Submitted to *Econometrica*

THE CONVERGENCE OF LEARNING ALGORITHMS IN BAYESIAN AUCTION GAMES

MARTIN BICHLER, STEPHAN B. LUNOWA, MATTHIAS OBERLECHNER, FABIAN R. PIEROTH, BARBARA WOHLMUTH
School of Computation, Information and Technology, Technical University of Munich, Germany

Many digital markets, such as display advertising exchanges, are run as repeated first- or second-price auctions and are increasingly automated by learning agents. Recent empirical work shows that simple learning algorithms converge to an equilibrium in such settings, yet the reasons for this convergence remain elusive. We model the equilibrium problem as an infinite-dimensional variational inequality and analyze the associated dynamical system induced by gradient-based learning. We show that known sufficient conditions for convergence – such as strict monotonicity or the Minty condition – do not hold. While the second-price auction admits a Minty-type solution, the first-price auction does not. To prove convergence in the latter, we construct a Lyapunov function in the space of piecewise linear bid functions. Our approach provides the first ex-ante convergence proof for learning in first- and second-price auctions and establishes a new framework for analyzing the asymptotic stability of learning dynamics in these games.

KEYWORDS: Bayes–Nash equilibrium, Variational inequality, Learning algorithm.

1. INTRODUCTION

Auction theory analyzes how goods are allocated and prices are determined in markets with self-interested participants acting in equilibrium. Auctions are typically modeled as

Martin Bichler, Stephan B. Lunowa, Matthias Oberlechner, Fabian R. Pieroth, Barbara Wohlmuth: wohlmuth@cit.tum.de

We would like to express our sincere gratitude to Mete Ahunbay, Paul Milgrom, Bary Pradelski, Marco Scarsini, and Stefan Weltge, as well as participants of the Tolouse Workshop on Learning in Games 2024, ALGA 2025, and AAAI 2025 for their feedback on this research. This project has received funding from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement no. 101198689), and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grant No. WO 671/11 and Project Number 277991500.

Bayesian games with continuous types and action spaces, with the Bayes–Nash equilibrium (BNE) serving as the central solution concept. While the Vickrey auction admits a dominant-strategy equilibrium, the first-price sealed-bid auction does not. In the symmetric independent private values model, the equilibrium in a first-price auction can be characterized by an ordinary differential equation (ODE), which has a closed-form solution (Vickrey 1961). However, once this canonical model is extended to allow for asymmetries or multiple objects, the equilibrium problem typically results in systems of non-linear differential equations, for which no general exact solution theory exists. More importantly, the information required to derive a Bayes–Nash equilibrium is often unavailable in real-world settings. Even in a Bayesian framework, players must know the common prior distribution, an assumption that is rarely satisfied in practice. This limitation was famously highlighted in the Wilson critique, which argued that game theory relies too heavily on assumptions of common knowledge (Wilson 1987).

In today's algorithmic markets, such as display advertising auctions or sponsored search, bidders delegate strategic bidding decisions to software agents. These agents adapt their strategies in response to observed prices using learning algorithms, without requiring knowledge of a common prior distribution (Liang et al. 2024, Kumar et al. 2024, Wang et al. 2023, Zhang et al. 2022, Tilli and Espinosa-Leal 2021). Bidders in display ad exchanges can easily face millions to billions of auctions per day, often competing against the same competitors for specific types of impressions. These ad exchanges use mostly first-price auctions (Despotakis et al. 2021). A repeated auction where bidders aim to maximize payoff and employ simple learning algorithms constitutes a stylized model of such algorithmic markets. However, why learning dynamics converge to the BNE of a first-price auction in such a market model is mainly unknown.

1.1. *Learning in games*

The question of learning in games more generally has a rich history. The literature examines what kind of outcome arises as a consequence of a relatively simple process of learning and adaptation (Brown 1951, Foster and Vohra 1997, Fudenberg and Levine 1998, Hart and Mas-Colell 2003, Young 2004, Sorin 2023, Foster and Hart 2023, Sorin 2024). The central idea is that, in repeated interaction, agents who adapt their actions independently may converge to a Nash equilibrium – even without prior knowledge of others' types. These systems of learning agents naturally give rise to dynamical systems (Papadimitriou and Piliouras 2019). Yet it is well established that uncoupled learning dynamics do not necessarily converge to equilibrium (Hart and Mas-Colell 2003).

Numerical analyses of matrix games show that gradient-based algorithms can oscillate, diverge, or even be chaotic (Sanders et al. 2018, Bielawski et al. 2021, Chotibut et al. 2020, Palaiopanos et al. 2017, Vlatakis-Gkaragkounis et al. 2023). Recently, Milionis et al. (2023) proved that there exist games for which all game dynamics fail to converge to a Nash equilibrium. On the other hand, there are also game classes, such as potential games or games with strictly dominated strategies, where certain learning algorithms do converge. Most of the literature is, however, confined to static and complete-information games, not to Bayesian games with continuous types and actions.

1.2. *Learning in Bayesian auction games*

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, 2023b). Equilibrium can be verified ex-post, but the reasons for the convergence of such learning algorithms in auction games have not been well understood so far.

If Bayesian auctions are indeed learnable, this has important implications for both theory and practice. First, it would enable the development of numerical solvers for models that have so far resisted analytical solutions. Second, if learning agents would not even converge to equilibrium in stylized repeated auctions, it would raise serious concerns about the efficiency of display advertising auctions and related applications. In this paper, we study convergence of learning algorithms in Bayesian auction games with continuous types and actions to equilibrium and introduce respective mathematical tools.

1.3. *Convergence analysis in Bayesian games*

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 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 et al. 2002). Interestingly, the link between Nash equilibria and VIs has been explored for traffic games (Patriksson and Rockafellar 2003) or Walrasian equilibrium (Jofré et al. 2007), but not for Bayesian games with continuous type and action space and non-smooth utility functions as is the case in auction theory (Ui 2016). Thus, 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 strict *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 (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 2025). Monotonicity is also central for guaranteed convergence in the literature on learning in games (Ratliff et al. 2013, Chasnov et al. 2019). Apart from this, a global Minty condition is sufficient for extragradient algorithms to converge to equilibrium (Strodiot et al. 2016, Song et al. 2020). This condition is also referred to as the Minty VI or dual VI of the Stampaccia-type VI (Ye 2022). In constant-sum games it is related to the well-known smoothness condition of games (Anagnostides and Sandholm 2023) introduced by (Roughgarden 2015). Not much is known about the convergence of learning algorithms in games beyond these two sufficient conditions.[1] Given the experimental evidence showing convergence of a variety of learning algorithms (Bichler et al. 2021, 2023b), it is thus natural to ask if the monotonicity or the Minty condition hold in auction games, and therefore guarantee convergence for a wide variety of learning algorithms.

Demonstrating the monotonicity of auction games is challenging. Whereas practical equilibrium learning algorithms employ some form of discretization, previous research has shown that in such discretized versions of auction games, the monotonicity condition is

---

[1]Several learning dynamics are known to converge to a Nash equilibrium in potential games (Monderer and Shapley 1996). The existence of a concave potential implies monotonicity (Mertikopoulos and Zhou 2019).

often violated (Bichler et al. 2023a). 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. Thus, 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 a function space, this could explain the convergence of algorithms also in discretized versions of the game, where the condition is violated (Glowinski et al. 1981).

We show that the second-price auction is not monotonous, but it satisfies the Minty-condition globally. However, the simple first-price auction satisfies neither of these known and sufficient conditions globally, and we need to introduce a new approach based on Lyapunov functions to show convergence of gradient-based learning algorithms in the space of piece-wise linear bid functions.

1.4. *Contributions*

Overall, our paper makes three contributions: First, introducing a novel proof based on the Gateaux derivative of the ex-ante utility function, 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).

Second, this proof technique for equilibrium problems in auctions and the resulting operator for the Gateaux derivative allow us to analyze the monotonicity and the Minty conditions. Our findings reveal that the first- and the second-price auctions are neither monotone nor pseudo- or quasi-monotone. Thus, we consider 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. A short version of these results were presented in Bichler et al. (2025). We provide insights into the nature of the Minty-violations and show that in spite of violations of this sufficient condition algorithms converge to equilibrium in simple parametric cases.

Our third and main result is the construction of a Lyapunov function for a piece-wise linear function space showing that the gradient flow leads to the BNE in the first-price auction independent of the starting point and independent of Minty violations. This proves the asymptotic stability of the BNE in the standard first-price auction and thus provides a convergence proof for this central auction format. Finding Lyapunov functions in games is challenging, because constraints such as non-negativity constraints on the slopes of individual pieces need to be considered.

Prior literature on learning in games is silent on Bayesian games with continuous type and action spaces. Our work introduces techniques that can be applied to the analysis of learning behavior in a wide range of game-theoretic models with incomplete information such as they are used for modeling contests or oligopoly competition.

2. PROBLEM SETTING AND VARIATIONAL INEQUALITY FORMULATION

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 a function space. This requires us to work with a set of strategies that exhibit sufficient well-behaved properties. Furthermore, we limit ourselves 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.

2.1. *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 $\boldsymbol{B} := \bigtimes_{i=1}^n B_i$ and $\boldsymbol{\mathcal{X}} := \bigtimes_{i=1}^n \mathcal{X}_i$. The goal of each bidder is to maximize their payoff, i.e., they consider their utility function

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

For this, they search for a strategy $\beta_i : \boldsymbol{\mathcal{X}} \to B_i$. 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, and we define $\boldsymbol{V} := \bigtimes_{i=1}^n V_i$ and $\boldsymbol{\mathcal{B}} := \bigtimes_{i=1}^n \mathcal{B}_i$. The random values X_i of all bidders $i = 1, \dots, n$ are distributed in $\mathcal{X}_i$ according to the atomless F_i. We denote by $U_i : \boldsymbol{V} \to \mathbb{R} : U_i(\boldsymbol{\beta}) := \mathbb{E}_{\boldsymbol{X}}[u_i(\boldsymbol{\beta}(\boldsymbol{X}), \boldsymbol{X})]$ the expected utility of bidder i for given strategies $\boldsymbol{\beta} \in \boldsymbol{\mathcal{B}}$. Here and in the following, we denote a vector of strategies by $\boldsymbol{\beta} := (\beta_1, \beta_2, \dots, \beta_n)$ while $\beta_i, \boldsymbol{\beta}^*_{-i} := (\beta^*_1, \dots, \beta^*_{i-1}, \beta_i, \beta^*_{i+1}, \dots, \beta^*_n)$ denotes a vector of strategies where the i-th one is replaced. The BNE for this auction game is then given by:

Problem 1—BNE: Find $\boldsymbol{\beta}^* = (\beta_1^*, \dots, \beta_n^*) \in \boldsymbol{\mathcal{B}}$ such that for all $i = 1, \dots, n$ there holds

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

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. Let V_i be 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. We emphasize that in the infinite-dimensional case, a linear functional is not necessarily continuous. 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. Let $DU_i(\boldsymbol{\beta})[d]$ denote the directional derivative of U_i at $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n) \in \boldsymbol{\mathcal{B}}$ with respect to β_i along $d \in V_i$, i.e.,

$$DU_i(\boldsymbol{\beta})[d] := \lim_{\varepsilon \to 0} \varepsilon^{-1} \big(U_i(\beta_i + \varepsilon d, \boldsymbol{\beta}_{-i}) - U_i(\boldsymbol{\beta}) \big) \qquad \forall d \in V_i. \tag{2}$$

The directional derivative is the Gateaux-derivative iff $DU_i(\boldsymbol{\beta}) \in V_i^*$, i.e., when the derivative is a continuous linear functional 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 $\boldsymbol{\beta}^* \in \boldsymbol{\mathcal{B}}$ such that for all $i = 1, \dots, n$ there holds

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

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 $\boldsymbol{\beta}^* \in \boldsymbol{\mathcal{B}}$ which satisfies for all $i = 1, \dots, n$

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

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 et al. 2002).

