

Beyond Monotonicity: On the Convergence of Learning Algorithms in Standard Auction Games

Martin Bichler, Stephan B. Lunowa, Matthias Oberlechner, Fabian R. Pjeroth, Barbara Wohlmuth

Technical University of Munich
School of Computation, Information and Technology,
Arcisstr. 21, 80333 Munich, Germany
{bichler, stephan.lunowa, oberlech, fabian.pjeroth, wohlmuth}@cit.tum.de

Abstract

Equilibrium problems in Bayesian auction games can be described as systems of differential equations. Depending on the model assumptions, these equations might be such that we do not have a rigorous mathematical solution theory. The lack of analytical or numerical techniques with guaranteed convergence for the equilibrium problem has plagued the field and limited equilibrium analysis to rather simple auction models such as single-object auctions. Recent advances in equilibrium learning led to algorithms that find equilibrium under a wide variety of model assumptions. Monotonicity and the Minty condition are the known sufficient conditions for learning algorithms to converge to an equilibrium in games. Not much is known about convergence of learning algorithms beyond these conditions. We analyze first- and second-price auctions where simple learning algorithms consistently converge to an equilibrium. The analysis is challenging, because these properties need to be shown in infinite dimensions. Interestingly, we show that neither monotonicity nor pseudo- or quasi-monotonicity holds for the respective variational inequalities (VIs). The second-price auction's equilibrium is a Minty-type solution, but the first-price auction is not. However, the analysis via infinite-dimensional VIs allows us to get ex-post guarantees for gradient-based algorithms. We show that the Bayes–Nash equilibrium is the unique solution to the VI within the class of uniformly increasing bid functions, which ensures that gradient-based algorithms attain the equilibrium in case of convergence, as also observed in numerical experiments.

Extended version — <https://arxiv.org/abs/2311.15398>

Introduction

The Nash equilibrium is the central solution concept in non-cooperative game theory (Nash et al. 1950). Nash's original idea was extended to games with incomplete information (Harsanyi 1967). Today, the analysis of Bayes–Nash equilibrium (BNE) is key to the study of auctions, contests, and many other economic models, where the strategic complexity for participants arises from their uncertainty about the preferences of competitors (Vickrey 1961).

Auction theory studies allocation and prices on markets with self-interested participants in equilibrium. Such predictions are important because most market institutions in

this field do not have simple dominant strategies for bidding truthfully. The celebrated Vickrey–Clarke–Groves (VCG) mechanism was shown to be unique in this respect in the widely used standard model with independent private values and quasi-linear (i.e., payoff maximizing) bidders (Green and Laffont 1977). However, for various reasons, the VCG mechanism is rarely used in practice (Ausubel, Milgrom et al. 2006). Therefore, it is essential to understand non-truthful equilibrium strategies in markets more generally.

Participants in auction games do not have complete but only distributional information about the valuations of competing market participants. The value of competing bidders is modeled as a draw from an atomless distribution function. An equilibrium bid function then determines how much they bid based on their value draw and their knowledge of the prior distribution. The first-price sealed-bid auction in the symmetric and independent prior values model provides a textbook example, where the first-order conditions lead to an ordinary differential equation that admits a closed-form solution for the BNE bid function (Krishna 2009).

Unfortunately, despite the enormous academic attention auction theory has received, BNE strategies are only known for very restricted, simple market models. For bidders with non-uniform or interdependent valuation distributions, multiple objects, or non-quasilinear (i.e., non-payoff-maximizing) utility functions, we typically do not know an explicit equilibrium bid function. Even for simple combinatorial auctions with two objects only, one needs to make strong assumptions to find closed-form solutions (Kokott, Bichler, and Paulsen 2019). The equilibrium problem often leads to a system of non-linear differential equations for which we have no exact mathematical solution theory. Beyond analytical solutions, numerical techniques for solving differential equations turned out to be challenging and have received only limited attention. Fibich and Gavish (Fibich and Gavish 2011) criticize the inherent numerical instability of standard techniques used in this field.

Equilibrium Learning

Equilibrium learning is an alternative numerical approach to finding equilibrium, as compared to standard techniques for solving differential equations. This literature examines what kind of equilibrium arises as a consequence of a relatively simple process of learning and adaptation, in which agents

are trying to maximize their payoff while learning about the actions of other agents (Fudenberg and Levine 1998; Hart and Mas-Colell 2003; Young 2004). Learning provides an intuitive, tractable model of how equilibrium emerges but it is well established that independent learning dynamics do not generally obtain a Nash equilibrium (Benaïm and Hirsch 1999). Numerical analyses of matrix games show that gradient-based algorithms can circle, diverge, or they are even chaotic (Sanders, Farmer, and Galla 2018; Bielawski et al. 2021; Chotibut et al. 2020; Palaiopanos, Panageas, and Piliouras 2017; Vlatakis-Gkaragkounis, Flokas, and Piliouras 2023). Recently, Milionis et al. (Milionis et al. 2023) proved that there exist games, for which all game dynamics fail to converge to a Nash equilibrium from all starting points. On the other hand, there are also game classes such as potential games or games with strictly dominated strategies, where learning algorithms do converge.

Learning in Auctions

In recent years, a number of learning algorithms were introduced for Bayesian auction games with a continuous type and action space, and they showed convergence on a wide variety of auction models ranging from simple single-object auctions in the independent private values model to interdependent valuations and models with multiple objects (Bichler et al. 2021; Bichler, Kohring, and Heidekrüger 2023). Equilibrium can be verified ex-post, but the reasons for the convergence of such learning algorithms in auction games have been a mystery so far. If auctions are indeed learnable, this has significant consequences for theory and applications alike. First, numerical solvers could be provided for models that could not be solved analytically so far. Second, if already stylized repeated auctions were not learnable, what would this mean for the growing economy of autobidding agents in display ad auctions and related applications?

What is so different in auction games compared to general-sum games, such that we see convergence in such a wide variety of models? Recently, it has been shown that a large class of learning algorithms converge in auction games with complete information, where the valuations of bidders are known a priori (Kolombus and Nisan 2022). However, the challenge in auctions is that we have incomplete information and Bayesian models with continuous type- and action spaces are quite different.

We draw on the field of operator theory and infinite-dimensional variational inequalities, which provides us with a new lens to analyze auction-theoretical models. Every Nash Equilibrium (NE) can be seen as a solution to a Stampacchia-type variational inequality (VI), and in some cases, the reverse is also true, for example with quasi-concave utility functions (Migot and Cojocaru 2020). This connection also holds for auction games and infinite-dimensional VIs (Cavazzuti, Pappalardo, and Passacantando 2002). Interestingly, the link between Nash equilibria and VIs has been explored for traffic games (Patriksson and Rockafellar 2003) or Walrasian equilibrium (Jofré, Rockafellar, and Wets 2007), but not for Bayesian games with continuous type and action space as is the case in auction theory. Now, auctions need to be modeled as infinite-

dimensional variational inequalities, which is different to applications in finite games.

