An empirical study of Fictitious Play for estimating Nash equilibria in first-price auctions with correlated values

Benjamin Heymann B.HEYMANN@CRITEO.COM
Criteo AI Lab, Fairplay joint team Paris, France

Panayotis Mertikopoulos PANAYOTIS.MERTIKOPOULOS@IMAG.FR
CNRS, Laboratoire d'Informatique de Grenoble (LIG)

Abstract

This study concerns the computation of the *Nash equilibria* of *first-price auctions* with *correlated values*. Although some equilibrium computation methods exist for auctions with *independent* values, the correlation of bidders' values introduces significant complications that render the existing methods unsatisfactory. Our empirical contribution is a step towards filling this gap. We report surprisingly good numerical convergence of *Fictitious Play* toward an ϵ-equilibrium for an extensive set of instances. By doing so, we extend the insights of [39] to the correlated setting. These preliminary results call for further investigations into the properties of fictitious play algorithms on first-price auctions.

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. The highest bidder then obtains the item and pays the corresponding bid to the seller. A common approach to auctions is to frame them as games. More precisely, because buyers' values are not necessarily known to other participants, it is usually assumed that each of these values is a random variable sampled from a probability distribution. In this narrative, each buyer is assumed to have formed their individual value – a *private* signal – before submitting their bid, so that the resulting strategic interactions can be recast in the form of a *Bayesian game*. A sizable part of the auction-theoretic literature makes an additional structural assumption that agents' values are independently distributed (an assumption we do not make).

Given the probability distributions of 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 [1] is of primary theoretical and practical importance for the study of auctions, and equilibrium numerical estimation 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 stringent limitation for real-life applications of auction theory (*e.g.*, in online ad auctions).

In parallel, Fictitious Play (FP) was first discussed in [5]. It has attracted considerable attention because (a) it is a simple procedure that has been used to justify the Nash equilibrium play [4], and (b) despite being slow, it has been used in AI to approximate Nash equilibrium [12, 39].

1. since the context is clear, we will use the term Nash equilibrium, or just equilibrium in this article

© B. Heymann & P. Mertikopoulos.

Contributions

In this work, we broaden the empirical results from [39] to accommodate auctions with correlated values, employing several technical adaptations (see Section 5). In all the examples we tested, the method approximated the equilibrium surprisingly well. Notably, the FP could be the first approach in the literature that remains agnostic to any correlations or dependencies between bidders' value distributions.

2. Numerical solutions for auctions

The literature on auctions is immense [26, 43] and is impossible to survey in a short paper. Therefore, we discuss below only the most relevant works we are aware of concerning the numerical computation of Nash equilibria in first-price auctions.

In his seminal paper [46], Vickrey derived a characterization of equilibria when bidders' values are independent and otherwise identical. Later, Plum [38] showed how to compute the equilibrium of 2-bidder auctions for some special cases when the price is a combination of the first and second prices. The first general numerical method for computing the equilibrium of a first-price auction with independent values appears in [28]. Theoretical analyses of the equilibrium structure are provided in particular in [27, 29, 41]. To study bidding rings, Bajari introduced several heuristics to compute the equilibrium of first-price auctions [1]. Several other computation methods for first-price auctions with independent values have been proposed since then [8, 9, 14, 22, 24], and more recently, [47], which addresses the cases of discrete value distributions. These methods rely on a first-order characterization of the players best replies to produce a system of ordinary differential equations, which are then solved using various methods. One of the main difficulties lies in the numerical instability of the solutions.

Research on first-price auctions received renewed impetus in 2019, when Google switched its display advertising marketplace to first-price auctions [18, 37]. This led to a surge of interest in new topics such as computational complexity [10], numerical approximation [40] or, more recently, neural networks to compute auction equilibrium [3]. We also mention that there is an active track of research that, in the wake of [36], aims to maximize the seller's revenue [6, 7, 17].