2.2. *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

$$B_i = B, \quad \mathcal{X}_i = \mathcal{X}, \quad X_i \sim_{iid.} F_i \equiv F, \quad u_i \equiv u \qquad \forall i = 1, \ldots, n.$$

Furthermore, private values ensure $\beta_i(\boldsymbol{x}) = \beta_i(x_i)$ for $i = 1, \ldots, n$, 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 $U(\boldsymbol{\beta}) := \mathbb{E}_{\boldsymbol{X}}[u(\boldsymbol{\beta}(\boldsymbol{X}), \boldsymbol{X})]$, and symmetric bids are denoted $\tilde{\boldsymbol{\beta}} := (\tilde{\beta}, \ldots, \tilde{\beta}) \in \boldsymbol{\mathcal{B}}$ 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 strictly increasing and 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, for example, an inner product or a natural dual product. Therefore, consider the Banach space

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

i.e., V 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

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

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. Altogether, this ensures the set $\mathcal{B}_\delta$ to be convex, closed, and bounded in V for any $\delta \geq 0$. Requiring a small positive derivative

($\beta' \geq \delta > 0$) is slightly more restrictive, but necessary in the following analysis to obtain upper bounds for the derivative of the inverse function $(\beta^{-1})' = (\beta' \circ \beta^{-1})^{-1} \leq \delta^{-1}$.

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 (with $\boldsymbol{\beta}^* = (\beta^*, \ldots, \beta^*)$)

$$U(\boldsymbol{\beta}^*) \geq U(\beta, \boldsymbol{\beta}^*_{-1}), \qquad \forall \beta \in \mathcal{B}_\delta. \tag{5}$$

A solution $\beta^* \in \mathcal{B}_\delta$ to the symmetric VI satisfies (with $\boldsymbol{\beta}^* = (\beta^*, \ldots, \beta^*)$)

$$DU(\boldsymbol{\beta}^*)[\beta - \beta^*] \leq 0 \qquad \forall \beta \in \mathcal{B}_\delta. \tag{6}$$

A solution $\beta^* \in \mathcal{B}_\delta$ to the symmetric MVI satisfies (with $\tilde{\boldsymbol{\beta}} = (\tilde{\beta}, \ldots, \tilde{\beta})$)

$$DU(\beta, \tilde{\boldsymbol{\beta}}_{-1})[\beta - \beta^*] \leq 0 \qquad \forall \beta, \tilde{\beta} \in \mathcal{B}_\delta. \tag{7}$$

Here $DU(\boldsymbol{\beta})[d]$ is the Gateaux-derivative defined by (2) of U at $\boldsymbol{\beta} \in \mathcal{B}_\delta$ with respect to β_1 along $d \in V$.

3. MONOTONICITY AND VARIATIONAL STABILITY

In this section, we analyze the symmetric second- and first-price sealed-bid auctions in the continuous setting using the mathematical tools presented above. Starting with the symmetric second-price sealed-bid auction, we first derive the Gateaux-derivative for the bidder's utility function and then use it to prove the existence of a unique BNE which satisfies the VI and MVI. In a second step, we show that the Gateaux-derivative is not (quasi-)monotone by a counterexample for the simple case of two bidders with independent uniform priors. Hence, this property cannot be used to explain the convergence of (certain) gradient-based learning algorithms, and only convergence of extragradient algorithms can be guaranteed due to the existence of the solution to the MVI.

For the symmetric first-price sealed-bid auction, we follow the same steps. Here however, we show the existence of a unique BNE which satisfies the VI, but not the MVI, i.e., there is no MVI solution. Together with the subsequent counterexample for (quasi-)monotonicity, this means that the classical variational stability criteria do not hold here and thus cannot explain the numerically observed convergence.

3.1. *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

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

Since every $\tilde{\beta} \in \mathcal{B}_\delta$ is an increasing function, it satisfies $\max_{j=2,\dots,n} \tilde{\beta}(x_j) = \tilde{\beta}\big(\max_{j=2,\dots,n} x_j\big)$. 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 bids $\tilde{\boldsymbol{\beta}} = (\tilde{\beta},\dots,\tilde{\beta})$ can then be reformulated as

$$\begin{aligned}
U(\beta,\tilde{\boldsymbol{\beta}}_{-1}) &= \int_{[0,1]^n} \chi_{\{\beta(x_1) > \max_{j=2,\dots,n} \tilde{\beta}(x_j)\}}\big(x_1 - \max_{j=2,\dots,n} \tilde{\beta}(x_j)\big)\mathrm{d}F(x_1)\cdots\mathrm{d}F(x_n)\\
&= \int_{[0,1]^n} \chi_{\{\beta(x_1) > \tilde{\beta}(\max_{j=2,\dots,n} x_j)\}}\big(x_1 - \tilde{\beta}(\max_{j=2,\dots,n} x_j)\big)\mathrm{d}F(x_1)\cdots\mathrm{d}F(x_n)\\
&= \int_0^1\int_0^1 \chi_{\{\beta(x)>\tilde{\beta}(y)\}}\big(x-\tilde{\beta}(y)\big)\mathrm{d}G(y)\mathrm{d}F(x).
\end{aligned}$$

This leads to the following expression for the derivative:

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

$$DU(\beta,\tilde{\boldsymbol{\beta}}_{-1})[d] = \int_0^1 d(x)\chi_{\{\beta(x)\in Im(\tilde{\beta})\}}\big(x-\beta(x)\big)\frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))}\mathrm{d}F(x). \tag{8}$$

PROOF: Proof.By definition, we have

$$\begin{aligned}
DU(\beta,\tilde{\boldsymbol{\beta}}_{-1})[d] &= \lim_{\varepsilon\to 0}\varepsilon^{-1}\Big(U(\beta+\varepsilon d,\tilde{\boldsymbol{\beta}}_{-1}) - U(\beta,\tilde{\boldsymbol{\beta}}_{-1})\Big)\\
&= \lim_{\varepsilon\to 0}\int_0^1 \underbrace{\int_0^1 \varepsilon^{-1}\big(\chi_{\{\beta(x)+\varepsilon d(x)>\tilde{\beta}(y)\}} - \chi_{\{\beta(x)>\tilde{\beta}(y)\}}\big)\big(x-\tilde{\beta}(y)\big)\mathrm{d}G(y)}_{I(x):=}\,\mathrm{d}F(x).
\end{aligned}$$

The direction $d \in V \subset AC([0,1])$ is uniformly bounded by $\|d\|_{L^\infty(0,1)} < \infty$, so that the integrand $I(x)$ is bounded F-a.e. as seen by the following calculation:

$$|I(x)| \le |\varepsilon^{-1}|\int_0^1 \chi_{\{\beta(x)+|\varepsilon|\|d\|_{L^\infty}\ge\tilde{\beta}(y)\ge\beta(x)-|\varepsilon|\|d\|_{L^\infty}\}}\mathrm{d}G(y)$$

$$
\begin{aligned}
&= |\varepsilon^{-1}| \int_{\tilde\beta(0)}^{\tilde\beta(1)} \chi_{\{\beta(x)+|\varepsilon|\|d\|_{L^\infty} \geq b \geq \beta(x)-|\varepsilon|\|d\|_{L^\infty}\}} \mathrm{d}(G \circ \tilde\beta^{-1})(b) \\
&\leq |\varepsilon^{-1}| L_{G\circ\tilde\beta^{-1}} \left|(\beta(x) + |\varepsilon|\|d\|_{L^\infty}) - (\beta(x) - |\varepsilon|\|d\|_{L^\infty})\right| \\
&= 2\|d\|_{L^\infty} L_{G\circ\tilde\beta^{-1}} \leq 2\delta^{-1}\|d\|_{L^\infty}\|g\|_{L^\infty(0,1)}.
\end{aligned}
$$

The third step follows since $G \circ \tilde\beta^{-1}$ is Lipschitz-continuous due to $0 \leq (\tilde\beta^{-1})' = (\tilde\beta'(\tilde\beta^{-1}(z)))^{-1} \leq \delta^{-1}$ (F-a.e.). Using the Lipschitz-continuity, we also obtain the F-a.e. point-wise convergence:

$$
\begin{aligned}
I(x) &= \int_0^1 \varepsilon^{-1}\left(\chi_{\{\beta(x)+\varepsilon d(x) > \tilde\beta(y)\}} - \chi_{\{\beta(x) > \tilde\beta(y)\}}\right)\left(x - \tilde\beta(y)\right)\mathrm{d}G(y) \\
&= \chi_{\{\beta(x)<\tilde\beta(1)\}}\varepsilon^{-1} \int_{\beta(x)}^{\beta(x)+\varepsilon d(x)} \left(x - b\right)\mathrm{d}(G \circ \tilde\beta^{-1})(b) \qquad F\text{-a.e.} \\
&\to \chi_{\{\beta(x)<\tilde\beta(1)\}} d(x)\left(x - \beta(x)\right)\frac{g(\tilde\beta^{-1}(\beta(x)))}{\tilde\beta'(\tilde\beta^{-1}(\beta(x)))} \qquad F\text{-a.e.}\,.
\end{aligned}
$$

By the dominated convergence theorem, we can interchange integration and limit to rewrite the original integral as

$$
\begin{aligned}
DU(\beta, \tilde{\boldsymbol\beta}_{-1})[d] &= \int_0^1 \lim_{\varepsilon\to 0} I(x)\mathrm{d}F(x) \\
&= \int_0^1 \chi_{\{\beta(x)<\tilde\beta(1)\}} d(x)\left(x - \beta(x)\right)\frac{g(\tilde\beta^{-1}(\beta(x)))}{\tilde\beta'(\tilde\beta^{-1}(\beta(x)))}\mathrm{d}F(x).
\end{aligned}
$$

This expression is bounded by $|DU(\beta, \tilde{\boldsymbol\beta}_{-1})[d]| \leq 2\delta^{-1}\|g\|_{L^\infty(0,1)}\|d\|_{L^\infty(\mathcal{X},F)}$ and obviously linear in d, so that $DU(\beta, \tilde{\boldsymbol\beta}_{-1}) \in V^*$ for all $\beta, \tilde\beta \in \mathcal{B}_\delta$. ■ *Q.E.D.*

3.1.1. *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 unique solution* $\beta^* = \mathrm{Id}$ *in the compact and convex set* $\mathcal{B}_\delta \subset V$ *for* $0 < \delta \leq 1$.