In the literature on variational inequalities, two sufficient conditions are known, for which some types of algorithms always converge to an equilibrium. They can be seen as a generalization of convexity in optimization. The *monotonicity* condition is the most well-known condition to guarantee convergence for VIs (Bauschke and Combettes 2017). Various first-order projection methods, as discussed by Tseng (Tseng 1995), converge to a unique solution of a monotone VI, and higher-order methods have also been developed (Adil et al. 2022; Lin and Jordan 2022). Monotonicity is also central for convergence guarantees in the recent literature on learning in games (Ratliff, Burden, and Sastry 2013; Chasnov et al. 2019). Apart from this, the Minty condition has drawn some attention, as extragradient algorithms are known to converge to equilibrium if this condition holds globally (Strodiot, Vuong, and Nguyen 2016; Song et al. 2020). This condition is also referred to as the Minty VI or dual VI of the Stampacchia-type VI (Ye 2022). In constant-sum games it is related to the celebrated smoothness condition of games (Anagnostides and Sandholm 2023) introduced by (Roughgarden 2015). It is natural to ask if the monotonicity and the Minty condition, two sufficient conditions, hold in auction games where learning algorithms seem to converge so consistently.

Contributions

Demonstrating the monotonicity of auction games is challenging. Whereas practical equilibrium learning algorithms employ some form of discretization, it is easy to check with SODA (Bichler, Fichtl, and Oberlechner 2023), that the monotonicity condition is sometimes violated. However, such non-monotonicities could arise due to the game's discretization. Violations in a discretized game might vanish in games with continuous types and actions, and the only way to understand whether monotonicity holds is to study these auctions in function space. We study whether monotonicity or the Minty condition is satisfied in infinite dimensions in a function space. If any of the two conditions were satisfied in function space, this would explain the convergence of algorithms also in discretized versions of the game, where the condition is violated (Glowinski, Lions, and Trémolières 1981).

We make several contributions: First, we recover the well-known symmetric equilibrium strategies for the first-price and the second-price sealed-bid auction in the symmetric independent private-values model (Krishna 2009), but with a new proof technique based on the Gateaux derivative of the ex-ante utility function, which is novel and useful in its own right. Second, this proof technique for equilibrium problems in auctions and the resulting operator for the Gateaux derivative allows us to analyze the monotonicity and the Minty conditions in infinite dimensions. Our findings reveal that the first- and the second-price auctions are neither monotone nor pseudo- or quasi-monotone. Thus, we look at the Minty condition for variational inequalities. While the dominant-strategy incentive-compatible second-price auction satisfies this condition, this is not the case for the first-price auction.

Third, the analysis of infinite-dimensional VIs allows us to derive an ex-post guarantee. This result is derived by showing the uniqueness of the VI solution (the BNE) for the two auction formats within the class of uniformly increasing bid functions, which ensures that every gradient-based learning algorithm must converge to this solution if it does converge to a pure-strategy profile within this function class. This is very helpful because one can certify an equilibrium upon convergence to a strategy profile, without costly exploration of deviating strategies.

Finally, we provide insights into the nature of the Minty-violations and show that such violations are without loss in simple parametric cases. In the space of piecewise-linear functions we show that the gradient flows lead to the BNE in the first-price auction independent of the starting point and independent of Minty violations. The learning literature has almost exclusively focused on complete-information games. Our analysis is the first to study convergence of learning algorithms in Bayesian auction games.

Related Literature

More than 60 years ago, Vickrey (Vickrey 1961) derived the Bayes–Nash equilibrium (BNE) strategy in a single-object first-price auction in the independent-private values model with symmetric bidders, uniform prior distributions, and quasilinear utility functions. The first-order conditions, together with the assumption of symmetric bidding behavior, lead to a linear ordinary differential equation, which has a closed-form solution for the BNE bidding strategy. Unfortunately, this is only the case for the symmetric and independent private values model. It turns out that deviations from this benchmark model lead to challenges in the equilibrium analysis. For example, even in the asymmetric independent private values model, where different bidders' values are drawn from different distributions, we get a system of non-linear differential equations; and no closed-form expression for the bidding strategies exist for general distributions (Hubbard and Paarsch 2014).

The first question is whether BNE always *exists* in such auction games. For finite, complete-information games, we know of the existence of a mixed Nash equilibrium (Nash et al. 1950). Glicksberg (Glicksberg 1952) extended the existence result to games with continuous and compact action sets. For Bayesian games with continuous action space, Jackson and Swinkles (Jackson and Swinkels 2005) provide assumptions for the existence of equilibrium in distributional strategies. For example, first-price and second-price single-unit auctions, all-pay auctions, double auctions, and multi-unit discriminatory or uniform price auctions were shown to have an equilibrium in distributional strategies, not necessarily in pure strategies. It is interesting to note that there are auction models where there is no BNE of the continuous game, but there are equilibria in the discretized game (Jackson et al. 2002). While there was significant progress, we do not have a complete theory when a BNE exists in general continuous-type and -action auction games (Carbonell-Nicolau and McLean 2018). Note that if one could prove the strict monotonicity of the game, this also proves the uniqueness of the equilibrium in the respective auction. Therefore,

we can already deduce that strict monotonicity cannot hold for the second-price auction considering asymmetric strategies, where several equilibria are known to exist (Krishna 2009, p. 118). However, we know that it is unique in the symmetric case.

Another question concerns the *computational complexity* in those cases where Nash equilibria do exist. Computing a Nash equilibrium in a finite, complete-information game is generally PPAD-hard (Daskalakis, Goldberg, and Papadimitriou 2009). Cai and Papadimitriou (Cai and Papadimitriou 2014) showed that finding an exact BNE in specific simultaneous auctions for multiple items is at least hard for PP, with PP being a complexity class higher than the polynomial hierarchy and close to PSPACE. Such problems can be considered intractable even for very small problem sizes. We know little about the complexity of finding BNE in other multi-item auctions. Recently, (Chen and Peng 2023) proved that there is a PTAS for approximating first-price auctions with a uniform tie-breaking rule, but that the approximation problem is PPAD-complete for other tie-breaking rules.

