Simultaneous Gradient Learning in First-Price Auctions

Janik Bürgermeister[*]
Department of Computer Science
Technical University of Munich
Munich, Germany

Julius Durmann
Department of Computer Science
Technical University of Munich
Munich, Germany

Martin Bichler
Department of Computer Science
Technical University of Munich
Munich, Germany

Mete Şeref Ahunbay
Department of Computer Science
University of Oxford
Oxford, UK

Bary Pradelski
Department of Computer Science
CNRS
Grenoble, France

Marco Scarsini
Dpt. of Econ. and Financial Markets
Luiss University
Rome, Italy

Abstract

We study equilibrium learning in discretized Bayesian games, focusing on first-price auctions as a testbed for gradient-based learning dynamics under private information. While the symmetric independent private values model admits a unique Bayesian correlated equilibrium (BCE) in the continuous limit, standard convergence guarantees for gradient-based learners do not apply: classical no-regret theory predicts convergence only to Bayesian coarse correlated equilibria (BCCE), which can remain large and far from the competitive equilibrium even under fine discretizations. To bridge this gap, we introduce the notion of Bayesian semicoarse correlated equilibrium (BSCCE), an equilibrium concept that refines BCCE by restricting deviations to probability-preserving transformations consistent with projected gradient dynamics. We provide a linear programming characterization of BSCCE with polynomially many variables and constraints and establish the inclusion hierarchy $BCE \subseteq BSCCE \subseteq BCCE$. Our main finding is that the BSCCE set contracts in many cases as the discretization of the Bayesian game is refined. Even with priors for which BCCE fails to concentrate, the diameter of the BSCCE polytope shrinks toward zero, indicating convergence to the unique continuous-game Bayes-Nash equilibrium. These results provide a principled explanation for the observed robustness of gradient-based learning in discretized Bayesian games, showing that these dynamics select a smaller equilibrium set than standard no-regret theory predicts.

1 Introduction

Learning in multi-agent systems fundamentally differs from single-agent learning: strategy updates are shaped by the adaptation of other agents. As a result, learning dynamics in games need not converge, may cycle, exhibit chaotic behavior, or converge to different outcomes depending on initial

[*]Corresponding author; janik.buergermeister@tum.de

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

conditions. Understanding the subset of games where commonly used learning algorithms converge to meaningful equilibrium predictions remains a central challenge at the intersection of machine learning, optimization, and game theory.

Somewhat surprisingly, recent papers observe convergence close to the Bayes-Nash equilibria (BNE) in economically relevant Bayesian games (Bichler et al., 2023, 2021). These methods build on gradient learning, and we will focus on its simplest variant, projected gradient ascent, in this paper.

We will deliberately restrict our attention to first-price single-item auctions. This game is a canonical example used in large-scale applications such as display advertisements. However, despite the uniqueness of the BNE, the empirical convergence of gradient-based algorithms has not yet been fully theoretically justified. Here, *we aim to provide an explanation for the convergence of projected gradient ascent in first-price auctions under increasingly finer discretizations.* The BSCCE, however, is a broader concept and can be applied to all games with private information.

Gradient Dynamics in Games. Motivated by their algorithmic simplicity, recent work studies gradient-based learning dynamics in games (Mertikopoulos & Sandholm, 2018; Mertikopoulos et al., 2024; Cai et al., 2024). While it is known that Nash equilibria are equilibria of the resulting dynamical system, strong assumptions, including strict monotonicity, are usually needed to guarantee uniqueness and the convergence of the dynamics to such an equilibrium in last-iterate, independent of initial conditions (Mertikopoulos & Zhou, 2018; Zhang & Nagurney, 1995). Similarly, Lyapunov-function-based arguments can be used to show stability (Mertikopoulos & Sandholm, 2016). In general, the dynamics of a game may also cycle, prohibiting last-iterate convergence statements in general.

Dynamical systems offer a powerful toolset, but the analysis is game-specific and often requires significant effort to use (see, e.g., Bichler et al. (2024)). In Bayesian games with private information, the picture is even less well understood: equilibrium strategies are functions that map types to actions, which further complicates the analysis. In particular, Bichler et al. (2024, 2025) have shown that the first-price auction satisfies neither the monotonicity nor the (weaker) Minty condition, which rules out many of the related convergence statements.