Proof: Proof.Using the expression (8) for DU, the symmetric VI (6) reads

$$0 \geq DU(\boldsymbol{\beta^*})[\beta - \beta^*] = \int_0^1 \left(\beta(x) - \beta^*(x)\right)\left(x - \beta^*(x)\right)\frac{g(x)}{\beta^{*\prime}(x)}\mathrm{d}F(x) \tag{9}$$

for all $\beta \in \mathcal{B}_\delta$. Obviously, $\beta^* = \mathrm{Id}$ satisfies the VI, and $\mathrm{Id} \in \mathcal{B}_\delta$ for $\delta \leq 1$. We further show that this is the only solution of the VI. Let $\beta^* \in \mathcal{B}_\delta$ be any solution of the VI and consider $\beta = \mathrm{Id}$. Then, (9) becomes

$$\int_{\mathcal{X}} \left|x - \beta^*(x)\right|^2 \frac{g(x)}{\beta^{*\prime}(x)}\mathrm{d}F(x) \leq 0.$$

Since $N = \{x \in \mathcal{X} \mid g(x) = \frac{\mathrm{d}}{\mathrm{d}x}((F(x))^{n-1}) = 0\}$ is a set of measure zero with respect to F, and $\beta^{*\prime} \geq \delta > 0$, this yields $\beta^*(x) = x$ for F-a.e. $x \in \mathcal{X}$.
On the other hand, the MVI (7) for arbitrary $\beta, \tilde{\beta} \in \mathcal{B}_\delta$ reads

$$0 \geq DU(\beta, \tilde{\boldsymbol{\beta}}_{-1})[\beta - \beta^*] = \int_0^1 \chi_{\{\beta(x) < \tilde{\beta}(1)\}}\left(\beta(x) - \beta^*(x)\right)\left(x - \beta(x)\right)\frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))}\mathrm{d}F(x).$$

Obviously, $\beta^* = \mathrm{Id}$ satisfies the MVI. Therefore, $\beta^* = \mathrm{Id}$ is the only strategy satisfying the necessary and sufficient condition for a BNE. ■ *Q.E.D.*

3.1.2. *Monotonicity*

Gradient-based learning for symmetric strategies uses the gradient operator $DU(\beta, \tilde{\boldsymbol{\beta}}_{-1})$ with $\beta = \tilde{\beta}$. Therefore, we are particularly interested whether the operator $DU(\boldsymbol{\beta})$ 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 ($F = \mathrm{Id}$), the operator DU turns out to be neither monotone nor pseudo- nor quasi-monotone.

The operator DU is monotone if it satisfies

$$\left(DU(\tilde{\boldsymbol{\beta}}) - DU(\boldsymbol{\beta})\right)[\tilde{\beta} - \beta] \leq 0 \qquad \forall \beta, \tilde{\beta} \in \mathcal{B}_\delta. \tag{10}$$

For pseudo-monotonicity we require (Khanh 2016)

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

while quasi-monotonicity requires (11) with a strict inequality sign in $DU(\boldsymbol{\beta})[\tilde{\beta}-\beta] < 0$. 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 $\delta < \frac{9}{100}$, $F = \mathrm{Id}$ and $n = 2$.

PROOF: Proof.Plugging (8) into the quasi-monotonicity condition (11) yields

$$\int_0^1 \frac{\big(\tilde{\beta}(x)-\beta(x)\big)\big(x-\beta(x)\big)}{\beta'(x)}\mathrm{d}x < 0 \;\;\Rightarrow\;\; \int_0^1 \frac{\big(\tilde{\beta}(x)-\beta(x)\big)\big(x-\tilde{\beta}(x)\big)}{\tilde{\beta}'(x)}\mathrm{d}x \le 0.$$

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

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

which are in $\mathcal{B}_\delta$ for $\delta \le \frac{9}{100}$ (due to $\tilde{\beta}' \ge \frac{9}{100}$) and yield a negative integral on the left-hand side, but a positive one on the right-hand side. Therefore, DU is not quasi-monotone and consequently neither pseudo-monotone nor monotone. ■ *Q.E.D.*

Note that even though this second-price auction leads to a non-monotone VI, a Minty-type solution exists which was recently shown to be sufficient to ensure convergence for a number of projection-type gradient-based algorithms (Strodiot et al. 2016, Song et al. 2020, Huang and Zhang 2023). In particular, Song et al. (2020) show that when a Minty-type solution exists, optimistic dual extrapolation converges to a solution of the MVI. This implies that at least the rather expensive optimistic dual extrapolation provably finds the BNE of symmetric second-price auctions with independent private values, see the Appendix for details. Moreover, the equivalence of BNE and VI solution implies that a gradient-based learning algorithm must reach the BNE whenever it does converge (within B_δ).

3.2. *First-price sealed-bid auction*

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

$$u(\boldsymbol{b},\boldsymbol{x}) = \chi_{\{b_1 > \max\limits_{j=2,\dots,n} b_j\}}\,(x_1 - b_1)\,.$$

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 bids $\tilde{\boldsymbol{\beta}} = (\tilde{\beta}, \dots, \tilde{\beta})$ can be reformulated as

$$\begin{aligned} U(\beta, \tilde{\boldsymbol{\beta}}_{-1}) &= \int_{[0,1]^n} \chi_{\{\beta(x_1) > \max_{j=2,\dots,n} \tilde{\beta}(x_j)\}} \big(x_1 - \beta(x_1)\big) \mathrm{d}F(x_1) \cdots \mathrm{d}F(x_n) \\ &= \int_0^1 \int_0^1 \chi_{\{\beta(x) > \tilde{\beta}(y)\}} \big(x - \beta(x)\big) \mathrm{d}G(y) \mathrm{d}F(x) \\ &= \int_0^1 \big(x - \beta(x)\big) \int_0^1 \chi_{\{\beta(x) > \tilde{\beta}(y)\}} \mathrm{d}G(y) \mathrm{d}F(x). \end{aligned}$$

The Gateaux-derivative at $(\beta, \tilde{\boldsymbol{\beta}}_{-1}) \in \boldsymbol{\mathcal{B}}_\delta$ along $d \in V$, can be computed as:

$$\begin{aligned} DU(\beta, \tilde{\boldsymbol{\beta}}_{-1})[d] &= \lim_{\varepsilon \to 0} \varepsilon^{-1} \big(U(\beta + \varepsilon d, \tilde{\boldsymbol{\beta}}_{-1}) - U(\beta, \tilde{\boldsymbol{\beta}}_{-1})\big) \\ &= \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \int_0^1 \big(x - \beta(x) - \varepsilon d(x)\big) \int_{\{\beta(x) + \varepsilon d(x) > \tilde{\beta}(y)\}} \mathrm{d}G(y) - \big(x - \beta(x)\big) \int_{\{\beta(x) > \tilde{\beta}(y)\}} \mathrm{d}G(y) \mathrm{d}F(x) \\ &= \lim_{\varepsilon \to 0} \int_0^1 \Bigg[\big(x - \beta(x)\big) \frac{1}{\varepsilon} \int_0^1 \chi_{\{\beta(x) + \varepsilon d(x) > \tilde{\beta}(y)\}} - \chi_{\{\beta(x) > \tilde{\beta}(y)\}} \mathrm{d}G(y) \\ &\qquad - d(x) \int_0^1 \chi_{\{\beta(x) + \varepsilon d(x) > \tilde{\beta}(y)\}} \mathrm{d}G(y) \Bigg] \mathrm{d}F(x). \end{aligned}$$

Analogously to the derivation in the previous section, we obtain for $\tilde{\beta} \in \mathcal{B}_\delta$ that the first term converges to $\chi_{\{\beta(x) < \tilde{\beta}(1)\}} (x - \beta(x)) d(x) (G \circ \tilde{\beta}^{-1})'(\beta(x))$, while the second one yields $\chi_{\{\beta(x) < \tilde{\beta}(1)\}} d(x) G(\tilde{\beta}^{-1}(\beta(x)))$, such that we obtain

$$DU(\beta, \tilde{\boldsymbol{\beta}}_{-1})[d] = \int_0^1 d(x) \chi_{\{\beta(x) < \tilde{\beta}(1)\}} \left[\big(x - \beta(x)\big) \frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))} - G(\tilde{\beta}^{-1}(\beta(x))) \right] \mathrm{d}F(x). \tag{12}$$

Hence, the operator $DU(\beta, \tilde{\boldsymbol{\beta}})$ is linear and in V^*.

3.2.1. *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 $\mathcal{B}_\delta$). Even in the simple case of two bidders ($n = 2$) with uniform priors ($F = \mathrm{Id}$), no solution to the MVI exists. So, in contrast to second-price auctions, these solution notions are different for first-price auctions.

LEMMA 3: *Assume that $f = F'$ satisfies*

$$\delta_0 := \frac{\inf_{x\in[0,1]} f(x)}{2 \sup_{x\in[0,1]} f(x)} > 0.$$

For $0 < \delta \le \delta_0$, $\beta^*(x) = \frac{1}{G(x)} \int_0^x y \mathrm{d}G(y)$ is the unique solution to the symmetric VI (6) in the interior of $\mathcal{B}_\delta$. This solution is the unique symmetric BNE of (5).

PROOF: Proof.The unique symmetric BNE is given by $\beta^*(x) = \frac{1}{G(x)} \int_0^x y \mathrm{d}G(y)$ (Krishna 2009, Chawla and Hartline 2013). Next, we show that this is the unique solution of the symmetric VI (6) in the interior of $\mathcal{B}_\delta$. Using (12), the symmetric VI (6) reads

$$\int_X \left[\big(x - \beta^*(x)\big) \frac{g(x)}{\beta^{*\prime}(x)} - G(x) \right] \big(\beta(x) - \beta^*(x)\big) \mathrm{d}F(x) \le 0 \qquad \forall \beta \in \mathcal{B}_\delta. \tag{13}$$

A solution in the interior of $\mathcal{B}_\delta$ must satisfy

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

This ODE can be rearranged to $\frac{\mathrm{d}}{\mathrm{d}x}(G(x)\beta^*(x)) = xg(x)$. Since F is Lipschitz-continuous, the right-hand side $xg(x) = (n-1)xf(x)(F(x))^{n-2}$ is integrable and depends only on x. Then, the unique solution in the interior of $\mathcal{B}_\delta$ is given by $\beta^*(x) = \frac{1}{G(x)} \int_0^x y \mathrm{d}G(y)$ due to $\beta(0) = 0$ for any $\beta \in \mathcal{B}_\delta$. Note that $\beta^* \in \mathcal{B}_\delta$ since

$$\beta^{*\prime}(x) = \frac{g(x)}{(G(x))^2} \int_0^x G(y) \mathrm{d}y = \frac{(n-1)f(x) \int_0^x (F(y))^{n-1} \mathrm{d}y}{n \int_0^x f(y)(F(y))^{n-1} \mathrm{d}y}$$

$$\ge \frac{\inf_{z\in[0,1]} f(z) \int_0^x (F(y))^{n-1} \mathrm{d}y}{2 \int_0^x \sup_{z\in[0,1]} f(z)(F(y))^{n-1} \mathrm{d}y} = \frac{\inf_{z\in[0,1]} f(z)}{2 \sup_{z\in[0,1]} f(z)} = \delta_0 \ge \delta > 0.$$