3. Fictitious play

The work that is closest in spirit to our own is mostly empirical work [39], where the authors rely on FP running on a discrete set of possible actions. That said, the setting of [39] still differs from our own in that bidders are assumed therein to be symmetric in most of the paper, and their values are further assumed to be continuously and *independently* distributed. In contrast, we consider both correlations between bidders (a crucial extension to capture real-life market behaviors) and atoms in the value distributions.

The original fictitious play process was pioneered concurrently by Brown [5] and Robinson [42]. This is one of the most widely studied procedures for learning in games, and it involves each player playing a best response to his/her beliefs about his/her opponents, given here by the empirical frequency of past play. Robinson first established the convergence of these beliefs to equilibrium in two-player zero-sum games [42]. Subsequently, the method has been shown to converge in 2×2 games [34], general N-player potential games [35], symmetric games with an interior evolutionarily stable strategy [20], and certain classes of supermodular games [16, 25, 31, 32]. Variants of FP 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* – FP, which was introduced by Fudenberg & Kreps [11] and shown by Hofbauer & Sandholm [21] to converge to an approximate equilibrium – a quantal response equilibrium to be exact – in the same classes of games as FP.[2]

At the same time, the literature on the convergence of (stochastic) FP should not be interpreted as suggesting that these processes converge to equilibrium in *all* games. Notable examples include Jordan's three-player matching pennies variant [23], as well as the counterexamples by Shapley [45] and Gaunersdorfer and Hofbauer [13]. Similarly, we should stress that we do not make any claims of global convergence to 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 practical use of FP, and call for further theoretical investigations.

4. Auction model

An auction is a Bayesian game where the players are the auction participants. Each player places a bid after seeing their random and private value. Then, the auction is cleared according to the first-price auction rules.[3]. When the values and the bids belong to a finite set, it is practical to take a perspective based on agents, by associating each pair (player, value) to an agent. Hence, we can regard the auction as a tuple $(\mathbb{A}, \mathcal{P}, v, \mu, \mathbb{B})$ such that $\mathbb{A}$ is a finite set of agents –one per (player, value) pair–, $\mathcal{P}$ maps back each agent from $\mathbb{A}$ to the player it represents, v maps each agent from $\mathbb{A}$ to a value, μ is a probability distribution on $2^{\mathbb{A}}$ that encodes the common knowledge Bayesian prior, and $\mathbb{B}$ is a grid of admissible bids. A policy γ_a for an agent a is then a probability distribution on the set of bids $\mathbb{B}$.

Let $S_a = \{S \in 2^{\mathbb{A}} \mid a \in S\}$ be the set of all possible combinations of agents that contains agent a. For a first-price auction, the agent's (expected) payoff writes

$$\pi_a(\gamma_a, \gamma_{-a}) = \sum_{S \in S_a} \sum_{b \in \mathbb{B}^{\mathbb{A}}} \hat{\pi}_a(b_a, b_{-a}) \mu(S) \gamma_{-a}(b_{-a}) \gamma_a(b_a), \tag{1}$$

where $\hat{\pi}_a(b_a, b_{-a})$ is the payoff of agent a when their bid is b_a and the other agents bids are concatenated in the vector b_{-a}. In the absence of ties, this is $\hat{\pi}_a(b_a, b_{-a}) = (v_a - b_a) \prod_{a' \in S \setminus a} \{b_a > b_{a'}\}$. A tie-breaking rule needs to be specified for the case of multiple highest bidders. It is typically the lexicographic order or a random allocation among the highest bidders. It seems empirically that giving a payoff of 0 to everybody in case of ties improves the convergence of FP. While this is a choice that is hard to justify for an application, the impact on the game ϵ-equilibrium[4] is arguably small in many cases, so that this tie-breaking rule can be used to approximate the real game with FP.

2. Stochastic FP is related – *but not in any way equivalent* – to the class of no-regret learning policies known as "follow the regularized leader" [44]; for a detailed discussion, we refer the reader to [15, 30], and references therein.
3. In case of equality of the highest bids, a tie-breaking rule is needed.
4. An ϵ-equilibrium in this setting is a strategy profile $\gamma^\star$ such that for all $a \in \mathbb{A}$ and all alternative strategy profile γ, $\pi_a(\gamma_a^\star, \gamma_{-a}^\star) \geq \pi_a(\gamma_a, \gamma_{-a}^\star) - \epsilon$.