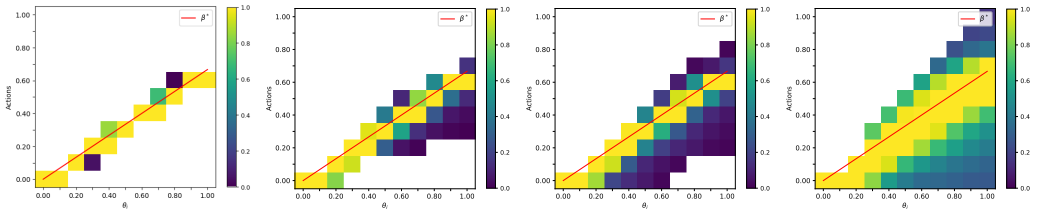
No-regret Learning and Bayesian (Coarse) Correlated Equilibria. As an alternative to dynamical systems, we draw on no-regret learning theory, shifting the focus from last-iterate to time-average convergence. Indeed, it is well known that no-regret algorithms converge to (Bayesian) coarse correlated equilibria (BCCE) (Hartline et al., 2015). Similarly, no-internal-/no-swap-regret algorithms converge to (Bayesian) correlated equilibria (BCE) (Gordon et al., 2008). However, the BCCE can be a large superset of the BNEs of a game, and this result says little about the convergence of no-regret algorithms to a BNE. Meanwhile, no-internal regret algorithms are much more complex, and BCEs may be too strict to capture the behavior of gradient learning algorithms.

In the context of auctions, Ahunbay & Bichler (2024) recently showed that the set of BCCEs collapses to the unique BNE in all-pay auctions with continuous types, which proves convergence of no-regret algorithms to an approximate BNE in discretized versions of the auction. However, the authors show that this explanation is insufficient in first-price auctions: the BCCE set remains large for non-concave priors even under arbitrarily fine discretizations and can contain outcomes far from the competitive equilibrium. Thus, standard no-regret guarantees do not explain why gradient-based learners consistently converge close to the BNE in this setting.

Our research relates to the work by Bergemann et al. (2017) which analyzes BCEs under varying levels of information provided to the bidders. They find a subset relation for the BCEs as information quality increases (modeled as "information structures"), meaning that as information quality increases, the BCE set shrinks. Slightly different from our setup, their model assumes that bidders do not observe their private valuations directly, but instead receive a signal chosen by the auctioneer.

Contributions. The large support of BCCE warrants a closer inspection of gradient algorithms in first-price auctions. The central contribution of this paper is to refine the equilibrium notion used to analyze gradient-based learning in Bayesian games. We introduce the *Bayesian Semicoarse Correlated Equilibrium (BSCCE)*, which captures the outcomes of gradient dynamics more precisely than prior concepts. BSCCE is motivated by the observation that gradient-based algorithms can realize only a restricted class of deviations, namely, tangential, probability-preserving transformations of mixed strategies, which correspond to gradients of potential functions.

BSCCE lies between BCE and BCCE as it restricts (expands) the class of deviations relative to BCE (BCCE). Figure 1 illustrates the maximal support across equilibrium concepts in a first-price auction.



(a) BNE (b) BCE (c) BSCCE (d) BCCE

Figure 1: Maximal conditional bid probabilities ($\sigma(a_i \mid \theta_i)$) in the equilibrium polytopes for a $\mathrm{Beta}(2,1)$ prior using uniform discretization in a first-price auction.

Formally, BSCCE extends the notion of Semicoarse Correlated Equilibrium (SCCE), originally developed for complete-information games (Ahunbay & Bichler, 2025), to Bayesian settings. However, this extension must be handled carefully, as the strategic normal form of a Bayesian game is exponential in the number of types. Thus, applying the SCCE to the strategic normal form is computationally infeasible. In contrast, our approach is based on the partial Bayesian approach Forges (1993), which allows a polynomial characterization of BSCCE in terms of the game size and thus makes BSCCE tractable in discretized games.

BSCCE is defined via linear incentive constraints against a class of deviations, restricted to affine, symmetric, probability-preserving deviations that capture the geometry of projected gradient ascent. This approach relates to the concept of Φ-equilibria, as introduced by Greenwald & Jafari (2003) and Lugosi & Stoltz (2007). Another related variant of Φ-equilibria are linear correlated equilibria (LCE), for which Daskalakis et al. (2025) provide a polynomial algorithm that has no-Φ-regret in convex games. The deviations leading to LCEs include those considered in this paper. Therefore, LCEs are a stricter notion of equilibrium ($\mathrm{LCE} \subseteq \mathrm{SCCE}$). We focus on (B)SCCE because of its relation to gradient learners, which has been established by Ahunbay, Mete (2025); Ahunbay & Bichler (2025).