Then, $\beta^* \in \mathcal{B}_\delta$ satisfies (13). ■ *Q.E.D.*

REMARK 1: Note that Lemma 3 holds for all $0 < \delta \leq \delta_0$. Hence, a limit argument shows that the BNE is the unique solution to the symmetric VI (13) in the interior of the class of uniformly increasing functions $B_{0+} := \bigcup_{\delta > 0} B_\delta$. 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 $DU(\boldsymbol{\beta})$ 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 4: *In the case of two bidders with uniform priors, i.e., for $n = 2$ and $F = \mathrm{Id}$, the unique BNE $\beta^*(x) = \frac{x}{2}$ according to Lemma 3 does not satisfy the symmetric MVI (7). In particular, the condition is also not satisfied locally for any open neighborhood of the BNE for $\delta \leq \frac{1}{5}$.*

PROOF: Proof.Using (12), the symmetric MVI (7) reads

$$\int_0^1 \left[\frac{x - \beta(x)}{\beta'(x)} - x\right]\big(\beta(x) - \beta^*(x)\big)\mathrm{d}F(x) \leq 0 \qquad \forall \beta \in \mathcal{B}_\delta. \tag{14}$$

Inserting $\beta^*(x) = \frac{x}{2}$ and the continuous, piece-wise linear and strictly increasing bid function

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

for arbitrary $n \in \mathbb{N}_0$ (with $\beta' \geq \frac{1}{5} = \delta$), 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 $\beta \longrightarrow \beta^*$ in V as $n \longrightarrow \infty$. ■ *Q.E.D.*

3.2.2. *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 $\mathcal{B}_\delta$. Again, we show that even in the most simple setting of two bidders ($n = 2$) with uniform priors ($F = \mathrm{Id}$), 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: Proof.For $F(x) = x$ and $n = 2$, and using (12), the quasi-monotonicity condition (11) reads

$$\int_0^1 \left(\tilde{\beta}(x) - \beta(x)\right)\left(\frac{x-\beta(x)}{\beta'(x)} - x\right)\mathrm{d}x < 0 \Rightarrow \quad \int_0^1 \left(\tilde{\beta}(x) - \beta(x)\right)\left(\frac{x-\tilde{\beta}(x)}{\tilde{\beta}'(x)} - x\right)\mathrm{d}x \leq 0. \quad (15)$$

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

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

which are in $\mathcal{B}_\delta$ for $\delta \leq \frac{1}{10}$ (due to $\tilde{\beta}' \geq \frac{1}{10}$) and yield a negative value for the left-hand side of (15), but a positive value for the right-hand side of (15). Therefore, DU is not quasi-monotone and consequently neither pseudo-monotone nor monotone. ■ *Q.E.D.*

4. ASYMPTOTIC STABILITY

We know that the first- and second-price sealed-bid auctions have a unique BNE and observe empirically that first-order learning algorithms converge consistently to this equilibrium. In the previous section, we demonstrated that neither the Minty condition nor various forms of monotonicity suffice to explain the convergence of learning algorithms in first-price sealed-bid auctions. The abstract framework of dynamical systems provides an alternative pathway by specifying a suitable Lyapunov function, and showing asymptotic stability. Unfortunately, identifying a Lyapunov function for gradient dynamics is challenging because we need to consider the projection onto the constrained set of bidding strategies B_δ. Projected dynamical systems of this sort have received relatively little attention (Nagurney and Zhang 1996, Souaiby et al. 2020).

In this section, we establish that a Lyapunov function satisfying additional boundary constraints guarantees convergence of long-term gradient dynamics globally, that is, for every starting position. To this end, we introduce an algorithm to compute candidate Lyapunov functions, demonstrating this approach for piecewise linear strategies with up to four pieces and proving convergence in the case of two pieces. Notably, the counterexamples constructed for Lemma 4 and Proposition 2 are confined to the space of piecewise linear functions. This indicates that the problem's complexity persists even when the strategy space is restricted to such functions. Therefore, we focus our analysis on piecewise linear strategies characterized by m pieces and provide a parameterization for these cases. Our numerical experiments, detailed in Section 4.3, confirm convergence in all configurations tested, including strategies with up to sixteen pieces. Such piecewise linear functions can approximate arbitrary non-linear bid functions arbitrarily well over compact domains.

4.1. *Analytically verifying a Lyapunov function for piece-wise linear strategies*

We can write continuous piecewise linear functions with m pieces in the following way:

$$\beta(x) = \sum_{k=1}^{m} b_k d_k(x) \quad \text{with } d_k(x) = \begin{cases} 0 & \text{for } 0 \le x \le \frac{k-1}{m} \\ x - \frac{k-1}{m} & \text{for } \frac{k-1}{m} < x \le \frac{k}{m} \\ \frac{1}{m} & \text{for } \frac{k}{m} < x \le 1. \end{cases} \tag{16}$$

The parameter b_k corresponds to the slope of the k-th piece. To ensure that the strategies satisfy our assumptions, i.e., are elements of $\mathcal{B}_\delta$, we have to assume that $b_k \ge \delta$ for all $k = 1, \dots, m$ and $\sum_{k=1}^{m} b_k \le m$. The set of all such strategies is given by the corresponding parameter set

$$\mathcal{B}_\delta^m := \left\{ \boldsymbol{b} \in \mathbb{R}^m : g_i(\boldsymbol{b}) \ge 0 \text{ for } i \in \{1, \dots, m+1\} \right\} \subset \mathbb{R}^m, \tag{17}$$

with $g_i(\boldsymbol{b}) = b_i - \delta \ge 0$ for $i \in \{1, \dots, m\}$ and $g_{m+1}(\boldsymbol{b}) = m - \sum_{k=1}^{m} b_k$. The gradient of the expected utility with respect to the parameters $\boldsymbol{b} \in \mathcal{B}_\delta^m$ is the Gateaux derivative of U at the corresponding piecewise linear strategy along d_k (cf. Equation (16)).

Proposition 3: In a first-price sealed bid auction with two players and uniformly distributed (i.i.d.) values, the partial derivative of the expected utility given symmetric piecewise linear strategies $\boldsymbol{b} \in \mathcal{B}_\delta^m$ with respect to the parameter b_i is given by

$$\begin{aligned} \left.\frac{\partial U(\boldsymbol{b}, \tilde{\boldsymbol{b}})}{\partial b_i}\right|_{\tilde{\boldsymbol{b}}=\boldsymbol{b}} &= \frac{1}{6m^3 b_i}\left(-3 \sum\nolimits_{k=1}^{i-1} b_k - (3i+1)b_i + 3i - 1\right) \\ &\quad + \sum_{j=i+1}^{m} \frac{1}{2m^3 b_j}\left(-2\sum\nolimits_{k=1}^{j-1} b_k - 2jb_j + 2j - 1\right). \end{aligned} \tag{18}$$

PROOF: Proof. The Gateaux derivative for the first-price sealed-bid auction is given by (cf. Equation (12))

$$DU(\beta,\tilde{\boldsymbol{\beta}}_{-1})[d] = \int_0^1 d(x)\chi_{\{\beta(x)<\tilde{\beta}(1)\}}\left[\left(x-\beta(x)\right)\frac{g(\tilde{\beta}^{-1}(\beta(x)))}{\tilde{\beta}'(\tilde{\beta}^{-1}(\beta(x)))} - G(\tilde{\beta}^{-1}(\beta(x)))\right]\mathrm{d}F(x).$$

Assuming that we have two players with symmetric strategies ($\beta=\tilde{\beta}$) and a uniform prior ($F(x)=x$), we get $G(x)=x$ and $g(x)=1$. This simplifies the derivative to

$$DU(\beta,\beta_{-1})[d] = \int_0^1 d(x)\left[\left(x-\beta(x)\right)\frac{1}{\beta'(x)} - x\right]\mathrm{d}x.$$

The partial derivative of a given piecewise linear strategy β with respect to a parameter b_i is given by $DU(\beta,\beta_{-1})[d_i]$. Using the definitions of β and d_i (cf. Equation (16)), we get

$$DU(\beta,\beta_{-1})[d_i] = \sum_{j=i}^{m}\int_{\frac{j-1}{m}}^{\frac{j}{m}} d_i(x)\left(\frac{x-\beta(x)}{\beta'(x)} - x\right)\mathrm{d}x.$$

By definition of d_i, the first $i-1$ terms are zero, and we find

$$= \int_{\frac{i-1}{m}}^{\frac{i}{m}} (x-\tfrac{i-1}{m})\left(\frac{x-\beta(x)}{\beta'(x)} - x\right)\mathrm{d}x + \sum_{j=i+1}^{m}\int_{\frac{j-1}{m}}^{\frac{j}{m}} \frac{1}{m}\left(\frac{x-\beta(x)}{\beta'(x)} - x\right)dx.$$

By definition of β, we have $\beta(x) = b_j(x-\frac{j-1}{m}) + \frac{1}{m}b_{<j}$ with $b_{<j} := \sum_{k=1}^{j-1} b_k$ and $\beta'(x) = b_j$ for values in the interval $[\frac{j-1}{m},\frac{j}{m}]$ with $j=1,\dots,m$. This gives us

$$\begin{aligned}
&= \frac{1}{b_i}\int_{\frac{i-1}{m}}^{\frac{i}{m}} (x-\tfrac{i-1}{m})(x-b_i(x-\tfrac{i-1}{m}) - \tfrac{1}{m}b_{<i} - b_i x)\mathrm{d}x \\
&\quad + \sum_{j=i+1}^{m}\frac{1}{mb_j}\int_{\frac{j-1}{m}}^{\frac{j}{m}} x - b_i(x-\tfrac{i-1}{m}) - \tfrac{1}{m}b_{<i} - b_i x\mathrm{d}x \\
&= \frac{1}{b_i}\int_{\frac{i-1}{m}}^{\frac{i}{m}} (1-2b_i)(x-\tfrac{i-1}{m})^2 + (x-\tfrac{i-1}{m})(\tfrac{i-1}{m}(1-b_i) - b_{<i})\mathrm{d}x \\
&\quad + \sum_{j=i+1}^{m}\frac{1}{mb_j}\int_{\frac{j-1}{m}}^{\frac{j}{m}} (1-2b_j)(x-\tfrac{j-1}{m}) + \tfrac{j-1}{m}(1+b_j) - b_{<j}\mathrm{d}x \\
&= \frac{1}{6m^3 b_i}\left(-3\sum\nolimits_{k=1}^{i-1} b_k - (3i+1)b_i + 3i - 1\right)
\end{aligned}$$

$$+\sum_{j=i+1}^{m}\frac{1}{2m^3 b_j}\left(-3\sum_{k=1}^{j-1} b_k - 2jb_j + 2j - 1\right).$$

■ *Q.E.D.*

To simplify notation, we write $\nabla U(\boldsymbol{b})$ for the gradient consisting of the partial derivatives defined in Equation (18).