A related stream in the literature focuses on *learning in games* (Fudenberg and Levine 1998). For example, fictitious play is a natural method by which agents iteratively search for a pure Nash equilibrium and play the best response to the empirical frequency of play of other players (Brown 1951). Several algorithms have been proposed based on best or better response dynamics for finite and simultaneous-move games, and the literature is large (Abreu and Rubinstein 1988; Hart and Mas-Colell 2000; Fudenberg and Levine 1998; Hart and Mas-Colell 2003; Young 2004). Learning dynamics do not always converge to equilibrium (Daskalakis et al. 2010; Vlatakis-Gkaragkounis et al. 2020). Learning algorithms can cycle, diverge, or be chaotic; even in zero-sum games, where the Nash equilibrium is tractable (Mertikopoulos, Papadimitriou, and Piliouras 2018; Bailey and Piliouras 2018; Cheung and Piliouras 2020). Sanders et al. (Sanders, Farmer, and Galla 2018) argue that chaos is typical behavior for more general large matrix games. Recent results have shown that learning dynamics do not converge in games with mixed Nash equilibria (Bailey and Piliouras 2018; Letcher et al. 2019), and that there are finite games for which any dynamics is bound to have starting points that do not end up at a Nash equilibrium (Milionis et al. 2023). Overall, the dynamics of matrix games can be arbitrarily complex and hard to characterize (Andrade, Frongillo, and Piliouras 2021).

It is well-known that no-external-regret learners converge to the set of (Bayesian) coarse correlated equilibria, but the convergence of learning algorithms to a Nash equilibrium is only known for specific types of games such as finite potential games (Foster and Vohra 1997; Hart and Mas-Colell 2000; Jafari et al. 2001; Stoltz and Lugosi 2007; Hartline, Syrgkanis, and Tardos 2015; Foster et al. 2016). In general, the convergence of no-regret learning algorithms to a Nash equilibrium depends on the specific structure of the game and the algorithm used. The convergence of learning algorithms to a Nash equilibrium in Bayesian auction games with continuous type and action space is largely unexplored.

We draw on the literature on *variational inequalities* (VI),

in particular infinite-dimensional VIs. Variational inequalities were initially introduced to study equilibrium problems in physical systems (Fichera 1964; Kikuchi and Oden 1988). Since then, the theory's applicability has been expanded to include various problems (Lions and Stampacchia 1967; Glowinski, Lions, and Trémolières 1981; Kinderlehrer and Stampacchia 2000). Notably, the search for equilibrium in a game can be reformulated as solving a related VI. Auction models with continuous type and action space can be described as infinite-dimensional variational inequalities (Cavazzuti, Pappalardo, and Passacantando 2002).

Monotonicity can be seen as a generalization of convexity in optimization, and it is the primary property to show the convergence of algorithms in VIs. For first-order methods, Nemirovski (Nemirovski 2004) proved that the extragradient method converges to a weak solution with a global rate of $O(1/t)$ if the operator F is monotone and Lipschitz-continuous. Other methods achieve the same goal (Nesterov 2007; Kotsalis, Lan, and Li 2022) and achieve the lower bound of Ouyang and Xu (Ouyang and Xu 2021). These are also known as "projection-type" methods. More recently, also higher-order methods were developed (Lin and Jordan 2022). There is also more recent work on non-monotone VIs, which usually assumes the existence of Minty-type solutions to the VI (Song et al. 2020; Huang and Zhang 2023). Huang and Zhang (Huang and Zhang 2023) show that the existence of Minty solutions turns out to be critical in establishing convergence for the projection-type methods for non-monotone VIs.

Preliminaries

This section lays the foundation for studying the equilibrium problem and its associated variational inequality (VI) in a function space. To begin our analysis, it is crucial to establish a derivative in function space. This requires us to work with a set of strategies that exhibit sufficient well-behaved properties. Furthermore, we narrow our focus to the symmetric setting with symmetric priors and strategies and the independent private values model. This choice simplifies our analysis and is sufficient to give answers to whether forms of monotonicity are the reason for the convergence of first-order methods in auction models.

Abstract Setting

Let $n \in \mathbb{N}$ be the number of bidders. For bidder i , the set of possible bids is called $B_i \subset \mathbb{R}$, and the set of valuations of bidder i is $\mathcal{X}_i \subset \mathbb{R}$. We define $\mathbf{B} := \times_{i=1}^n B_i$ and $\mathcal{X} := \times_{i=1}^n \mathcal{X}_i$. The goal of each bidder is to maximize their payoff, i.e., they consider their utility function

$$u_i : \mathbf{B} \times \mathcal{X} \rightarrow \mathbb{R} : u_i(\mathbf{b}, \mathbf{x}).$$

For this, they search for a strategy $\beta_i : \mathcal{X} \rightarrow B_i$. Let the vector space V_i contain all possible strategies β_i , while the subset $\mathcal{B}_i \subset V_i$ contains all admissible strategies (see Equation 5), and we define $\mathbf{V} := \times_{i=1}^n V_i$ and $\mathcal{B} := \times_{i=1}^n \mathcal{B}_i$. The random values X_i of all bidders $i = 1, \dots, n$ are distributed in \mathcal{X}_i according to an atomless probability distribution F_i with full support over \mathcal{X}_i . We denote by $U_i : \mathbf{V} \rightarrow \mathbb{R} :$

$U_i(\boldsymbol{\beta}) := \mathbb{E}_{\mathbf{X}}[u_i(\boldsymbol{\beta}(\mathbf{X}), \mathbf{X})]$ the expected utility of bidder i for given strategies $\boldsymbol{\beta}$. Note that here and in the following, we denote by $\boldsymbol{\beta}$ a vector $(\beta_1, \beta_2, \dots, \beta_n)$ and by $\beta_i, \boldsymbol{\beta}_{-i}$ the vector $(\beta_1^*, \dots, \beta_{i-1}^*, \beta_i, \beta_{i+1}^*, \dots, \beta_n^*)$. The Bayesian Nash equilibrium (BNE) for this auction game is then given by:

Problem 1 (BNE). Find $\boldsymbol{\beta}^* \in \mathcal{B}$ such that $\forall i \in [n]$

$$U_i(\boldsymbol{\beta}^*) \geq U_i(\beta_i, \boldsymbol{\beta}_{-i}^*) \quad \forall \beta_i \in \mathcal{B}_i. \quad (1)$$

Under certain conditions (to be elaborated in the following), the equilibrium condition can be reformulated as a variational inequality. For this, we need to consider the expected utility function's derivative in a function space. This demands some regularity of the underlying function space. Assume that V_i is a Banach space and denote by $V_i^* := \mathcal{L}(V_i, \mathbb{R})$ its dual space consisting of all continuous linear functionals on V_i . Note that not all linear functionals are continuous for general infinite-dimensional spaces. Furthermore, assume that $\mathcal{B}_i \subset V_i$ is convex and closed. Following the standard procedure in the literature (Lions and Stampacchia 1967; Kinderlehrer and Stampacchia 2000), we derive the so-called Gateaux-derivative of U_i , which can be understood as the generalization of the (linear) directional derivative in normed spaces. So, let $DU_i(\boldsymbol{\beta})[d]$ denote the directional derivative of U_i at $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n) \in \mathcal{B}$ with respect to β_i along $d \in V_i$, i.e., for all $d \in V_i$:

$$DU_i(\boldsymbol{\beta})[d] := \lim_{\varepsilon \rightarrow 0} \varepsilon^{-1} (U_i(\beta_i + \varepsilon d, \boldsymbol{\beta}_{-i}) - U_i(\boldsymbol{\beta})). \quad (2)$$