We apply BSCCEs to discrete first-price auctions with different priors. Our results show convergence of gradient-based algorithms in many first-price auctions as the discretization is refined. This includes auctions with non-concave priors where the BCCE set remains large (Ahunbay, Mete, 2025).

Taken together, these results justify the empirical success of gradient-based learning in discretized Bayesian games, selecting a smaller equilibrium set than standard no-regret theory predicts. A full theory for the convergence of gradient algorithms remains an open task, but our proposed concept provides important insights on the convergence of gradient-based learning algorithms.

2 The Model

We consider strategic interactions under incomplete information. We begin by defining the standard game-theoretic concepts for an incomplete information setting. Let $\Delta(\mathcal{H})$ denote the set of probability distributions over a measurable space $(\mathcal{H}, \mathcal{B}(\mathcal{H}))$, where $\mathcal{B}(\mathcal{H})$ is the Borel σ-algebra associated with the metric space $\mathcal{H}$.

Definition 2.1 (Bayesian Game)**.** A *Bayesian game* is a tuple $\mathcal{G} = (\mathcal{N}, \Theta, \mathcal{A}, u, \pi)$, where

- $\mathcal{N} = \{1, \ldots, n\}$ is a finite set of players.
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is the type space.
- $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ is the space of action profiles.
- $u = (u_1, \ldots, u_n)$ are utility functions, where $u_i : \Theta_i \times \mathcal{A} \to \mathbb{R}$ specifies the payoff of player i.
- $\pi \in \Delta(\Theta)$ is the joint prior distribution of types.

When discussing distributions such as $\sigma(\theta, a)$, we denote the marginal distributions by omitting the marginalized variables. For example, $\sigma(a)$ denotes the marginal distribution over actions. We

denote conditional distributions and expectations with vertical bars; for instance, $\sigma(a_{-i} \mid \theta_i, a_i)$ is the conditional probability of the opponents' actions a_{-i} given player i's action a_i and type θ_i.

A *pure strategy* for player i is a mapping $s_i : \Theta_i \to \mathcal{A}_i$. A *mixed strategy* $x_i : \Theta_i \to \Delta(\mathcal{A}_i)$ maps types to distributions over actions.

The expected utility of player i given a strategy profile x and type θ_i (*interim utility*) is $u_i(x_i, x_{-i} \mid \theta_i)$:

$$\mathbb{E}_{\theta_{-i}, a \sim x}[u_i(\theta_i, a) \mid \theta_i] = \sum_{\theta_{-i} \in \Theta_{-i}} \pi(\theta_{-i} \mid \theta_i) \sum_{a \in \mathcal{A}} \left(\prod_{j \in \mathcal{N}} x_j(a_j \mid \theta_j) \right) u_i(\theta_i, a).$$

This definition can be extended to infinite type- and action spaces by replacing the corresponding sums with integrals. Observe that the interim utility is linear in x_i, the player's own action distribution.

Definition 2.2 (Bayes-Nash Equilibrium (BNE))**.** A strategy profile is a *Bayes-Nash Equilibrium* if it maximizes interim utility for almost all types. More specifically, a *distributional* profile $x^* = (x_1^*, \ldots, x_n^*)$, where $x_i^* : \Theta_i \to \Delta(\mathcal{A}_i)$, is a BNE if for all $i \in \mathcal{N}$ and almost all $\theta_i \in \Theta_i$:

$$u_i(x^* | \theta_i) \geq u_i(a_i', x_{-i}^* | \theta_i) \quad \forall a_i' \in \mathcal{A}_i.$$

There are multiple approaches to define Bayesian correlated equilibria (BCEs) (Forges, 1993). We will rely on the "partial Bayesian approach", which assumes that all players observe their types and a recommendation for a corresponding action. In contrast to the "strategic form" and "agentic form" approaches, recommendations are correlated on the level of individual actions and may depend on the realized types of *all* players. For this, we need the following concepts. Let $\sigma \in \Delta(\Theta \times \mathcal{A})$ denote a distribution over types and actions.