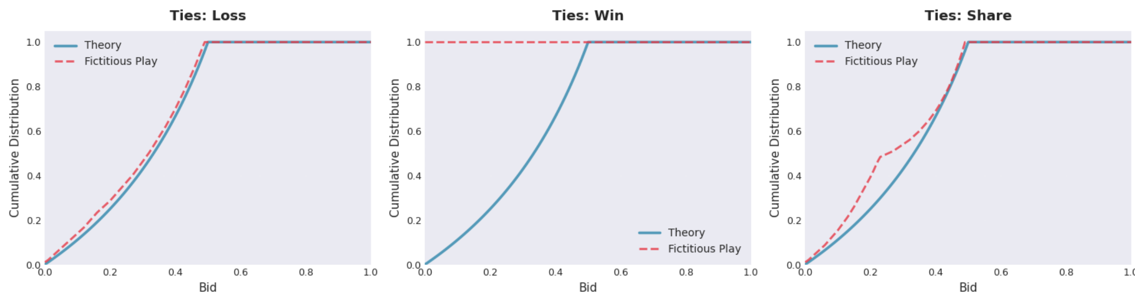

Figure 1: We consider a two-player setting where values are sampled independently and uniformly from $\{0, 1\}$. The blue curve corresponds to the theoretical solution when bids can be in $[0, 1]$ (continuous case). The tie-breaking rule we propose — no player receives the item when bids are equal — results in a reasonable approximation of the equilibrium (left). In the middle, we tested giving an item to each highest bidder, and on the right, to split the item in half.

5. Experiments

FP can be defined as follows for any stage $k > 0$

$$b_a^{(k)} \in \arg\max_{b \in \mathbb{B}} \pi_a(\delta_b, \gamma_{-a}^{(k)}), \quad \forall a \in \mathbb{A}, \tag{2}$$

$$\gamma_a^{(k+1)} = (1 - \eta_k)\gamma_a^{(k)} + \eta_k \delta_{b_a^{(k)}}, \forall a \in \mathbb{A}, \tag{3}$$

where $\eta_k > 0$ is a sequence of learning rates, and $(\gamma_a^0)_{a \in \mathbb{A}}$ an initial strategy profile, and δ_b is a mass of 1 at b. We take $\eta_k = {}^1\!/_{k+1}$ in most of the experiments. The algorithm and code for the experiments were implemented in Julia [2] and can be found on this Github repository `https://github.com/BenHey/FP4FPA`. Other experiments are reported in a earlier version of the paper [19].

We recommend a tie-breaking rule, in which no player receives the item when bids are equal. Figure 1 provides an ablation study illustrating this approach. We consider a two-player setting where values are sampled independently and uniformly from $\{0, 1\}$. When bids are constrained to $[0, 1]$, the equilibrium bidding policy can be derived analytically: players with value zero should bid 0, while players with value 1 should randomize their bids according to the cumulative distribution function $F(b) = \min(1/(1 - b) - 1, 1)$. We compare three tie-breaking rules. Only the one we propose, on the left, results in a convergence close to the continuous, theoretical equilibrium.

We tested FP on a symmetric continuous environment, for which there exists a method to compute a numerical solution [33]. The results are shown in Figure 2. Then, we tested FP on a batch of 20 randomly generated discrete environments with correlated values. The results are shown in Figure 3. For the two experiments, FP converges to ϵ-equilibria with a small ϵ. Finally, Figure 4 presents results on several auction instances, each characterized by a list of VALUES (one per agent) and a list of SCENARIOS. Since we allow repetition of scenarios, their elements can be treated as equiprobable without loss of generality. To visually assess convergence quality, we overlay each agent's (renormalized) policy density with their corresponding payoffs. At Nash equilibrium, an agent's payoff should be maximal on the support of their policy, making any deviations immediately apparent in these plots.