The directional derivative is the Gateaux-derivative iff $DU_i(\boldsymbol{\beta}) \in V_i^*$, i. e., when the derivative is a continuous linear operator in the direction $d \in V_i$. If the Gateaux-derivative exists, a *necessary condition* for a BNE is the following (Stampacchia-type) VI (Kinderlehrer and Stampacchia 2000):

Problem 2 (VI). Find $\boldsymbol{\beta}^* \in \mathcal{B}$ such that $\forall i \in [n]$

$$DU_i(\boldsymbol{\beta}^*)[\beta_i - \beta_i^*] \leq 0 \quad \forall \beta_i \in \mathcal{B}_i. \quad (3)$$

A *sufficient condition*, also referred to as the dual VI of the Stampacchia formulation, is given by the (Minty-type) VI:

Problem 3 (MVI). Find $\boldsymbol{\beta}^* \in \mathcal{B}$ which satisfies $\forall i \in [n]$

$$DU_i(\boldsymbol{\beta})[\beta_i - \beta_i^*] \leq 0 \quad \forall \boldsymbol{\beta} \in \mathcal{B}. \quad (4)$$

In general, solutions of the Minty-type VI (4) are a subset of the BNEs given by (1), which are, in turn, a subset of solutions of the Stampacchia-type VI (3). Vice versa, solutions of the Stampacchia-type VI (3) are also BNEs if U_i is pseudoconvex in β_i for all $\boldsymbol{\beta}_{-i}$, and BNEs are, in turn, solutions of the Minty-type VI (4) if U_i is quasiconvex in β_i for all $\boldsymbol{\beta}_{-i}$ (Cavazzuti, Pappalardo, and Passacantando 2002).

Symmetric and Independent Private Value Auctions

In the following sections, we consider second- and first-price sealed-bid auctions under the assumption of (complete) symmetry and identically independently distributed private values. (Complete) symmetry implies

$$B_i = B, \quad \mathcal{X}_i = \mathcal{X}, \quad X_i \sim_{iid} F_i \equiv F, \quad u_i \equiv u$$

for all $i = 1, \dots, n$. Furthermore, private values ensure $\beta_i(\mathbf{x}) = \beta_i(x_i)$ for $i = 1, \dots, n$, i.e., the strategy of each bidder i depends only on the knowledge of their own valuation $X_i = x_i$. The ex-ante utility is denoted $U(\beta) := \mathbb{E}_{\mathbf{X}}[u(\beta(\mathbf{X}), \mathbf{X})]$, and symmetric strategies are denoted $\tilde{\beta} := (\tilde{\beta}_1, \dots, \tilde{\beta}_n) \in \mathcal{B}$ for $\tilde{\beta} \in \mathcal{B}$.

In the following, we assume $\mathcal{X} = [0, 1]$ (without loss of generality) and $F \in C^{0,1}([0, 1])$, i.e., the cumulative probability function is Lipschitz-continuous. To analyze the VI we have to define an appropriate set of admissible strategies. This set should be sufficiently general to allow for strategies that may be considered as sensible for the underlying problem. Additionally, it needs to provide adequate structure, e.g., an inner product or a natural dual product. Therefore, consider the Banach space

$$V_i = V := \{\beta \in L^1(0, 1; F) \mid \beta' \in L^1(0, 1; F)\},$$

i.e., $V = W^{1,1}(0, 1; F)$ consists of F -integrable functions with F -integrable weak derivatives. Note that $V \subseteq AC([0, 1])$, where the latter space denotes all absolutely continuous functions on $[0, 1]$. For small $\delta > 0$ we define

$$\mathcal{B}_\delta := \{\beta \in V : 0 \leq \beta \leq 1 \text{ F-a.e.}, \\ 0 < \delta \leq \beta' \text{ F-a.e.}, \text{ and } \beta(0) = 0\}. \quad (5)$$

Note that the restriction $0 \leq \beta \leq 1$ is natural because only positive bids are feasible, and bidding more than the maximal valuation 1 implies a non-positive payoff. Similarly, it is natural to assume the bids β to be increasing in valuation. Requiring a small positive derivative ($\beta' \geq \delta > 0$) is slightly more restrictive, but together with the other assumptions, this ensures the set \mathcal{B}_δ to be convex, closed, and bounded in V .

In this setting, the BNE (1), VI (3) and MVI (4) simplify to deviations in a single strategy:

Problem 4 (Symmetric BNE, VI and MVI). *A symmetric BNE $\beta^* \in \mathcal{B}_\delta$ satisfies*

$$U(\beta^*) \geq U(\beta, \beta_{-1}^*), \quad \forall \beta \in \mathcal{B}_\delta. \quad (6)$$

A solution $\beta^ \in \mathcal{B}_\delta$ to the symmetric VI satisfies*

$$DU(\beta^*)[\beta - \beta^*] \leq 0 \quad \forall \beta \in \mathcal{B}_\delta. \quad (7)$$

A solution $\beta^ \in \mathcal{B}_\delta$ to the symmetric MVI satisfies*

$$DU(\beta, \tilde{\beta}_{-1})[\beta - \beta^*] \leq 0 \quad \forall \beta, \tilde{\beta} \in \mathcal{B}_\delta. \quad (8)$$

Here $DU(\beta)[d]$ is the Gateaux-derivative of U at $\beta \in \mathcal{B}_\delta$ with respect to β_1 along $d \in V$.

In the following two sections, we use these mathematical tools to analyze the second- and first-price sealed-bid auction in the continuous setting. Thereto, we first derive the Gateaux-derivative of the bidder's utility function, and use it to analyze the BNE, VI and MVI and their potentially different solutions in detail for these applications. Subsequently, we analyze whether the Gateaux-derivative is (quasi-)monotone, since this property would ensure convergence of (certain) gradient-based learning algorithms. In particular, we discuss counterexamples for the simplest case of two bidders with independent uniform priors.