Definition 2.3 (Prior Consistency)**.** σ satisfies *prior consistency* with respect to a prior π if the marginal distribution of types induced by σ coincides with the exogenous prior π.

$$\sum_{a \in \mathcal{A}} \sigma(\theta, a) = \sigma(\theta) = \pi(\theta) \quad \forall \theta \in \Theta.$$

Definition 2.4 (Conditional Independence)**.** σ satisfies *conditional independence* if

$$\sigma(\theta_{-i} \mid \theta_i, a_i) = \sigma(\theta_{-i} \mid \theta_i) \quad \forall i.$$

This condition ensures that a player's recommendation does not reveal information about opponents.

Definition 2.5 (BCE and BCCE)**.** A distribution $\sigma \in \Delta(\Theta \times \mathcal{A})$ satisfying prior consistency is

a *Bayesian correlated equilibrium (BCE)* if, for all players i, types θ_i, and actions a_i, a_i',

$$\sum_{\theta_{-i}, a_{-i}} \sigma(\theta, a) \big[u_i(a_i, a_{-i} | \theta_i) - u_i(a_i', a_{-i} | \theta_i) \big] \geq 0; \tag{1}$$

a *Bayesian coarse correlated equilibrium (BCCE)* if, for all players i, types θ_i, and actions a_i',

$$\sum_{\theta_{-i}, a} \sigma(\theta, a) \big[u_i(a | \theta_i) - u_i(a_i', a_{-i} | \theta_i) \big] \geq 0.$$

Optionally, one may impose conditional independence.

The *Incentive compatibility* constraint (1) for a BCE must hold for each type, and means that no player wants to deviate from their recommended action a_i to another action a_i' after observing their type θ_i. The hierarchy of the introduced equilibrium concepts is well-known:

$$BNE \subseteq BCE \subseteq BCCE.$$

Conditional independence is optional across equilibrium concepts. The inclusion hierarchy is preserved whether conditional independence is imposed for all equilibrium concepts or for none. We adopt the former as the default.

Single-Item, First-Price Auctions. We focus on the standard symmetric Independent Private Values (IPV) model. The set of bidders is $\mathcal{N} = \{1, \dots, n\}$. Each bidder i draws a private valuation θ_i independently from a probability mass function (or density) π with corresponding cumulative distribution function F with support on $[0, \bar{\theta}]$. W.l.o.g. we can set $\bar{\theta} = 1$. Bidders submit bids $a_i \in \mathcal{A}_i = [0, 1]$. The restriction $\mathcal{A}_i = [0, 1]$ is w.l.o.g. as rational bidders never overbid.

The auction is defined by an allocation rule $q_i(a)$. The utility of bidder i with valuation θ_i is:

$$u_i(a|\theta_i) = q_i(a) \cdot (\theta_i - a_i). \tag{2}$$

We assume a standard highest-wins allocation rule with uniform tie-breaking:

$$q_i(a) = \begin{cases} 1 & \text{if } a_i > \max_{j \neq i} a_j \\ \frac{1}{|\{j : a_j = \max a\}|} & \text{if } a_i = \max_j a_j \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

In the *first-price auction*, the winner pays their bid.

Under regularity assumptions (Chawla & Hartline, 2013), the first-price auction has a unique BNE β^{FP} in the symmetric IPV model, which is given by the expectation of the second-highest valuation conditioned on the bidder having the highest valuation (i.e., the highest of the others):

$$\beta_i^{\text{FP}}(\theta_i) = \mathbb{E}[Y_1 \mid Y_1 < \theta_i], \tag{4}$$

where Y_1 is the highest order statistic of the $n-1$ other bidders' valuations. For example, if valuations are uniformly distributed ($\theta_i \sim U[0, 1]$), this simplifies $\beta_i^{\text{FP}}(\theta_i) = \frac{n-1}{n}\theta_i$.

Equilibria in the Discretized First-Price Auction. To find equilibria in more general settings, where theoretical predictions are unavailable, learning algorithms discretize the game. However, this discretization alters the game already in the symmetric IPV model and breaks the existence of the unique pure symmetric BNE (Rasooly & Gavidia-Calderon, 2022). Consequently, learning algorithms are no longer expected to converge to a singleton but rather to a set of mixed strategies. Figure 1 illustrates the set of possible BNEs and BCEs for an example. We are interested in how the limit set of the learning dynamics, in particular the projected gradient dynamics, changes as the discretization becomes finer.