EXAMPLE: In the two-dimensional case, i.e., piecewise linear strategies with two pieces, the projected gradient dynamics for $\boldsymbol{b} \in \mathcal{B}_\delta^2$ and some $\delta > 0$ are given by

$$\nabla U(\boldsymbol{b}) = \begin{pmatrix} -\frac{2b_1+3}{16b_2} + \frac{1}{24b_1} - \frac{1}{3} \\ \frac{-3b_1-7b_2+5}{48b_2} \end{pmatrix}. \tag{19}$$

We study the following dynamical system

$$\dot{\boldsymbol{b}} = \Pi_{\mathcal{T}C_{\mathcal{B}_\delta^m}(\boldsymbol{b})}\left(\nabla U(\boldsymbol{b})\right), \tag{20}$$

where $\dot{\boldsymbol{b}}$ denotes the time derivative and $\Pi_{\mathcal{T}C_{\mathcal{B}_\delta^m}(\boldsymbol{b})}$ the projection onto the tangent cone of $\mathcal{B}_\delta^m$ at $\boldsymbol{b}$. We denote the trajectory at time t for an initial starting point $\boldsymbol{b}^o \in \mathcal{B}_\delta^m$ by $\boldsymbol{b}(t, \boldsymbol{b}^o)$ and have $\boldsymbol{b}(0, \boldsymbol{b}^o) = \boldsymbol{b}^o$.

To show global asymptotic stability of the system described by Equation (20) in the equilibrium point $\boldsymbol{b}^* = \left(\frac{1}{2}, \ldots, \frac{1}{2}\right)^T$, we leverage the formalism introduced by Souaiby et al. (2020) for projected dynamical systems. Recall that the system is stable in $\boldsymbol{b}^*$ if for every $\varepsilon > 0$ there exists $\kappa > 0$ such that for $\boldsymbol{b}^o \in \mathcal{B}_\delta^m$ with $\|\boldsymbol{b}^* - \boldsymbol{b}^o\| \leq \kappa$, we have $\|\boldsymbol{b}(t, \boldsymbol{b}^o) - \boldsymbol{b}^*\| \leq \varepsilon$ for all $t \geq 0$. The point $\boldsymbol{b}^*$ is further globally asymptotically stable if it is stable and for all $\boldsymbol{b}^o \in \mathcal{B}_\delta^m$, we have $\lim_{t\to\infty} \|\boldsymbol{b}(t, \boldsymbol{b}^o) - \boldsymbol{b}^*\| = 0$.

A function h qualifies as a Lyapunov function for the constrained system if it satisfies the following condition:

DEFINITION 1—Constrained Lyapunov function (Souaiby et al. 2020): The system given by Equation (20) has a continuously differentiable global Lyapunov function $h : \mathbb{R}^m \to \mathbb{R}$ with respect to $\mathcal{B}_\delta^m$ and $\boldsymbol{b}^*$ if there exist class $\mathcal{K}$ functions[2] $\overline{\alpha}, \underline{\alpha}$, which are additionally unbounded, such that

[2] We say that a function $\alpha : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ is of class $\mathcal{K}$ if it is continuous, strictly increasing, and $\alpha(0) = 0$.

- $\underline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right) \le h(\boldsymbol{b}) \le \overline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right)$, for all $\boldsymbol{b} \in \mathcal{B}_\delta^m$;
- $\langle \nabla h(\boldsymbol{b}), \Pi_{\mathcal{TC}_{\mathcal{B}_\delta^m}(\boldsymbol{b})} \nabla U(\boldsymbol{b})\rangle \le -w\,\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$ with $w > 0$, for all $\boldsymbol{b} \in \mathcal{B}_\delta^m$.

Note that in the original definition of Souaiby et al. (2020), the inner product in the second inequality is bounded by some general class $\mathcal{K}$ function α. For simplicity, we focus on the specific function of the form $\alpha(\|\boldsymbol{b}-\boldsymbol{b}^*\|) = w\,\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$ with $w > 0$. The projection onto the tagent cone complicates the application of the results of standard dynamical systems to demonstrate convergence, necessitating specific additional assumptions. If the Lyapunov function's gradient also points inward from the boundary, we have the following statement, which is a special case of Proposition 1 by Souaiby et al. (2020).

Proposition 4—(Souaiby et al. 2020): Consider the system defined by Equation (20). Assume that there exists a continuously differentiable function h that satisfies the following conditions:

(i) $h(\boldsymbol{b}^*) = 0$, and $\underline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right) \le h(\boldsymbol{b}) \le \overline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right)$ for every $\boldsymbol{b} \in \mathcal{B}_\delta^m$, and some $\underline{\alpha}, \overline{\alpha} \in \mathcal{K}$.

(ii) $\langle \nabla h(\boldsymbol{b}), \nabla U(\boldsymbol{b})\rangle \le -w\,\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$ for some $w > 0$.

(iii) If $\boldsymbol{b}$ is such that $g_i(\boldsymbol{b}) = 0$, for some $i \in \{1, \dots, M\}$, then $\langle h(\boldsymbol{b}), \nabla g_i(\boldsymbol{b})\rangle \le 0$.

Then h is a Lyapunov function for system (20) and $\boldsymbol{b}^*$ is globally asymptotically stable.

For $m = 2$ pieces, we construct a Lyapunov function that satisfies the required conditions, thereby proving stability in this setting.

Theorem 5: *The system $\dot{\boldsymbol{b}} = \Pi_{\mathcal{TC}_{\mathcal{B}_\delta^m}(\boldsymbol{b})}\left(\nabla U(\boldsymbol{b})\right)$ with $m = 2$ and $0 < \delta \le \frac{1}{2}$ has a globally asymptotically stable equilibrium $\boldsymbol{b}^* = (\frac{1}{2}, \frac{1}{2})$. The Lyapunov function for the system is given by*

$$h(b_1, b_2) := \frac{1}{2}\begin{pmatrix} b_1 - \frac{1}{2} \\ b_2 - \frac{1}{2}\end{pmatrix}^T \begin{pmatrix} 20 & 0 \\ 0 & 52\end{pmatrix} \begin{pmatrix} b_1 - \frac{1}{2} \\ b_2 - \frac{1}{2}\end{pmatrix}. \tag{21}$$

Figure 1 illustrates the gradient field for two pieces alongside the candidate's level sets. The red regions highlight where the Minty condition fails. While the gradient flows take a detour, all trajectories converge to the Bayes–Nash equilibrium. This demonstrates the advantage of Lyapunov functions over the Minty condition in capturing these dynamics, providing a robust proof of convergence.

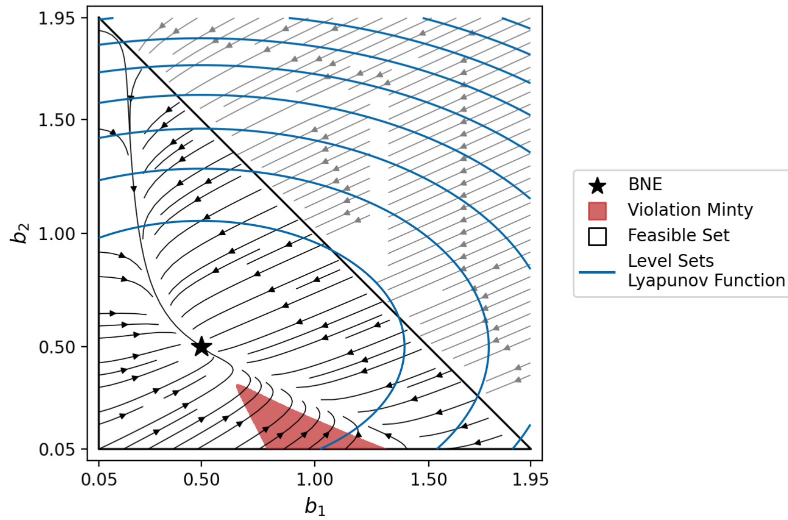

FIGURE 1.—Gradient field of the first-price sealed-bid auction with $n = 2$ bidders, $m = 2$ pieces and $\delta = \frac{1}{20}$.

Before we show that the function h from Theorem 5 is indeed a Lyapunov function for the system, we want to briefly describe how one can find a suitable candidate.

Assume that there exists a Lyapunov function of the form $h(\boldsymbol{b}) = \frac{1}{2}(\boldsymbol{b} - \boldsymbol{b}^*)^T H(\boldsymbol{b} - \boldsymbol{b}^*)$ with a matrix $H \in \mathbb{R}^{m \times m}$. To find a suitable candidate, we solve an LP with the parameters of H as variables and the conditions (ii) and (iii) of Proposition 4 at finitely many points as constraints. To that end, we discretize the feasible set $\mathcal{B}_\delta^m$ by laying a uniform grid over the space and selecting all points that satisfy the constraints. In addition to these interior points, we also explicitly include additional points located on the boundary of the feasible set to ensure that stability conditions are appropriately captured near the edges. The set of these discrete points is denoted by B. This discretization allows us to impose the conditions as linear constraints in a finite-dimensional linear program for finding a suitable matrix H:

$$
\begin{aligned}
& \max \gamma \\
\text{s.t.} \quad & \langle H(\boldsymbol{b} - \boldsymbol{b}^*), \nabla U(\boldsymbol{b}) \rangle \leq -w \|\boldsymbol{b} - \boldsymbol{b}^*\|_2^2 - \gamma \qquad && \forall b \in B \\
\forall i \in \{1, \dots, m+1\} \quad & \langle H(\boldsymbol{b} - \boldsymbol{b}^*), \nabla g_i(b) \leq 0 && \forall b \in B : g_i(b) = 0 \\
& \gamma \in \mathbb{R},\ H \in \mathbb{R}^{m \times m}.
\end{aligned}
$$

The parameter $w > 0$ is some strictly positive constant, for example, $w = \frac{1}{10}$. If we find a solution such that $\gamma = 0$, we have a candidate that satisfies the conditions at the discrete points. We added the additional constraint $H \in \mathbb{Z}^{m\times m}$ which allows us to prove the properties analytically as we will do in the following for two pieces. This avoids numerical issues due to the fact that fractional values need to be rounded.

PROOF: Proof of Theorem 5. To show that $\boldsymbol{b}^*$ is globally asymptotically stable and that h is a Lyapunov function for system (20) with $m = 2$ and $\delta = \frac{1}{20}$, we have to check the conditions of Proposition 4:

(i) For the first condition, we need to find functions $\underline{\alpha}$, and $\overline{\alpha}$ such that

$$h(\boldsymbol{b}^*) = 0, \text{ and } \underline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right) \leq h(\boldsymbol{b}) \leq \overline{\alpha}\left(\|\boldsymbol{b}-\boldsymbol{b}^*\|\right) \text{ for every } \boldsymbol{b} \in \mathcal{B}_\delta^m \tag{22}$$

Since h is given by $h = 10(b_1 - b_1^*)^2 + 26(b_2 - b_2^*)^2$ we can define the functions $\underline{\alpha}$ and $\overline{\alpha}$ with $\underline{\alpha}(\|\boldsymbol{b}-\boldsymbol{b}^*\|) := 10\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$ and $\overline{\alpha}(\|\boldsymbol{b}-\boldsymbol{b}^*\|) := 26\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$. The inequality obviously holds and the squared distance is of class $\mathcal{K}$.