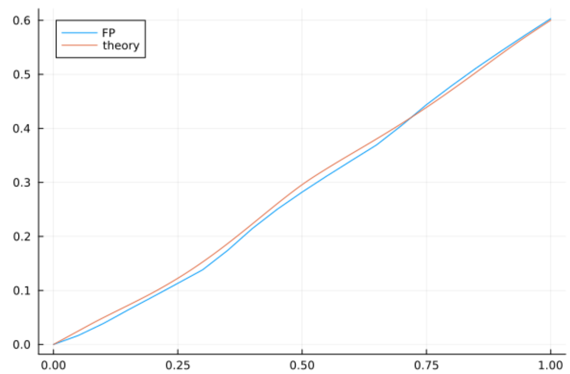

Figure 2: We consider an auction with two bidders, whose correlated values are distributed in $[0, 1]$. The generative model first tosses a coin c valued in $\{L, H\}$. Then, if $c = L$, the values are uniformly and independently sampled in $[0, 2/3]$. Otherwise, if $c = H$, the values are sampled uniformly and independently in $[1/3, 1]$. We can approximate the theoretical solution using numerical integration [33]. We obtain $\epsilon = 1.2 \cdot 10^{-5}$

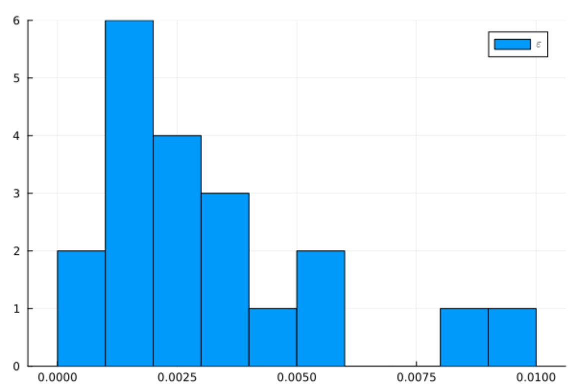

Figure 3: Distribution of ϵ for the 20 experiments on random environments.