3 Bayesian Semicoarse Correlated Equilibrium

Most learning dynamics are not guaranteed to converge to BNEs but only to relaxed notions thereof. On the one hand, BCEs are a relatively tight superset of BNEs, but they lack convergence guarantees for gradient-type algorithms. On the other hand, the set of BCCEs is too large in the case of First Price auctions to derive meaningful convergence statements (Ahunbay & Bichler, 2024). Therefore, we need to move on to another, relaxed notion of equilibrium. We propose the Bayesian Semicoarse Correlated Equilibrium (BSCCE) for this purpose, which lies between BCE and BCCE.

3.1 Semicoarse Correlated Equilibrium (SCCE)

SCCEs were introduced by Ahunbay & Bichler (2025) to capture the convergence set of gradient-based learners in complete-information games while maintaining computational feasibility. It can be understood as a specific instance of Φ-equilibria (Greenwald & Jafari, 2003; Lugosi & Stoltz, 2007) that includes CEs but is (sometimes) strictly smaller than the CCE polytope. Ahunbay & Bichler (2025) already applied this solution concept to complete-information versions of Bertrand competition and first-price auctions.

Let us recall the definition of SCCEs and provide an intuitive explanation. Gradient-based dynamics satisfy no-regret guarantees with respect to a restricted class of deviations that correspond to gradient fields of potential functions (Ahunbay, Mete, 2025). By restricting to gradients of quadratic functions, we obtain a solution concept that is representable as an LP and thus computationally tractable, and it is a superset of the outcomes of gradient dynamics.

A probability distribution $\sigma \in \Delta(\mathcal{A})$ is an SCCE if it is stable against infinitesimal deviations in the direction $Qx_i + q_i$ defined by a pair (Q_i, q_i) for each player i, where

- $Q_i \in \mathbb{R}^{|\mathcal{A}_i| \times |\mathcal{A}_i|}$ is a symmetric matrix defining probability flow between actions, and
- $q_i \in \mathbb{R}^{|\mathcal{A}_i|}$ is a vector providing an offset.

Since the utility function u_i is linear in player i's action, restricting attention to infinitesimally small deviations entails no loss of generality, as any local deviation can be extended globally by linearity.

These deviations capture the form of update directions that arise under gradient steps which are projected back into the simplex. Furthermore, to avoid projections in the analysis and to tailor the solution concept to mixed extensions of normal-form games, Ahunbay & Bichler (2025) restrict the (Q_i, q_i) according to three conditions:

1. *Tangency:* $\sum_{a_i' \in \mathcal{A}_i} (Q_i(a_i', a_i) + q_i(a_i')) = 0$ for all a_i. This ensures probability is merely reallocated between actions, not created or destroyed.

2. *Non-negativity:* $Q_i(a_i', a_i) + q_i(a_i') \geq 0$ for all $a_i' \neq a_i$. This condition ensures that, at pure strategies, admissible deviation directions lie within the tangent cone of the probability simplex, so that infinitesimal deviations preserve feasibility.

3. *Symmetry:* $Q_i = Q_i^\top$. This ensures that deviations are gradient fields of quadratic functions. Projected gradient ascent can only move along gradients of scalar functions. Without symmetry, the deviation direction of $Q_i x_i + q_i$ can exhibit rotational flow, which is impossible for gradient fields.

The incentive constraints are infinitesimal deviations in the direction $Q_i x_i + q_i$ evaluated for pure strategies (i.e., x_i is a canonical basis vector). Thus, the pairs (Q_i, q_i) define deviations $d_i(a', a) = Q_i(a', a) + q_i(a')$ on actions $a', a \in \mathcal{A}_i$. We denote the *set of admissible deviations*, i.e., where d_i satisfies the conditions above, by $\mathcal{D}$.

In summary, we get the following definition of SCCEs.