(ii) $\langle \nabla h(\boldsymbol{b}), \nabla U(\boldsymbol{b})\rangle \leq -w\,\|\boldsymbol{b}-\boldsymbol{b}^*\|^2$ for some $w > 0$.

To show this inequality, we first rewrite the inner product in the following way.

$$\begin{aligned}
&\langle \nabla h(\boldsymbol{b}), \nabla U(\boldsymbol{b})\rangle \\
&= (20b_1 - 10)\cdot\left(-\frac{b_1}{8b_2} + \frac{3}{16b_2} + \frac{1}{24b_1} - \frac{1}{3}\right) + (52b_2 - 26)\cdot\left(\frac{-3b_1 - 7b_2 + 5}{48b_2}\right) \\
&= \frac{1}{b_1b_2}\left(-\frac{5b_1^3}{2} - \frac{119b_1^2b_2}{12} + \frac{53b_1^2}{8} - \frac{91b_1b_2^2}{12} + \frac{107b_1b_2}{8} - \frac{55b_1}{12} - \frac{5b_2}{12}\right) \\
&= \frac{1}{b_1b_2}\left(g_1(\boldsymbol{b})\cdot\sigma_1(\boldsymbol{b}) + g_2(\boldsymbol{b})\cdot\sigma_2(\boldsymbol{b}) + g_3(\boldsymbol{b})\cdot\sigma_3(\boldsymbol{b})\right)
\end{aligned}$$

where g_j are the constraints of the feasible set as defined in Equation (17), and the σ_j are polynomials.[3] The polynomials are given by $\sigma_j(\boldsymbol{b}) = \frac{1}{2}(\boldsymbol{b}-\boldsymbol{b}^*)^T\Sigma_j(\boldsymbol{b}-\boldsymbol{b}^*)$ for

[3]Such a decomposition is known to exist when the polynomial is positive (or negative in our case) (Putinar 1993), but the degree of the polynomials σ_i remains unknown a priori. While semidefinite programming can numerically verify whether such a decomposition exists for a given degree, we found the exact parameters through systematic trial-and-error.

$j = 1, 2, 3$ with

$$\Sigma_1 = \begin{pmatrix} -\frac{11}{2} & -\frac{955}{108} \\ -\frac{955}{108} & -\frac{15463}{1026} \end{pmatrix}, \quad \Sigma_2 = \begin{pmatrix} -\frac{143}{54} & -\frac{35}{108} \\ -\frac{35}{108} & -\frac{21}{38} \end{pmatrix}, \quad \Sigma_3 = \begin{pmatrix} -\frac{1}{2} & 0 \\ 0 & -\frac{21}{38} \end{pmatrix}.$$

It is easy to verify that these quadratic polynomials are strictly concave with unique maximizers at $\boldsymbol{b} = \boldsymbol{b}^*$ with $\sigma_j(b^*) = 0$ for $j = 1, 2, 3$.

Let $\lambda^* < 0$ be the largest eigenvalue of Σ_j with $j \in \{1, 2, 3\}$ (note that all eigenvalues are negative). Then we have $\sigma_j(\boldsymbol{b}) \leq \lambda^* \|b - b^*\|_2^2 \leq 0$ for all $j = 1, 2, 3$ and we can write

$$\begin{aligned} \langle \nabla h(\boldsymbol{b}), \nabla U(\boldsymbol{b}) \rangle &\leq \frac{\lambda^*}{b_1 b_2} \cdot \|\boldsymbol{b} - \boldsymbol{b}^*\|_2^2 \cdot (g_1(\boldsymbol{b}) + g_2(\boldsymbol{b}) + g_3(\boldsymbol{b})) \\ &\leq \frac{\lambda^*}{4} \cdot \|\boldsymbol{b} - \boldsymbol{b}^*\|_2^2 \cdot (2 - 2\delta) = -w \|\boldsymbol{b} - \boldsymbol{b}^*\|_2^2., \end{aligned}$$

where we set $w := -\frac{\lambda^*}{4}(2 - 2\delta)$. Note that $\lambda^* = -\frac{10553}{1026} + 5\frac{\sqrt{17026937}}{2052} \leq -0.231$, which guarantees $w > 0$ for all $0 < \delta \leq \frac{1}{2}$.

(iii) If $g_i(\boldsymbol{b}) = 0$, for some $i \in \{1, \ldots, M\}$, then $\langle \nabla h(\boldsymbol{b}), \nabla g_i(\boldsymbol{b}) \rangle \leq 0$. We show this case by case:

- If $g_1(\boldsymbol{b}) = 0$, then $b_1 = \delta$ and $\langle \nabla h(\delta, b_2), \nabla g_1(\delta, b_2) \rangle = 20(\delta - \frac{1}{2}) \leq 0$.
- If $g_2(\boldsymbol{b}) = 0$, then $b_2 = \delta$ and $\langle \nabla h(b_1, \delta), \nabla g_2(b_1, \delta) \rangle = 52(\delta - \frac{1}{2}) \leq 0$.
- If $g_3(\boldsymbol{b}) = 0$, then $b_2 = 2 - b_1$ and $\langle \nabla h(b_1, 2 - b_1), \nabla g_3(b_1, 2 - b_1) \rangle = 32b_1 - 68 \leq 0$ for all $b_1 \leq 2$.

Therefore, the conditions of Proposition 4 are satisfied. The function h is a Lyapunov function for (20) with $m = 2$ and $0 < \delta \leq \frac{1}{2}$, and b^* is globally asymptotically stable. ■ *Q.E.D.*

4.2. *Numerical Lyapunov function construction via sum-of-squares decomposition*

The approach described in Section 4.1 to prove convergence for two pieces, also works for more pieces. However, proving the required conditions analytically becomes very laborious. Instead, we use the method recently proposed by Souaiby et al. (2020, Algorithm 2) to numerically compute Lyapunov functions. This method relies on sum-of-squares (SOS) decompositions of polynomials to verify positivity. It can be used to find a polynomial Lyapunov function such that conditions (ii) and (iii) of Proposition 4 are SOS, i.e., they can be expressed as sums of squared polynomials. The resulting problem can be formulated and solved using semidefinite programming (SDP). The existence of such SOS decompositions of positive polynomials over the feasible set is ensured by Putinar's Positivstellensatz (Putinar 1993). However, the required polynomial degree of the functions used in the SOS decomposition is not known a priori. Therefore, we iteratively increase the degree and solve the corresponding SDP to check whether such a decomposition exists, assuming a suitable Lyapunov function exists.

We slightly adapt the original algorithm to fit the structure of our problem. To avoid complications from rational expressions in the vector field, we multiply it by the product of all variables, ensuring that the resulting conditions involve only polynomials. This does not change the condition (ii), as the variables are bounded and strictly positive in the original problem. Additionally, to avoid numerical issues, we do not only compute a SOS decomposition of the inner product in (ii) (cf. Souaiby et al. (2020)), but of the inner product minus a scaled ($w = \frac{1}{10}$) squared distance to the equilibrium instead. This way, we prevent that almost constant Lyapunov functions are approximate solutions of the SDP. The problem, adapted from Souaiby et al. (2020), with an SOS lower bound, i.e., we minimize γ such that some objective $f(x) - \gamma$ is still positive, i.e., admits a SOS decomposition over the feasible set, is given by

$$
\begin{aligned}
&\min_{\gamma \in \mathbb{R}} \gamma \quad \text{s.t.} \\
&h(\boldsymbol{b}) - \gamma - \sum_{i=1}^{m+1} \sigma_i(\boldsymbol{b}) g_i(\boldsymbol{b}) \in \text{SOS} \\
&-\langle \nabla h(\boldsymbol{b}), \Pi_{i=1}^{m} b_i \cdot \nabla U(\boldsymbol{b}) \rangle - w \|\boldsymbol{b} - \boldsymbol{b}^*\|_2^2 - \gamma - \sum_{i=1}^{m+1} \sigma_{0,i}(\boldsymbol{b}) g_i(\boldsymbol{b}) \in \text{SOS} \\
\forall j = 1, \ldots, m+1: \quad &\langle \nabla h(\boldsymbol{b}), \nabla g_j(\boldsymbol{b}) \rangle - \gamma - \sum_{i=1, i \neq j}^{m+1} \sigma_{j,i}(\boldsymbol{b}) g_i(\boldsymbol{b}) - \gamma_j g_j(\boldsymbol{b}) \in \text{SOS} \\
&\forall j = 0, \ldots, m+1,\ i = 1, \ldots, m+1: \quad \sigma_i, \sigma_{j,i} \in \text{SOS}
\end{aligned}
$$

where all $\sigma_i, \sigma_{i,j}$, and γ_j are polynomials of degree d and h is a polynomial of degree d_L.

We implement the method in MATLAB using YALMIP with its SOS module (Löfberg 2009), together with the conic solver MOSEK. The code is available upon request. The results reported in Table 4.2 show that we can find quadratic Lyapunov functions for $m = 2, 3$ pieces and a polynomial Lyapunov function of degree 4 for $m = 4$. The degrees of the polynomial used to find the SOS decompositions grow with the dimension of the underlying problem.

Numerical results for computing lyapunov functions.

# Pieces m	Deg. Lyap.	Deg. Poly.	Status	Objective	Iterations	Runtime
2	2	2	OPTIMAL	6.8e-09	10	0.01s
3	2	4	OPTIMAL	-1.4e-10	21	0.11s
4	4	8	OPTIMAL	-4.0e-05	21	14.38s

TABLE I

THE SDP IS FORMULATED FOR DIFFERENT NUMBER OF PIECES m, MAXIMAL DEGREES FOR THE LYAPUNOV FUNCTION, AND MAXIMAL DEGREES FOR THE POLYNOMIALS USED FOR THE SOS DECOMPOSITION. THE FEASIBLE SET IS GIVEN BY $\mathcal{B}_\delta^m$ WITH $\delta = \frac{1}{20}$. WE ONLY REPORT RESULTS, WHERE THE SDP IS SOLVED WITH A SUFFICIENTLY LOW OBJECTIVE VALUE, INDICATING THAT A LYAPUNOV FUNCTION HAS BEEN FOUND. ADDITIONALLY TO THE OBJECTIVE, WE REPORT THE NUMBER OF ITERATIONS, STATUS, AND RUNTIME OF THE MOSEK SOLVER.

4.3. *Stability analysis: results for higher dimensions*

The theoretical results presented above provide a priori guarantees that projected gradient dynamics converge toward equilibrium in first-price auctions with up to three pieces. In this section, we extend the analysis empirically by presenting experiments that demonstrate convergence behavior for instances with up to sixteen pieces.

For each experimental run, we sample an initial strategy uniformly at random over the feasible set $\mathcal{B}_\delta^m$. This is achieved by normalizing i.i.d. samples drawn from a suitable exponential distribution, following the method described by Devroye (1986, page 208).

We run a projected gradient ascent algorithm to optimize the utility. The learning rate at time step t is set to $\eta_t = \frac{1}{t^{0.05}}$. To ensure comparability across different dimensions, we normalize the distance to equilibrium by a factor of $\frac{1}{\sqrt{m}}$. Each experiment runs for a maximum of 100,000 steps or until the norm of the gradient falls below 10^{-10}, whichever occurs first. We run 100 trials for each configuration and report the mean and standard deviation.