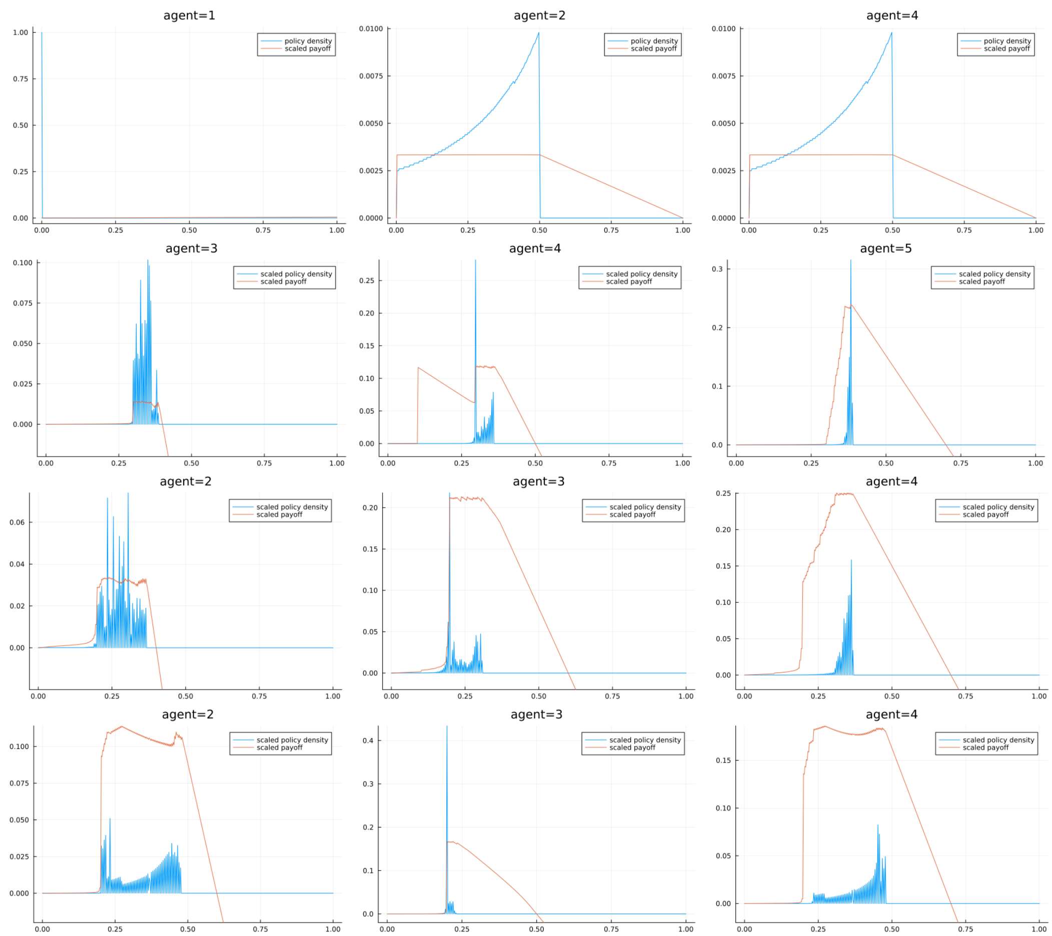

Figure 4: Visual diagnostic of Nash equilibrium convergence. For each agent and auction instance, we superimpose the normalized policy density with the corresponding payoff function. The equilibrium condition—that payoffs are maximal on the policy support—can be verified visually. The four rows correspond to different auction configurations: (Row 1) 4 agents with values $[0.0, 1.0, 0.0, 1.0]$ and scenarios $[1, 3]$, $[1, 4]$, $[2, 3]$, $[2, 4]$; (Row 2) 5 agents with values $[0.2, 0.3, 0.4, 0.5, 0.7]$ and scenarios $[1, 4]$, $[2, 4]$, $[3, 4]$, $[3, 5]$, $[3, 5]$; (Row 3) 4 agents with values $[0.2, 0.4, 0.6, 0.7]$ and scenarios $[1, 3]$, $[1, 3]$, $[1, 4]$, $[2, 3]$, $[2, 4]$, $[2, 4]$; (Row 4) 4 agents with values $[0.2, 0.6, 0.5, 0.7]$ and scenarios $[1, 3]$, $[1, 3]$, $[1, 4]$, $[2, 3]$, $[2, 4]$, $[2, 4]$.

References

[1] Patrick Bajari. Comparing competition and collusion: a numerical approach. *Economic Theory*, 18(1):187–205, 2001.

[2] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. URL https://doi.org/10.1137/141000671.

[3] Martin Bichler, Maximilian Fichtl, S Heiderkrueger, Nils Kohring, and Paul Sutterer. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Machine Intelligence*, 2021.

[4] Felix Brandt, Felix Fischer, and Paul Harrenstein. On the Rate of Convergence of Fictitious Play.

[5] George W. Brown. Iterative solutions of games by fictitious play. In T. C. Coopmans, editor, *Activity Analysis of Productions and Allocation*, 374-376. Wiley, 1951.

[6] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 243–252, 2014.

[7] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[8] Gadi Fibich and Nir Gavish. Numerical simulations of asymmetric first-price auctions. *Games and Economic Behavior*, 73(2):479–495, 2011.

[9] Gadi Fibich and Nir Gavish. Asymmetric first-price auctions—a dynamical-systems approach. *Mathematics of Operations Research*, 37(2):219–243, 2012. doi: 10.1287/moor.1110.0535. URL https://doi.org/10.1287/moor.1110.0535.

[10] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. *arXiv preprint arXiv:2103.03238*, 2021.

[11] Drew Fudenberg and David M. Kreps. Learning mixed equilibria. *Games and Economic Behavior*, 5(320-367), 1993.

