctions even when the values are correlated: fictitious bidding. Compared to other methods in the literature, the proposed fictitious bidding is inherently discrete, as it is introduced for auctions with values and bids living in a discrete set. As a result, our heuristic does not rely on the traditional continuous approximation of the auction, which can be viewed in itself as an improvement, since applications usually take place in discrete environments. The heuristic displayed also good result for mixtures of first and second-price auctions. Overall, our experimental results suggest there is a strong convergence theorem underneath; clarifying under which conditions this is true (and why) is a central question for future research on the topic.

References

- [1] Y. Cai, C. H. Papadimitriou, Simultaneous Bayesian auctions and computational complexity, in: EC' 14: Proceedings of the 15th ACM Conference on Electronic Commerce, 2014.
- [2] V. Krishna, Auction theory, Academic press, 2009.
- [3] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, *The Journal of finance* 16 (1) (1961) 8–37.
- [4] M. Plum, Characterization and computation of nash-equilibria for auctions with incomplete information, *International Journal of Game Theory* 20 (4) (1992) 393–418.
- [5] R. C. Marshall, M. J. Meurer, J.-F. Richard, W. Stromquist, Numerical analysis of asymmetric first price auctions, *Games and Economic Behavior* 7 (2) (1994) 193–220.
- [6] B. Lebrun, First price auctions in the asymmetric n bidder case, *International Economic Review* 40 (1) (1999) 125–142.
URL <http://www.jstor.org/stable/2648842>
- [7] E. Maskin, J. Riley, Uniqueness of equilibrium in sealed high-bid auctions, *Games and Economic Behavior* 45 (2) (2003) 395–409.
- [8] P. J. Reny, S. Zamir, On the existence of pure strategy monotone equilibria in asymmetric first-price auctions, *Econometrica* 72 (4) (2004) 1105–1125.
- [9] P. Bajari, Comparing competition and collusion: a numerical approach, *Economic Theory* 18 (1) (2001) 187–205.
- [10] W. R. Gayle, J. F. Richard, Numerical solutions of asymmetric, first-price, independent private values auctions, *Computational Economics* 32 (2008) 245–278. doi:10.1007/s10614-008-9125-7.
- [11] G. Fibich, N. Gavish, Numerical simulations of asymmetric first-price auctions, *Games and Economic Behavior* 73 (2) (2011) 479–495.
- [12] T. R. Kaplan, S. Zamir, Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case, *Economic Theory* 50 (2) (2012) 269–302.
- [13] G. Fibich, N. Gavish, Asymmetric first-price auctions—a dynamical-systems approach, *Mathematics of Operations Research* 37 (2) (2012) 219–243. arXiv:<https://doi.org/10.1287/moor.1110.0535>, doi:10.1287/moor.1110.0535.
URL <https://doi.org/10.1287/moor.1110.0535>
- [14] T. P. Hubbard, H. J. Paarsch, On the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm, in: *Handbook of computational economics*, Vol. 3, Elsevier, 2014, pp. 37–115.