Second-Price Sealed-Bid Auction

For symmetric second-price sealed-bid auctions with risk-neutral bidders, the utility function of a bidder is given by

$$u(\mathbf{b}, \mathbf{x}) = \chi_{\{b_1 > \max_{j=2, \dots, n} b_j\}} \left(x_1 - \max_{j=2, \dots, n} b_j \right)$$

Since every $\tilde{\beta} \in \mathcal{B}_\delta$ is an increasing function, it satisfies $\max_{j=2, \dots, n} \tilde{\beta}(x_j) = \tilde{\beta}(\max_{j=2, \dots, n} x_j)$. Using $Y := \max_{j=2, \dots, n} X_j \sim G := F^{n-1} \in C^{0,1}([0, 1])$ with derivative $g := G'$, the ex-ante utility U against symmetric strategies $\tilde{\beta} = (\tilde{\beta}_1, \dots, \tilde{\beta}_n)$ can then be reformulated as

$$U(\beta, \tilde{\beta}_{-1}) = \int_0^1 \int_0^1 \chi_{\{\beta(x) > \tilde{\beta}(y)\}} (x - \tilde{\beta}(y)) dG(y) dF(x).$$

This leads to the following expression for the derivative:

Lemma 1. *The Gateaux-derivative at $(\beta, \tilde{\beta}_{-1}) \in \mathcal{B}_\delta$ along $d \in V$ is given by*

$$DU(\beta, \tilde{\beta}_{-1})[d] = \int_0^1 d(x) \chi_{\{\beta(x) \in \text{Im}(\tilde{\beta})\}} (x - \beta(x)) \cdot \frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))} dF(x). \quad (9)$$

Proof sketch. Using the definition of the directional derivative (2), we get the expression above and observe that the expression is linear in d and bounded. Therefore, it is equivalent to the Gateaux derivative. \square

Existence and Uniqueness

For symmetric second-price sealed-bid auctions with independent private values, we can show that a unique BNE exists and coincides with the (unique) solution of the VI and of the MVI. Therefore, these notions are equivalent in this particular case, even though we show in the following section that the Gateaux-derivative is not monotone, nor pseudo- nor quasi-monotone.

Lemma 2. *The symmetric BNE, VI and MVI problems have the (F-a.e.) unique solution $\beta^* = \text{Id}$ in the compact and convex set $\mathcal{B}_\delta \subset V$ for $0 < \delta \leq 1$.*

Proof sketch. Using the expression (9) for DU , we derive the symmetric VI (7)

$$0 \geq DU(\beta^*)[\beta - \beta^*] \\ = \int_0^1 (\beta(x) - \beta^*(x)) (x - \beta^*(x)) \frac{g(x)}{\beta^{*'}(x)} dF(x) \quad (10)$$

for all $\beta \in \mathcal{B}_\delta$. We observe that $\beta^* = \text{Id}$ satisfies the VI, showing that it is a solution. Furthermore, we observe that the VI with $\beta = \text{Id}$ can only be satisfied by $\beta^* = \text{Id}$, showing uniqueness. \square

Monotonicity

Gradient-based learning for symmetric strategies uses the gradient operator $DU(\beta, \tilde{\beta}_{-1})$ with $\beta = \tilde{\beta}$. Therefore, we are interested whether the operator $DU(\beta)$ is (quasi-)monotone in $\beta \in \mathcal{B}_\delta$. This condition would ensure convergence for extra-gradient methods (Khanh 2016). However, we show that even in the most simple setting with two bidders ($n = 2$) and uniform priors ($F = \text{Id}$), the operator DU is neither monotone nor pseudo- nor quasi-monotone.

The operator DU is monotone if it satisfies

$$(DU(\tilde{\beta}) - DU(\beta))[\tilde{\beta} - \beta] \leq 0 \quad \forall \beta, \tilde{\beta} \in \mathcal{B}_\delta. \quad (11)$$

For pseudo-monotonicity we require (Khanh 2016)

$$DU(\beta)[\tilde{\beta} - \beta] \leq 0 \quad \Rightarrow \quad DU(\tilde{\beta})[\tilde{\beta} - \beta] \leq 0. \quad (12)$$

for all $\beta, \tilde{\beta} \in \mathcal{B}_\delta$, while quasi-monotonicity requires (12) with strict inequality on the left-hand side. Note that monotonicity implies pseudo-monotonicity, which, in turn, implies quasi-monotonicity.

Proposition 1. *The operator DU is neither monotone, nor pseudo- nor quasi-monotone (for $\delta < \frac{9}{100}$, $F = \text{Id}$ and $n = 2$).*

Proof. A counterexample is given by the piece-wise linear and continuous functions

$$\beta(x) = \frac{61x}{100}, \quad \tilde{\beta}(x) = \begin{cases} x & \text{if } 0 \leq x \leq \frac{1}{3}, \\ \frac{9x}{100} + \frac{91}{300} & \text{if } \frac{1}{3} < x \leq \frac{2}{3}, \\ \frac{63x}{100} - \frac{17}{300} & \text{if } \frac{2}{3} < x \leq 1, \end{cases}$$

which yield a negative integral on the left-hand side, but a positive one on the right-hand side in (12). Therefore, DU is not quasi-monotone and consequently neither pseudo-monotone nor monotone. \square

Even though this second-price auction leads to a non-monotone VI, a Minty-type solution exists, ensuring the convergence of projection type methods (Song et al. 2020; Huang and Zhang 2023), see also the extended version.

First-Price Sealed-Bid Auction

For symmetric first-price sealed-bid auctions with risk-neutral bidders, we have

$$u(\mathbf{b}, \mathbf{x}) = \chi_{\{b_1 > \max_{j=2, \dots, n} b_j\}} (x_1 - b_1)$$

As for the second-price auction, using $Y := \max_{j=2, \dots, n} X_j \sim G := F^{n-1}$, the ex-ante utility U against symmetric strategies $\tilde{\beta} = (\tilde{\beta}, \dots, \tilde{\beta})$ can be reformulated as

$$U(\beta, \tilde{\beta}_{-1}) = \int_0^1 (x - \beta(x)) \int_0^1 \chi_{\{\beta(x) > \tilde{\beta}(y)\}} dG(y) dF(x).$$

Lemma 3. *The Gateaux-derivative at $(\beta, \tilde{\beta}_{-1}) \in \mathcal{B}_\delta$ along $d \in V$ is given by*

$$DU(\beta, \tilde{\beta}_{-1})[d] = \int_0^1 d(x) \chi_{\{\beta(x) < \tilde{\beta}(1)\}} \quad (13)$$

$$\left[(x - \beta(x)) \frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))} - G(\tilde{\beta}^{-1}(\beta(x))) \right] dF(x).$$

The proof is analogous to the derivation in the previous section and can again be found in the extended version.

Uniqueness of BNE and Non-Existence of MVI

For symmetric first-price sealed-bid auctions with independent private values, we can show that the unique BNE coincides with a solution of the VI (which is unique in the interior of \mathcal{B}_δ). Even in the simple case of two bidders ($n = 2$) with uniform priors ($F = \text{Id}$), no solution to the MVI exists. So, in contrast to second-price auctions, these notions are different for first-price auctions.