Definition 3.1 (Semicoarse Correlated Equilibrium (SCCE), Ahunbay & Bichler (2025))**.** A distribution σ on the joint action profiles is a *semicoarse correlated equilibrium* if, for every agent i and every $d \in \mathcal{D}$, the following incentive compatibility constraint holds:

$$\sum_{a \in A} \sigma(a) \left[\sum_{a_i' \in \mathcal{A}_i} d(a_i', a_i) u_i(a_i', a_{-i}) \right] \leq 0.$$

Although there are uncountably many incentive compatibility constraints as the set of pairs (Q_i, q_i) forms a polyhedral cone, this cone admits a finite representation via its extreme rays. This yields a finite set of constraints, as detailed in Appendix A. Moreover, the resulting linear program has size polynomial in the cardinalities of the type and action spaces and thus is computationally tractable.

3.2 Extension to Bayesian Games

For the extension of SCCEs to Bayesian games, we adopt the *partial Bayesian* framework in line with the above definition of BCEs. This choice is motivated by the structure of learning dynamics in auction games: gradient-based algorithms operate on *type-contingent mixed actions* $x_{i,t}(\cdot \mid \theta_i)$ and update them using local, interim utility gradients. They do not optimize over or maintain distributions over complete strategy functions $s_i : \Theta_i \to \mathcal{A}_i$. The partial Bayesian formulation mirrors this behavior by defining equilibrium objects directly as joint distributions over types and actions.

A Bayesian Semicoarse Correlated Equilibrium (BSCCE) is therefore defined as a joint distribution $\sigma \in \Delta(\Theta \times \mathcal{A})$. The definition inherits the structural conditions of Bayesian correlated equilibrium (Definition 2.5), such as prior consistency and conditional independence, but replaces the standard incentive constraints with semicoarse incentive constraints tailored to gradient-based deviations.

We build on the SCCE framework and allow for admissible deviations that are *type-specific*: for each player i and type $\theta_i \in \Theta_i$, deviations are generated by a pair $(Q_i^{\theta_i}, q_i^{\theta_i})$ defining a tangential, probability-preserving transformation of the action simplex. For notational simplicity, we denote such deviations by $d^{\theta_i}(a_i', a_i) = Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i')$.

Definition 3.2 (Bayesian Semicoarse Correlated Equilibrium (BSCCE))**.** A *Bayesian Semicoarse Correlated Equilibrium (BSCCE)* is a joint probability distribution $\sigma \in \Delta(\Theta \times \mathcal{A})$ that satisfies:

1. *Prior consistency.*

2. *Conditional independence.*

3. *Semicoarse incentive compatibility.* No player can gain by admissible deviations. Formally, we require that, for every player i, every type $\theta_i \in \Theta_i$, and every deviation $d^{\theta_i} \in \mathcal{D}$

$$\mathop{\mathbb{E}}_{(\theta_{-i},a)\sim\sigma}\Big[\sum_{a_i'\in\mathcal{A}_i} d^{\theta_i}(a_i', a_i)\, u_i(a_i', a_{-i} \mid \theta_i) \mid \theta_i\Big] \le 0.$$

The conditioning of expectations affects the player's beliefs about other players and their actions.

The BSCCE constraints are all linear. We will add a linear objective that measures the distance to a given distribution (the BNE strategy) and obtain a linear program, which we provide in Appendix A.

Relation to other equilibrium concepts. Equilibrium notions can be ordered by the class of admissible deviations they allow. Allowing a richer class of deviations induces more incentive constraints and therefore yields a smaller equilibrium set. The Bayes-Nash equilibrium (BNE) is the most restrictive notion, as it allows all unilateral deviations at the level of complete strategies and requires that the actions be sampled independently for each player. Bayesian correlated equilibrium (BCE) relaxes this requirement by allowing players to condition deviations on correlated action recommendations, while still enforcing obedience to all such deviations. BSCCE further relaxes the equilibrium notion by restricting attention to a *proper subset* of deviations. Bayesian coarse correlated equilibrium (BCCE) imposes the weakest incentive constraints by allowing only unconditional deviations, and therefore yields the largest equilibrium set. Accordingly, the equilibrium concepts satisfy the inclusion chain

$$\text{BNE} \subseteq \text{BCE} \subseteq \text{BSCCE} \subseteq \text{BCCE}. \tag{5}$$

4 Bayesian Gradient Ascent Learners

To relate BSCCE to learning dynamics, we study gradient-based algorithms in discretized Bayesian games. Our goal is not to model all possible feedback mechanisms, but to identify a class of learning dynamics whose limiting behavior can be characterized by semicoarse incentive constraints.