- [15] Z. Wang, W. Shen, S. Zuo, Bayesian nash equilibrium in first-price auction with discrete value distributions, in: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '20, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2020, p. 1458–1466.
- [16] B. Heymann, How to bid in unified second-price auctions when requests are duplicated, *Operations Research Letters* 48 (4) (2020) 446–451.
- [17] R. Paes Leme, B. Sivan, Y. Teng, Why do competitive markets converge to first-price auctions?, in: Proceedings of The Web Conference 2020, 2020, pp. 596–605.
- [18] A. Filos-Ratsikas, Y. Giannakopoulos, A. Hollender, P. Lazos, D. Poças, On the complexity of equilibrium computation in first-price auctions, arXiv preprint arXiv:2103.03238.
- [19] I. Rasooly, C. Gavidia-Calderon, The importance of being discrete: on the inaccuracy of continuous approximations in auction theory (2021). [arXiv:2006.03016](https://arxiv.org/abs/2006.03016).
- [20] M. Bichler, M. Fichtl, S. Heiderkrueger, N. Kohring, P. Sutterer, Learning equilibria in symmetric auction games using artificial neural networks, *Nature Machine Intelligence*.
- [21] R. B. Myerson, Optimal auction design, *Mathematics of operations research* 6 (1) (1981) 58–73.
- [22] J. D. Hartline, T. Roughgarden, Simple versus optimal mechanisms, in: Proceedings of the 10th ACM conference on Electronic commerce, 2009, pp. 225–234.
- [23] P. Dhangwatnotai, T. Roughgarden, Q. Yan, Revenue maximization with a single sample, *Games and Economic Behavior* 91 (2015) 318–333.
- [24] R. Cole, T. Roughgarden, The sample complexity of revenue maximization, in: Proceedings of the forty-sixth annual ACM symposium on Theory of computing, 2014, pp. 243–252.
- [25] G. W. Brown, Iterative solutions of games by fictitious play, in: T. C. Coopmans (Ed.), *Activity Analysis of Productions and Allocation*, 374–376, Wiley, 1951.
- [26] J. Robinson, An iterative method for solving a game, *Annals of Mathematics* 54 (1951) 296–301.
- [27] K. Miyasawa, On the convergence of learning processes in a 2×2 non-zero sum game, *Research memorandum* 33, Princeton University (1961).
- [28] D. Monderer, L. S. Shapley, Fictitious play property for games with identical interests, *Journal of Economic Theory* 68 (1996) 256–265.
- [29] J. Hofbauer, Stability for the best response dynamics, mimeo (1995).
- [30] P. Milgrom, J. Roberts, Rationalizability, learning, and equilibrium in games with strategic complementarities, *Econometrica* 58 (6) (1990) 1255–1277.
- [31] P. Milgrom, J. Roberts, Adaptive and sophisticated learning in normal form games, *Games and Economic Behavior* 3 (1991) 82–100.
- [32] V. Krishna, Learning in games with strategic complementarities, mimeo (1992).
- [33] S. Hahn, The convergence of fictitious play in games with strategic complementarities, *Economics Letters* 99 (304–306).
- [34] D. Fudenberg, D. M. Kreps, Learning mixed equilibria, *Games and Economic Behavior* 5 (320–367).
- [35] J. Hofbauer, W. H. Sandholm, Stable games and their dynamics, *Journal of Economic Theory* 144 (4) (2009) 1665–1693.
- [36] S. Shalev-Shwartz, Online learning and online convex optimization, *Foundations and Trends in Machine Learning* 4 (2) (2011) 107–194.
- [37] P. Mertikopoulos, W. H. Sandholm, Learning in games via reinforcement and regularization, *Mathematics of Operations Research* 41 (4) (2016) 1297–1324.
- [38] S. Hadikhani, R. Laraki, P. Mertikopoulos, S. Sorin, Learning in nonatomic games, Part I: Finite action spaces and population games, <https://arxiv.org/abs/2107.01595> (2021).

- [39] J. S. Jordan, Three problems in learning mixed strategy Nash equilibria, *Games and Economic Behavior* 5 (3) (1993) 368–386.
- [40] L. S. Shapley, Some topics in two-person games, in: *Advances in Game Theory*, no. 52 in *Annals of Mathematics Studies*, Princeton University Press, 1964.
- [41] A. Gaunersdorfer, J. Hofbauer, Fictitious play, Shapley polygons, and the replicator equation, *Games and Economic Behavior* 11 (2) (1995) 279–303.
- [42] B. Lebrun, First price auctions in the asymmetric n bidder case, *International Economic Review* 40 (1) (1999) 125–142.
- [43] S. Athey, Single crossing properties and the existence of pure strategy equilibria in games of incomplete information, *Econometrica* 69 (4) (2001) 861–889.
- [44] Z. Wang, W. Shen, S. Zuo, Bayesian nash equilibrium in first-price auction with discrete value distributions, in: *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 1458–1466.
- [45] R. Aumann, Mixed and behavior strategies in infinite extensive games, *advances in game theory*, *Annals of Mathematics Studies* 627–650.
- [46] J. Bezanson, A. Edelman, S. Karpinski, V. B. Shah, Julia: A fresh approach to numerical computing, *SIAM review* 59 (1) (2017) 65–98.
URL <https://doi.org/10.1137/141000671>