Lemma 4. *Assume that $f = F'$ is Lipschitz-continuous and satisfies*

$$\delta_0 := \frac{\inf_{x \in [0,1]} f(x)}{\sup_{x \in [0,1]} f(x)} > 0.$$

For $0 < \delta \leq \delta_0$, $\beta^*(x) = \frac{1}{G(x)} \int_0^x y dG(y)$ is the unique solution to the symmetric VI (7) in the interior of \mathcal{B}_δ . This solution is the unique symmetric BNE of (6).

Proof sketch. We start by deriving the symmetric variational inequality (VI) in terms of β^*

$$\int_{\mathcal{X}} \left[(x - \beta^*(x)) \frac{g(x)}{\beta^{*'}(x)} - G(x) \right] (\beta(x) - \beta^*(x)) dF(x) \leq 0,$$

for all $\beta \in \mathcal{B}_\delta$. A solution in the interior of \mathcal{B}_δ must satisfy

$$(x - \beta^*(x)) \frac{g(x)}{\beta^{*'}(x)} - G(x) = 0 \quad F\text{-a.e.}$$

Finally, we use the Picard–Lindelöf theorem to analyze the above ODE and derive the unique solution. The detailed proof can be found in the extended version. \square

Remark 1. *Note that Lemma 4 holds for all $0 < \delta \leq \delta_0$. Hence, a limit argument shows that the BNE is the unique solution to the symmetric VI in the interior of the class of uniformly increasing functions $B_{0+} := \bigcup_{\delta > 0} \mathcal{B}_\delta$. If a solution β to the VI at the boundary ∂B_{0+} exists, its derivative β' must approach zero at some point $x \in [0, 1]$, such that the expression $DU(\beta)$ might be ill-defined. In particular, this implies that a gradient-based learning algorithm must reach the BNE if it does converge (within B_{0+}).*

Lemma 5. *In the case of two bidders with uniform priors, i.e., for $n = 2$ and $F = \text{Id}$, the unique BNE $\beta^*(x) = \frac{x}{2}$ according to Lemma 4 does not satisfy the symmetric MVI (8). In particular, the condition is also not satisfied locally for any open neighborhood of the BNE (for $\delta \leq \frac{1}{5}$).*

Proof. Using (13), the symmetric MVI (8) reads

$$\int_0^1 \left[\frac{x - \beta(x)}{\beta'(x)} - x \right] (\beta(x) - \beta^*(x)) dF(x) \leq 0, \quad (14)$$

for all $\beta \in \mathcal{B}_\delta$. Using $\beta^*(x) = \frac{x}{2}$ and the continuous, piece-wise linear and strictly increasing bid function

$$\beta(x) = \begin{cases} \frac{x}{2} & \text{for } x \leq \frac{n}{n+2}, \\ \frac{n}{2(n+2)} + \frac{4}{5} \left(x - \frac{n}{n+2} \right) & \text{for } \frac{n}{n+2} < x \leq \frac{n+1}{n+2}, \\ \frac{n}{2(n+2)} + \frac{3-n}{5(n+2)} + \frac{x}{5} & \text{for } \frac{n+1}{n+2} < x, \end{cases}$$

for arbitrary $n \in \mathbb{N}_0$, we obtain on the left-hand side of (14) a positive value which then contradicts (14). In particular, this variational stability condition is even violated locally, since $\beta \rightarrow \beta^*$ in V as $n \rightarrow \infty$. \square

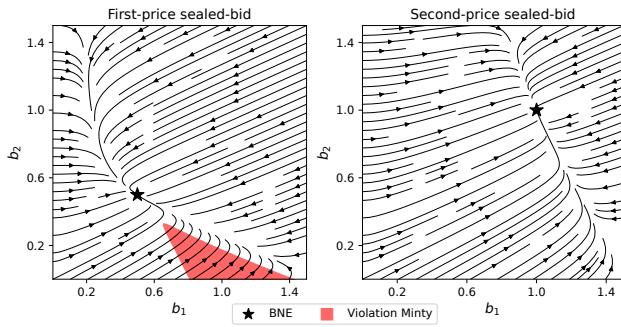


Figure 1: Gradients for piecewise linear bid functions with two pieces.

Monotonicity

As before, we are interested in the situation $\beta = \tilde{\beta}$ along which gradient-based learning takes place, and study whether the operator DU is (quasi-)monotone in \mathcal{B}_δ . Again, we show that even in the most simple setting of two bidders ($n = 2$) with uniform priors ($F = \text{Id}$), the operator DU turns out to be neither monotone, nor pseudo- nor quasi-monotone.

Proposition 2. *The operator DU is neither monotone, nor pseudo- nor quasi-monotone (for $0 < \delta \leq \frac{1}{10}$).*

Proof. For $F(x) = x$ and $n = 2$, and using (9), we can have a counterexample given by the piece-wise linear and continuous functions

$$\beta(x) = \frac{61x}{100}, \quad \tilde{\beta}(x) = \begin{cases} x & \text{for } x \leq \frac{1}{3}, \\ \frac{x}{10} + \frac{3}{10} & \text{for } \frac{1}{3} < x \leq \frac{2}{3}, \\ \frac{63x}{100} - \frac{4}{75} & \text{for } \frac{2}{3} < x, \end{cases}$$

which yield a negative integral on the left-hand side, but a positive one on the right-hand side in (12). Therefore, DU is not quasi-monotone and consequently neither pseudo-monotone nor monotone. Note that $\tilde{\beta}' \geq \frac{1}{10} \geq \delta$. \square

Numerical Analysis

In order to better understand the violations of the Minty condition in first-price auctions, we analyze symmetric learning in the space of piecewise linear functions. More specifically, we analyze bid functions consisting of two linear functions in a first- and second-price sealed-bid auction with two bidders having a uniform prior distribution. These bid functions are parametrized by the slopes b_1, b_2 of the two pieces, i.e.,

$$\beta(x) = \begin{cases} b_1 x & \text{if } x \leq \frac{1}{2} \\ \frac{b_1}{2} + b_2(x - \frac{1}{2}) & \text{if } \frac{1}{2} < x. \end{cases} \quad (15)$$

This low-parameter environment allows us to examine the resulting vector field and those areas where the Minty condition is violated (see Figure 1). The unique Bayes–Nash equilibrium (marked with a star) is to have a slope of $\frac{1}{2}$ for the first-price auction and a slope of 1 for the second-price auction for both linear functions. The areas colored red mark where the Minty condition is violated. The figure shows that

violations of the Minty condition, a sufficient condition for convergence, are without loss in this analysis. The laminar flows take a turn, but no matter with which parameter combination one starts, the gradient flow always leads to the Bayes–Nash equilibrium. Traversing the red area means that the strategies analyzed on the way move away from the Bayes–Nash equilibrium, but they move closer again once the algorithm steps outside this area.¹

Conclusions

Learning in games has received much recent attention in the literature. It is well known, that learning algorithms do not always converge to an equilibrium in games, but they do converge in some types such as potential games. Recent advances in equilibrium learning showed that learning algorithms converge in a wide variety of auction games. The reasons for these observations are not well understood. We draw on the connection between auction games and infinite-dimensional variational inequalities, which has not been explored so far. In particular, there are sufficient conditions for which it has been shown that independent optimization algorithms find a solution to the variational inequality. Monotonicity can be seen as a generalization of convexity in optimization, and it provides a condition for first-order optimization methods to converge to a solution of the variational inequality. Our analysis shows that neither the second- nor the first-price auctions are monotonous. There are even counterexamples for the weaker pseudo- and quasi-monotonicity conditions. More recent literature on non-monotone variational inequality uses the Minty condition to show the convergence of extragradient algorithms (cf. the extended version). In the first-price auction, this condition is also not satisfied, even when assuming a uniform prior.