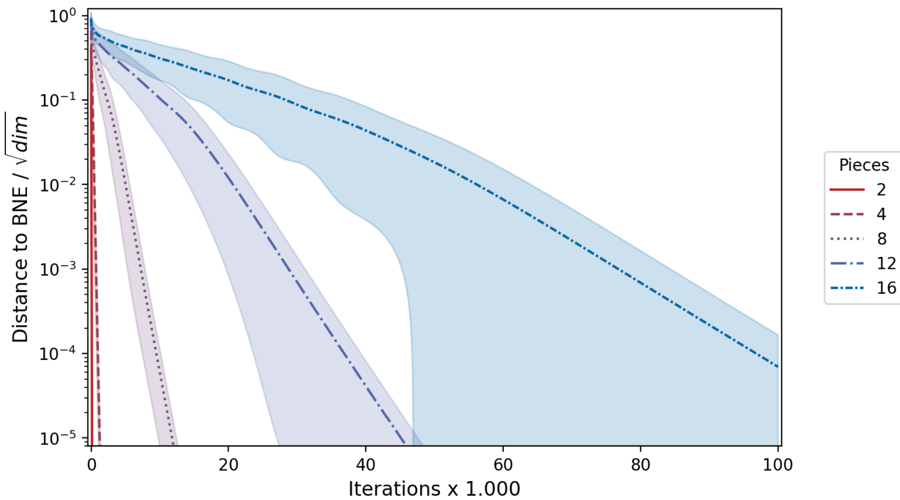

FIGURE 2.—Distance to BNE for gradient ascent algorithm for different number of pieces. The lines display the mean, whereas the shaded areas display the standard deviations over 100 different starting points.

Figure 2 summarizes the empirical results for up to sixteen pieces. We observe that the distance to equilibrium consistently decreases, although the rate of convergence is poor in higher dimensions. This aligns with the observation made in the two-dimensional case, where the gradient flow reaches a flat region quickly so that convergence towards equilibrium is slow (see Figure 1). These findings provide additional empirical support for the hypothesis that simple gradient-based algorithms can learn equilibrium strategies in first-price auctions, even in high-dimensional settings. They also reinforce our conjecture that theoretical guarantees may extend to these more complex cases.

5. 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 sufficient condition for first-order optimization methods to converge to the unique solution of a monotone 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 inequalities uses the Minty condition to show the convergence of extragradient algorithms (cf. Appendix A). In the first-price auction, this condition is also not satisfied, even when assuming a simple uniform prior. These findings highlight the need to go beyond traditional conditions, such as monotonicity or the Minty condition, to understand the convergence behavior of learning algorithms. This requires finding a Lyapunov function for the corresponding dynamical system.

Finding Lyapunov functions in games is challenging in general, but it is particularly difficult in games because constraints on the action space need to be considered. Our main result is the construction of such a Lyapunov function for piece-wise linear bid functions showing that the gradient flows lead to the BNE in the first-price auction independent of the initial condition. This proves the asymptotic stability of the BNE and thus provides a convergence proof for this central auction format and this generic set of bid functions. Establishing the convergence of learning algorithms to equilibrium in other auction games would offer a strong justification for using equilibrium as a predictive concept – one that does not depend on agents having prior information and unbounded rationality.

REFERENCES

Adil, Deeksha, Brian Bullins, Arun Jambulapati, and Sushant Sachdeva, "Optimal Methods for Higher-Order Smooth Monotone Variational Inequalities.," *arXiv preprint arXiv:2205.06167*, 2022. [4]

Anagnostides, Ioannis and Tuomas Sandholm, "On the Interplay between Social Welfare and Tractability of Equilibria," *arXiv preprint arXiv:2310.16976*, 2023. [4]

Bauschke, Heinz H. and Patrick L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces* CMS Books in Mathematics, Cham: Springer International Publishing, 2017. [4]

Bichler, M., M. Fichtl, S. Heidekrueger, N. Kohring, and P. Sutterer, "Learning equilibria in symmetric auction games using artificial neural networks," *Nature Machine Intelligence*, 2021, pp. 687–695. [3, 4]

Bichler, Martin, Maximilian Fichtl, and Matthias Oberlechner, "Computing Bayes-Nash equilibrium in auction games via gradient dynamics," in "ACM Conference on Economics and Computation" 2023. [5]

——— **, Nils Kohring, and Stefan Heidekrüger**, "Learning Equilibria in Asymmetric Auction Games," *INFORMS Journal on Computing*, 2023, *35* (3), 523–542. [3, 4]

——— **, Stephan B. Lunowa, Matthias Oberlechner, Fabian R. Pieroth, and Barbara Wohlmuth**, "Beyond Monotonicity: On the Convergence of Learning Algorithms in Standard Auction Games," *Proceedings of the AAAI Conference on Artificial Intelligence*, Apr. 2025, *39* (13), 13649–13657. [5]

Bielawski, Jakub, Thiparat Chotibut, Fryderyk Falniowski, Grzegorz Kosiorowski, Michał Misiurewicz, and Georgios Piliouras, "Follow-the-regularized-leader routes to chaos in routing games," in "International Conference on Machine Learning" PMLR 2021, pp. 925–935. [3]

Brown, George W, "Iterative solution of games by fictitious play," *Activity analysis of production and allocation*, 1951, *13* (1), 374–376. [2]

Cavazzuti, E., M. Pappalardo, and M. Passacantando, "Nash Equilibria, Variational Inequalities, and Dynamical Systems," *Journal of Optimization Theory and Applications*, September 2002, *114* (3), 491–506. [4, 8]

Chasnov, Benjamin, Lillian J Ratliff, Eric Mazumdar, and Samuel A Burden, "Convergence Analysis of Gradient-Based Learning with Non-Uniform Learning Rates in Non-Cooperative Multi-Agent Settings," *arXiv preprint arXiv:1906.00731*, 2019. [4]

Chawla, Shuchi and Jason D Hartline, "Auctions with unique equilibria," in "Proceedings of the fourteenth ACM conference on Electronic commerce" 2013, pp. 181–196. [15]

Chotibut, Thiparat, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras, "The route to chaos in routing games: When is price of anarchy too optimistic?," *Advances in Neural Information Processing Systems*, 2020, *33*, 766–777. [3]

Despotakis, Stylianos, Ramamoorthi Ravi, and Amin Sayedi, "First-price auctions in online display advertising," *Journal of Marketing Research*, 2021, *58* (5), 888–907. [2]

Devroye, Luc, *Non-Uniform Random Variate Generation*, New York, NY: Springer, 1986. [26]

Foster, Dean P and Rakesh V Vohra, "Calibrated learning and correlated equilibrium," *Games and Economic Behavior*, 1997, *21* (1-2), 40. [2]

——— **and Sergiu Hart**, ""Calibeating": Beating forecasters at their own game," *Theoretical Economics*, 2023, *18* (4), 1441–1474. [2]

Fudenberg, Drew and David K Levine, *The Theory of Learning in Games*, Vol. 2, MIT press, 1998. [2]

Glowinski, Roland, Jacques Louis Lions, and Raymond Trémolières, *Numerical Analysis of Variational Inequalities*, Vol. 8 of *Studies in mathematics and its applications*, Elsevier, 1981. [5]

Hart, Sergiu and Andreu Mas-Colell, "Uncoupled dynamics do not lead to Nash equilibrium," *American Economic Review*, 2003, *93* (5), 1830–1836. [2, 3]

Huang, Kevin and Shuzhong Zhang, "Beyond Monotone Variational Inequalities: Solution Methods and Iteration Complexities," *arXiv preprint arXiv:2304.04153*, 2023. [13]

Jofré, Alejandro, R Terry Rockafellar, and Roger JB Wets, "Variational inequalities and economic equilibrium," *Mathematics of Operations Research*, 2007, *32* (1), 32–50. [4]

Khanh, Pham Duy, "A Modified Extragradient Method for Infinite-Dimensional Variational Inequalities," *Acta Mathematica Vietnamica*, 2016, *41* (2), 251–263. [12]

Kinderlehrer, D. and G. Stampacchia, *An Introduction to Variational Inequalities and Their Applications*, SIAM, 2000. [7]

Krishna, Vijay, *Auction Theory*, Academic Press, 2009. [5, 15]

Kumar, Rachitesh, Jon Schneider, and Balasubramanian Sivan, "Strategically-Robust Learning Algorithms for Bidding in First-Price Auctions," *arXiv preprint arXiv:2402.07363*, 2024. [2]

Liang, Jason Cheuk Nam, Haihao Lu, and Baoyu Zhou, "Online ad procurement in non-stationary autobidding worlds," *Advances in Neural Information Processing Systems*, 2024, *36.* [2]

Lin, Tianyi and Michael I Jordan, "Perseus: A Simple and Optimal High-Order Method for Variational Inequalities," *Mathematical Programming*, 2025, *209*, 609–650. [4]

Lions, J. L. and G. Stampacchia, "Variational Inequalities," *Communications on Pure and Applied Mathematics*, 1967, *20* (3), 493–519. [7]

Löfberg, Johan, "Pre- and post-processing sum-of-squares programs in practice," *IEEE Transactions on Automatic Control*, 2009, *54* (5), 1007–1011. [25]

Mertikopoulos, Panayotis and Zhengyuan Zhou, "Learning in games with continuous action sets and unknown payoff functions," *Mathematical Programming*, 2019, *173* (1-2), 465–507. [4]

Migot, Tangi and Monica-G. Cojocaru, "A parametrized variational inequality approach to track the solution set of a generalized Nash equilibrium problem," *European Journal of Operational Research*, 2020, *283* (3), 1136–1147. [4]

Milionis, Jason, Christos Papadimitriou, Georgios Piliouras, and Kelly Spendlove, "An impossibility theorem in game dynamics," *Proceedings of the National Academy of Sciences*, 2023, *120* (41), e2305349120. [3]

Monderer, Dov and Lloyd S Shapley, "Potential games," *Games and economic behavior*, 1996, *14* (1), 124–143. [4]

Nagurney, Anna and Ding Zhang, "Projected Dynamical Systems," in Anna Nagurney and Ding Zhang, eds., *Projected Dynamical Systems and Variational Inequalities with Applications*, Boston, MA: Springer US, 1996, pp. 9–43. [17]

Palaiopanos, Gerasimos, Ioannis Panageas, and Georgios Piliouras, "Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos," *Advances in Neural Information Processing Systems*, 2017, *30.* [3]

Papadimitriou, Christos and Georgios Piliouras, "Game dynamics as the meaning of a game," *ACM SIGecom Exchanges*, 2019, *16* (2), 53–63. [3]

Patriksson, Michael and R Tyrrell Rockafellar, "Sensitivity analysis of aggregated variational inequality problems, with application to traffic equilibria," *Transportation Science*, 2003, *37* (1), 56–68. [4]

Putinar, Mihai, "Positive Polynomials on Compact Semi-algebraic Sets," *Indiana University Mathematics Journal*, 1993, *42* (3), 969–984. [23, 25]