[12] Sam Ganzfried and Tuomas Sandholm. Computing an approximate jam/fold equilibrium for 3-player no-limit texas hold'em tournaments. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 919–925, 2008.

[13] Andrea Gaunersdorfer and Josef Hofbauer. Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11(2):279–303, 1995.

[14] Wayne Roy Gayle and Jean Francois Richard. Numerical solutions of asymmetric, first-price, independent private values auctions. *Computational Economics*, 32:245–278, 2008. ISSN 09277099. doi: 10.1007/s10614-008-9125-7.

[15] Saeed Hadikhanloo, Rida Laraki, Panayotis Mertikopoulos, and Sylvain Sorin. Learning in nonatomic games, Part I: Finite action spaces and population games. https://arxiv.org/abs/2107.01595, 2021.

[16] Sunku Hahn. The convergence of fictitious play in games with strategic complementarities. *Economics Letters*, 99(304-306), 2008.

[17] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.

[18] Benjamin Heymann. How to bid in unified second-price auctions when requests are duplicated. *Operations Research Letters*, 48(4):446–451, 2020.

[19] Benjamin Heymann and Panayotis Mertikopoulos. A heuristic for estimating nash equilibria in first-price auctions with correlated values. *arXiv preprint arXiv:2108.04506*, 2021.

[20] Josef Hofbauer. Stability for the best response dynamics. mimeo, 1995.

[21] Josef Hofbauer and William H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693, July 2009.

[22] Timothy P Hubbard and Harry J Paarsch. On the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm. In *Handbook of computational economics*, volume 3, pages 37–115. Elsevier, 2014.

[23] James S. Jordan. Three problems in learning mixed strategy Nash equilibria. *Games and Economic Behavior*, 5(3):368–386, July 1993.

[24] Todd R Kaplan and Shmuel Zamir. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

[25] Vijay Krishna. Learning in games with strategic complementarities. mimeo, 1992.

[26] Vijay Krishna. *Auction theory*. Academic press, 2009.

[27] Bernard Lebrun. First price auctions in the asymmetric n bidder case. *International Economic Review*, 40(1):125–142, 1999. ISSN 00206598, 14682354. URL http://www.jstor.org/stable/2648842.

[28] Robert C Marshall, Michael J Meurer, Jean-Francois Richard, and Walter Stromquist. Numerical analysis of asymmetric first price auctions. *Games and Economic Behavior*, 7(2):193–220, 1994.

[29] Eric Maskin and John Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395–409, 2003.

[30] Panayotis Mertikopoulos and William H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, November 2016.

[31] Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, November 1990.

[32] Paul Milgrom and John Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior*, 3:82–100, 1991.

[33] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.

[34] K. Miyasawa. On the convergence of learning processes in a 2×2 non-zero sum game. Research memorandum 33, Princeton University, 1961.

[35] Dov Monderer and Lloyd S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68:256–265, 1996.

[36] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[37] Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*, pages 596–605, 2020.

[38] Michael Plum. Characterization and computation of nash-equilibria for auctions with incomplete information. *International Journal of Game Theory*, 20(4):393–418, 1992.

[39] Zinovi Rabinovich, Victor Naroditskiy, Enrico H Gerding, and Nicholas R Jennings. Computing pure bayesian-nash equilibria in games with finite actions and continuous types. *Artificial Intelligence*, 195:106–139, 2013.

[40] Itzhak Rasooly and Carlos Gavidia-Calderon. The importance of being discrete: on the inaccuracy of continuous approximations in auction theory, 2021.

[41] Philip J Reny and Shmuel Zamir. On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica*, 72(4):1105–1125, 2004.

[42] Julia Robinson. An iterative method for solving a game. *Annals of Mathematics*, 54:296–301, 1951.

[43] Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.

[44] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

[45] Lloyd S. Shapley. Some topics in two-person games. In *Advances in Game Theory*, number 52 in Annals of Mathematics Studies. Princeton University Press, 1964.

[46] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[47] Zihe Wang, Weiran Shen, and Song 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, page 1458–1466, Richland, SC, 2020. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450375184.