However, both auctions have a unique solution to the variational inequality, which means, that every gradient-based learning algorithm must attain the Bayes–Nash equilibrium (BNE) if it does converge. This is a helpful property as costly equilibrium verification can be avoided and learning itself is often fast. There is still a possibility that gradient dynamics might not converge to a unique solution but to a limit cycle, or diverge. Studying games as a dynamical system provides a potential remedy. In general, finding limit cycles of dynamical systems is a PSPACE-complete problem (Papadimitriou and Vishnoi 2015). However, our numerical explorations illustrate that the gradient dynamics do converge in the first- and second-price auctions, and they do not cycle. Whether limit cycles or divergence are possible in other auction games, remains an open research question to be studied in the future.

Acknowledgments

This project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under

¹Let us note that the utility of the agents increases with lower slopes. Utility does not increase along the trajectories of the gradient such that the utility function is no Lyapunov function of the dynamical system even in the symmetric case.

Grant No. BI 1057/9; WO 671/11; and Project Number 277991500.

References

- Abreu, D.; and Rubinstein, A. 1988. The structure of Nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6): 1259–1281.
- Adil, D.; Bullins, B.; Jambulapati, A.; and Sachdeva, S. 2022. Optimal Methods for Higher-Order Smooth Monotone Variational Inequalities. *arXiv preprint arXiv:2205.06167*.
- Anagnostides, I.; and Sandholm, T. 2023. On the Interplay between Social Welfare and Tractability of Equilibria. *arXiv preprint arXiv:2310.16976*.
- Andrade, G. P.; Frongillo, R.; and Piliouras, G. 2021. Learning in Matrix Games can be Arbitrarily Complex. *arXiv preprint arXiv:2103.03405*.
- Ausubel, L. M.; Milgrom, P.; et al. 2006. The Lovely but Lonely Vickrey Auction. *Combinatorial auctions*, 17: 22–26.
- Bailey, J. P.; and Piliouras, G. 2018. Multiplicative Weights Update in Zero-Sum Games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 321–338. ACM.
- Bauschke, H. H.; and Combettes, P. L. 2017. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Cham: Springer International Publishing.
- Benaïm, M.; and Hirsch, M. W. 1999. Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games. *Games and Economic Behavior*, 29: 36–72.
- Bichler, M.; Fichtl, M.; Heidekrueger, S.; Kohring, N.; and Sutterer, P. 2021. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Machine Intelligence*, 687–695.
- Bichler, M.; Fichtl, M.; and Oberlechner, M. 2023. Computing Bayes-Nash equilibrium in auction games via gradient dynamics. In *ACM Conference on Economics and Computation*.
- Bichler, M.; Kohring, N.; and Heidekrüger, S. 2023. Learning Equilibria in Asymmetric Auction Games. *INFORMS Journal on Computing*.
- Bielawski, J.; Chotibut, T.; Falniowski, F.; Kosiorowski, G.; Misiurewicz, M.; and Piliouras, G. 2021. Follow-the-regularized-leader routes to chaos in routing games. In *International Conference on Machine Learning*, 925–935. PMLR.
- Brown, G. W. 1951. Iterative Solution of Games by Fictitious Play. *Activity Analysis of Production and Allocation*, 13(1): 374–376.
- Cai, Y.; and Papadimitriou, C. 2014. Simultaneous Bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 895–910.
- Carbonell-Nicolau, O.; and McLean, R. P. 2018. On the Existence of Nash Equilibrium in Bayesian Games. *Mathematics of Operations Research*, 43(1): 100–129.
- Cavazzuti, E.; Pappalardo, M.; and Passacantando, M. 2002. Nash Equilibria, Variational Inequalities, and Dynamical Systems. *Journal of Optimization Theory and Applications*, 114(3): 491–506.
- Chasnov, B.; Ratliff, L. J.; Mazumdar, E.; and Burden, S. A. 2019. Convergence Analysis of Gradient-Based Learning with Non-Uniform Learning Rates in Non-Cooperative Multi-Agent Settings. *arXiv preprint arXiv:1906.00731*.
- Chen, X.; and Peng, B. 2023. Complexity of Equilibria in First-Price Auctions under General Tie-Breaking Rules. *arXiv preprint arXiv:2303.16388*.
- Cheung, Y. K.; and Piliouras, G. 2020. Chaos, extremism and optimism: Volume analysis of learning in games. *Advances in Neural Information Processing Systems*, 33: 9039–9049.
- Chotibut, T.; Falniowski, F.; Misiurewicz, M.; and Piliouras, G. 2020. The route to chaos in routing games: When is price of anarchy too optimistic? *Advances in Neural Information Processing Systems*, 33: 766–777.
- Daskalakis, C.; Frongillo, R.; Papadimitriou, C. H.; Pierrakos, G.; and Valiant, G. 2010. On learning algorithms for Nash equilibria. In *International Symposium on Algorithmic Game Theory*, 114–125. Springer.
- Daskalakis, C.; Goldberg, P.; and Papadimitriou, C. 2009. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1): 195–259.
- Fibich, G.; and Gavish, N. 2011. Numerical simulations of asymmetric first-price auctions. *Games and Economic Behavior*, 73(2): 479–495.
- Fichera, G. 1964. Problemi elastostatici con vincoli unilaterali: Il problema di Signorini con ambigue condizioni al contorno. *Atti Accad. Naz. Lincei, Mem., Cl. Sci. Fis. Mat. Nat., VIII. Ser., Sez. I*, 7: 91–140.
- Foster, D. J.; Li, Z.; Lykouris, T.; Sridharan, K.; and Tardos, E. 2016. Learning in games: Robustness of fast convergence. In *Proceedings of Advances in Neural Information Processing Systems*, 4734–4742.
- Foster, D. P.; and Vohra, R. V. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2): 40.
- Fudenberg, D.; and Levine, D. K. 1998. *The theory of learning in games*, volume 2. MIT press.
- Glicksberg, I. L. 1952. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1): 170–174.
- Glowinski, R.; Lions, J. L.; and Trémoières, R. 1981. *Numerical Analysis of Variational Inequalities*, volume 8 of *Studies in mathematics and its applications*. Elsevier.
- Green, J.; and Laffont, J.-J. 1977. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods. *Econometrica*, 45(2): 427–438.
- Harsanyi, J. C. 1967. Games with incomplete information played by “Bayesian” players, I–III Part I. The basic model. *Management science*, 14(3): 159–182.
- Hart, S.; and Mas-Colell, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5): 1127–1150.
- Hart, S.; and Mas-Colell, A. 2003. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5): 1830–1836.
- Hartline, J.; Syrgkanis, V.; and Tardos, E. 2015. No-Regret Learning in Bayesian Games. In *Advances in Neural Information Processing Systems* 28, 3061–3069. Curran Associates, Inc.
- Huang, K.; and Zhang, S. 2023. Beyond Monotone Variational Inequalities: Solution Methods and Iteration Complexities. *arXiv preprint arXiv:2304.04153*.
- Hubbard, T. P.; and Paarsch, H. J. 2014. Chapter 2 - On the Numerical Solution of Equilibria in Auction Models with Asymmetries within the Private-Values Paradigm. In Schmedders, K.; and Judd, K. L., eds., *Handbook of Computational Economics. Vol. 3*, 37–115. Elsevier.
- Jackson, M. O.; Simon, L. K.; Swinkels, J. M.; and Zame, W. R. 2002. Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, 70(5): 1711–1740.