Ratliff, Lillian J., Samuel A. Burden, and S. Shankar Sastry, "Characterization and computation of local Nash equilibria in continuous games," in "2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)" 2013, pp. 917–924. [4]

Roughgarden, Tim, "Intrinsic robustness of the price of anarchy," *Journal of the ACM (JACM)*, 2015, *62* (5), 1–42. [4]

Sanders, James BT, J Doyne Farmer, and Tobias Galla, "The prevalence of chaotic dynamics in games with many players," *Scientific Reports*, 2018, *8* (1), 1–13. [3]

Song, Chaobing, Zhengyuan Zhou, Yichao Zhou, Yong Jiang, and Yi Ma, "Optimistic Dual Extrapolation for Coherent Non-monotone Variational Inequalities," in H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, eds., *Advances in Neural Information Processing Systems*, Vol. 33 Curran Associates, Inc. 2020, pp. 14303–14314. [4, 13, 34, 35, 36]

Sorin, Sylvain, "Continuous time learning algorithms in optimization and game theory," *Dynamic Games and Applications*, 2023, *13* (1), 3–24. [2]

———, "No-regret algorithms in on-line learning, games and convex optimization," *Mathematical Programming*, 2024, *203* (1), 645–686. [2]

Souaiby, Marianne, Aneel Tanwani, and Didier Henrion, "Computation of Lyapunov Functions under State Constraints Using Semidefinite Programming Hierarchies," *IFAC-PapersOnLine*, January 2020, *53* (2), 6281–6286. [17, 20, 21, 24, 25]

Strodiot, Jean Jacques, Phan Tu Vuong, and Thi Thu Van Nguyen, "A class of shrinking projection extragradient methods for solving non-monotone equilibrium problems in Hilbert spaces," *Journal of Global Optimization*, 2016, *64* (1), 159–178. [4, 13]

Tilli, Tuomo and Leonardo Espinosa-Leal, "Multi-armed bandits for bid shading in first-price real-time bidding auctions," *Journal of Intelligent & Fuzzy Systems*, 2021, *41* (6), 6111–6125. [2]

Tseng, Paul, "On linear convergence of iterative methods for the variational inequality problem," *Journal of Computational and Applied Mathematics*, 1995, *60* (1-2), 237–252. [4]

Ui, Takashi, "Bayesian Nash Equilibrium and Variational Inequalities," *Journal of Mathematical Economics*, March 2016, *63*, 139–146. [4]

Vickrey, William, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, 1961, *16* (1), 8–37. [2]

Vlatakis-Gkaragkounis, Emmanouil-Vasileios, Lampros Flokas, and Georgios Piliouras, "Chaos persists in large-scale multi-agent learning despite adaptive learning rates," *arXiv preprint arXiv:2306.01032*, 2023. [3]

Wang, Qian, Zongjun Yang, Xiaotie Deng, and Yuqing Kong, "Learning to bid in repeated first-price auctions with budgets," in "International Conference on Machine Learning" PMLR 2023, pp. 36494–36513. [2]

Wilson, Robert, *Game-theoretic analyses of trading processes* Econometric Society Monographs, Cambridge University Press, [2]

Ye, Minglu, "An infeasible projection type algorithm for nonmonotone variational inequalities," *Numerical Algorithms*, 2022, *89* (4), 1723–1742. [4]

Young, H Peyton, *Strategic Learning and its Limits*, OUP Oxford, 2004. [2]

Zhang, Wei, Yanjun Han, Zhengyuan Zhou, Aaron Flores, and Tsachy Weissman, "Leveraging the hints: Adaptive bidding in repeated first-price auctions," *Advances in Neural Information Processing Systems*, 2022, *35*, 21329–21341. [2]

APPENDIX A: Minty condition and the Optimistic Dual Extrapolation Algorithm

In Lemma 2, we have shown that a unique solution to the MVI exists for the symmetric, second-price sealed-bid auction. In the following, we discuss the applicability of the optimistic dual extrapolation algorithm, which was shown to converge to the solution of the MVI by Song et al. (2020). To apply their results, we first need to reformulate the problem in a Hilbert space setting. To this end, we use the Hilbert space $\mathcal{H} = H^1((0,1);F)$ with inner product

$$(v,w)_{\mathcal{H}} := \int_0^1 v(x)w(x) + v'(x)w'(x)\mathrm{d}F(x)$$

and induced norm $\|v\|_{\mathcal{H}} := \sqrt{(v,v)_{\mathcal{H}}}$. Analogously to the set $\mathcal{B}_\delta \subset V$ for a Banach space V, we consider in the Hilbert space $\mathcal{H}$ the closed and convex set

$$\mathcal{W}_\delta := \{h \in \mathcal{H} \;:\; 0 < \delta \le h' \le \delta^{-1} \; F\text{-a.e., and } 0 \le h \le 1 \; F\text{-a.e.}\}$$

for $0 < \delta \le 1$. Note that $\mathcal{W}_\delta \subsetneq \mathcal{B}_\delta$, i.e. this setting is slightly more restrictive that that of Section 3.1. We proceed by showing the assumptions of Song et al. (2020, Assumptions 1, 2 and 3):

1. The operator $DU : \mathcal{W}_\delta \to \mathcal{H}^*$ defined by (8) is Lipschitz-continuous.
2. There are constants $0 < \gamma \le 1$ and $C > 0$ such that the norms satisfy

$$\begin{aligned} \tfrac{1}{2}\|w\|_{\mathcal{H}}^2 &\ge \tfrac{1}{2}\|v\|_{\mathcal{H}}^2 + D\big(\tfrac{1}{2}\|v\|_{\mathcal{H}}^2\big)[w-v] + \tfrac{\gamma}{2}\|w-v\|_{\mathcal{H}}^2, \\ \Big\|D\big(\tfrac{1}{2}\|w\|_{\mathcal{H}}^2\big)\Big\|_{\mathcal{H}^*} &\le C\|w\|, \end{aligned}$$

   for all $v, w \in \mathcal{H}$. Here, $D(\frac{1}{2}\|v\|_{\mathcal{H}}^2) \in \mathcal{H}^*$ denotes the Gateaux-derivative of $v \mapsto \frac{1}{2}\|v\|_{\mathcal{H}}^2$.
3. There exists a solution $\beta^* \in \mathcal{W}_\delta$ to the symmetric MVI for the second-price sealed-bid auction.

The latter directly follows from Lemma 2 for $\delta \le 1$, as $\beta^* = \mathrm{Id} \in \mathcal{W}_\delta \subset \mathcal{H}$. The norms satisfy Assumption 2, due to the Hilbert space setting. In particular, we have $D(\frac{1}{2}\|v\|_{\mathcal{H}}^2) = (v,\cdot)_{\mathcal{H}} \in \mathcal{H}^*$, such that we can use $\gamma = C = 1$. Finally, the operator $DU : \mathcal{W}_\delta \to \mathcal{H}^*$ is

Lipschitz-continuous with Lipschitz-constant $L_{DU} \leq 2\delta^{-2}\|g\|_{L^\infty}$ since

$$\begin{aligned}
\left|\left(DU(\boldsymbol{\beta}) - DU(\tilde{\boldsymbol{\beta}})\right)[d]\right| &= \left|\int_0^1 d(x)\left[\frac{x-\beta(x)}{\beta'(x)} - \frac{x-\tilde{\beta}(x)}{\tilde{\beta}'(x)}\right] g(x)\mathrm{d}F(x)\right| \\
&\leq \|d\|_{L^2}\|g\|_{L^\infty}\left\|\frac{\left(x-\beta(x)\right)\tilde{\beta}'(x) - \left(x-\tilde{\beta}(x)\right)\beta'(x)}{\beta'(x)\tilde{\beta}'(x)}\right\|_{L^2} \\
&\leq \|d\|_{\mathcal{H}}\|g\|_{L^\infty}\left\|\frac{\tilde{\beta}(x)-\beta(x)}{\beta'(x)}\right\|_{L^2} \\
&\quad + \|d\|_{\mathcal{H}}\|g\|_{L^\infty}\left\|\frac{\left(x-\tilde{\beta}(x)\right)\left(\tilde{\beta}'(x)-\beta'(x)\right)}{\beta'(x)\tilde{\beta}'(x)}\right\|_{L^2} \\
&\leq \|d\|_{\mathcal{H}}\|g\|_{L^\infty}\left(\delta^{-1}\|\beta-\tilde{\beta}\|_{L^2} + \delta^{-2}\|\beta'-\tilde{\beta}'\|_{L^2}\right) \\
&\leq \|d\|_{\mathcal{H}} 2\delta^{-2}\|g\|_{L^\infty}\|\beta-\tilde{\beta}\|_{\mathcal{H}}
\end{aligned}$$

for $0 < \delta \leq 1$.

Following the argumentation (Song et al. 2020, Theorem 1), now applied to the infinite-dimensional $\mathcal{H}$ instead of $\mathbb{R}^n$, we obtain the convergence of the optimistic dual extrapolation (Song et al. 2020, Algorithm 1), given here in Algorithm 1. In particular, we obtain:

Proposition 5: After K iterations, Algorithm 1 returns a $\tilde{\beta}_K \in \mathcal{W}_\delta$ such that

$$\sup_{\beta\in\mathcal{W}_\delta}\left|DU(\tilde{\boldsymbol{\beta}}_K)[\tilde{\beta}_K - \beta]\right| \leq \left(2\sqrt{2}+16\right)\left(1+\delta^{-2}\right)L_{DU}K^{-1/2}.$$

Here, we optimized the error constant by choosing the maximal free parameter α of Song et al. (2020) and using that $\left\|\beta-\tilde{\beta}\right\|_{\mathcal{H}} \leq \sqrt{1+\delta^{-2}}$ for all $\beta, \tilde{\beta} \in \mathcal{W}_\delta$. Note that in Algorithm 1 the proximal mapping $P_v : \mathcal{H}^* \to \mathcal{H}$ for $v \in \mathcal{H}$ is given by

$$P_v(F) := \arg\min_{z\in\mathcal{W}_\delta}\left(F[z] + \frac{1}{2}\|z-v\|^2_{\mathcal{H}}\right).$$

This requires solving a *monotone* VI on $\mathcal{W}_\delta$ for each evaluation.

Co-editor [Name Surname; will be inserted later] handled this manuscript.

Algorithm 1 Optimistic dual extrapolation algorithm (Song et al. (2020))

Input: Constant $c = \left(4\sqrt{2}L_{DU}\right)^{-1} > 0$.

1: $\beta_0 = z_0 \in W_\delta$, $G_0 = 0 \in \mathcal{H}^*$
2: **for** $k = 1, 2, \dots, K$ **do**
3: $\quad \beta_k = P_{z_{k-1}}\left(-cDU(\boldsymbol{\beta_{k-1}})\right)$
4: $\quad G_k = G_{k-1} - cDU(\boldsymbol{\beta_k})$
5: $\quad z_k = P_{\beta_0}\left(G_k\right)$
6: **end for**
7: $\tilde{\beta}_K = \arg\min_{\beta_k : 1 \le k \le K} \left(\|\beta_k - z_{k-1}\| + \|\beta_{k-1} - z_{k-1}\|\right)$