- Jackson, M. O.; and Swinkels, J. M. 2005. Existence of equilibrium in single and double private value auctions. *Econometrica*, 73(1): 93–139.
- Jafari, A.; Greenwald, A.; Gondek, D.; and Ercal, G. 2001. On no-regret learning, fictitious play, and nash equilibrium. In *Proceedings of International Conference on Machine Learning*, volume 1, 226–233.
- Jofré, A.; Rockafellar, R. T.; and Wets, R. J. 2007. Variational inequalities and economic equilibrium. *Mathematics of Operations Research*, 32(1): 32–50.
- Khanh, P. D. 2016. A Modified Extragradient Method for Infinite-Dimensional Variational Inequalities. *Acta Mathematica Vietnamica*, 41(2): 251–263.
- Kikuchi, N.; and Oden, J. T. 1988. *Contact Problems in Elasticity*. Society for Industrial and Applied Mathematics.
- Kinderlehrer, D.; and Stampacchia, G. 2000. *An Introduction to Variational Inequalities and Their Applications*. SIAM.
- Kokott, G.-M.; Bichler, M.; and Paulsen, P. 2019. The beauty of Dutch: Ex-post split-award auctions in procurement markets with diseconomies of scale. *European Journal of Operational Research*, 278(1): 202–210.
- Kolumbus, Y.; and Nisan, N. 2022. Auctions between regret-minimizing agents. In *Proceedings of the ACM Web Conference 2022*, 100–111.
- Kotsalis, G.; Lan, G.; and Li, T. 2022. Simple and optimal methods for stochastic variational inequalities, I: Operator extrapolation. *SIAM Journal on Optimization*, 32(3): 2041–2073.
- Krishna, V. 2009. *Auction Theory*. Academic Press.
- Letcher, A.; Balduzzi, D.; Racanière, S.; Martens, J.; Foerster, J. N.; Tuyls, K.; and Graepel, T. 2019. Differentiable Game Mechanics. *Journal of Machine Learning Research*, 20(84): 1–40.
- Lin, T.; and Jordan, M. I. 2022. Perseus: A Simple and Optimal High-Order Method for Variational Inequalities. *University of California, Berkeley*.
- Lions, J. L.; and Stampacchia, G. 1967. Variational Inequalities. *Communications on Pure and Applied Mathematics*, 20(3): 493–519.
- Mertikopoulos, P.; Papadimitriou, C.; and Piliouras, G. 2018. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2703–2717. SIAM.
- Migot, T.; and Cojocaru, M.-G. 2020. A parametrized variational inequality approach to track the solution set of a generalized nash equilibrium problem. *European Journal of Operational Research*, 283(3): 1136–1147.
- Milionis, J.; Papadimitriou, C.; Piliouras, G.; and Spendlove, K. 2023. An impossibility theorem in game dynamics. *Proceedings of the National Academy of Sciences*, 120(41): e2305349120.
- Nash, J. F.; et al. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49.
- Nemirovski, A. 2004. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1): 229–251.
- Nesterov, Y. 2007. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2-3): 319–344.
- Ouyang, Y.; and Xu, Y. 2021. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming*, 185(1-2): 1–35.
- Palaiopoulos, G.; Panageas, I.; and Piliouras, G. 2017. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. *Advances in Neural Information Processing Systems*, 30.
- Papadimitriou, C. H.; and Vishnoi, N. K. 2015. On the computational complexity of limit cycles in dynamical systems. *arXiv preprint arXiv:1511.07605*.
- Patriksson, M.; and Rockafellar, R. T. 2003. Sensitivity analysis of aggregated variational inequality problems, with application to traffic equilibria. *Transportation Science*, 37(1): 56–68.
- Ratliff, L. J.; Burden, S. A.; and Sastry, S. S. 2013. Characterization and computation of local Nash equilibria in continuous games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 917–924.
- Roughgarden, T. 2015. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)*, 62(5): 1–42.
- Sanders, J. B.; Farmer, J. D.; and Galla, T. 2018. The prevalence of chaotic dynamics in games with many players. *Scientific Reports*, 8(1): 1–13.
- Song, C.; Zhou, Z.; Zhou, Y.; Jiang, Y.; and Ma, Y. 2020. Optimistic Dual Extrapolation for Coherent Non-monotone Variational Inequalities. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 14303–14314. Curran Associates, Inc.
- Stoltz, G.; and Lugosi, G. 2007. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1): 187–208.
- Strodiot, J. J.; Vuong, P. T.; and Nguyen, T. T. V. 2016. A class of shrinking projection extragradient methods for solving non-monotone equilibrium problems in Hilbert spaces. *Journal of Global Optimization*, 64(1): 159–178.
- Tseng, P. 1995. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2): 237–252.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1): 8–37.
- Vlatakis-Gkaragkounis, E.-V.; Flokas, L.; Mertikopoulos, P.; and Piliouras, G. 2020. No-regret learning and mixed Nash equilibria: They do not mix. In *Annual Conference on Neural Information Processing Systems*.
- Vlatakis-Gkaragkounis, E.-V.; Flokas, L.; and Piliouras, G. 2023. Chaos persists in large-scale multi-agent learning despite adaptive learning rates. *arXiv:2306.01032*.
- Ye, M. 2022. An infeasible projection type algorithm for nonmonotone variational inequalities. *Numerical Algorithms*, 89(4): 1723–1742.
- Young, H. P. 2004. *Strategic learning and its limits*. OUP Oxford.