Expected-Utility Formulation. A key modeling choice is how learning agents process private information. Rather than modeling stochastic type realizations and noisy gradient feedback, we adopt an *expected-utility formulation* that treats the Bayesian game as a deterministic, high-dimensional game over type-contingent mixed actions. Specifically, each player i maintains a mixed action $x_i(\cdot \mid \theta_i) \in \Delta(\mathcal{A}_i)$ for every type $\theta_i \in \Theta_i$. The player's decision variable is therefore the vector

$$x_i \;\in\; X_i := \prod_{\theta_i\in\Theta_i} \Delta(\mathcal{A}_i),$$

and payoffs are given by the ex-ante expected utility

$$u_i(x) \;=\; \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{a\sim x(\cdot|\theta)}[u_i(a|\theta_i)].$$

Under this formulation, the Bayesian game can be viewed as a complete-information normal-form game with a structured action space X_i. The utility functions are smooth and multilinear in the mixed actions, and gradients with respect to $x_i(\cdot \mid \theta_i)$ are well defined.

Bayesian Projected Gradient Ascent. We consider players who apply projected gradient ascent to maximize their ex-ante expected utility. Given a step size sequence $(\eta_t)_{t\ge1}$, the update rule for player i is

$$x_{i,t+1} = \Pi_{X_i}\big[x_{i,t} + \eta_t \nabla_{x_i} u_i(x_{1,t}, \ldots, x_{n,t})\big], \tag{6}$$

where Π_{X_i} denotes Euclidean projection. Because the feasible set X_i is a Cartesian product of simplices, the projection decomposes across types. We assume that step-sizes satisfy $\eta_t = \omega(1/t)$ and $\sum_{t=1}^{T} \eta_t = o(T)$.[2]

This deterministic gradient dynamic avoids stochasticity from sampled type realizations and allows us to directly apply no-regret and stability arguments developed for gradient-based learning in complete-information games.

[2]Both requirements are in Bachmann-Landau notation.

Connection to Semicoarse Equilibria. The expected-utility view of our Bayesian game allows us to extend the convergence result of Ahunbay & Bichler (2025, Proposition 2.1), which shows that projected gradient ascent converges to the set of SCCEs in time-average under complete information, to our incomplete-information setting.

Corollary 4.1 (BPGA implies BSCCE for time-average limit points)**.** *Consider a finite Bayesian game with finite type sets Θ_i and finite action sets $\mathcal{A}_i$. Let each player i control a type-contingent mixed action vector $x_i \in X_i$ and let the (ex-ante) expected utility be $u_i(x)$. Let $(x_t)_{t\geq 1}$ be generated by Bayesian projected gradient ascent (BPGA). Define the time-average distribution on $\Theta \times A$ by*

$$\sigma_T(\theta, a) := \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j). \tag{7}$$

Then every limit point σ of $(\sigma_T)_{T\geq 1}$ is a BSCCE.

We provide a proof in Appendix B. This establishes BSCCE as a natural equilibrium concept for characterizing the long-run behavior of gradient-based learners in discretized Bayesian games.

Stochastic learning and sampled feedback. The BPGA dynamics studied above should be understood as an analytical proxy for the stochastic learning algorithms commonly used in practice. In typical auction-learning implementations, agents observe batched realized types and payoffs and update their mixed strategies using stochastic gradient estimates rather than exact interim gradients. Under standard assumptions, such stochastic updates track the corresponding expected-utility gradient dynamics in the limit (Benveniste et al., 2012). This suggests that BSCCE remains the correct limiting object for stochastic gradient dynamics. However, we do not claim that the stochastic case is fully resolved here. A full stochastic-gradient analysis of auction games is an important next step, but it is beyond the scope of this paper.

5 Convergence Analysis

In this section, we characterize the size of the Equilibrium sets for first-price auctions with two players in the IPV model. Since we study discretized auctions which have an ordered action space, we are not only interested in how much probability can be shifted between two feasible equilibria, but also how far this probability can be shifted.

Metrics. Given a possibly mixed strategy $\sigma \in \Delta(\Theta \times \mathcal{A})$ and a pure strategy $\beta^* : \Theta \to \mathcal{A}$, their squared Wasserstein-2 distance is

$$W_2^2(\sigma, \beta^*) = \frac{1}{|\mathcal{N}|}\sum_{\theta,a} \sigma(\theta, a)||a - \beta^*(\theta)||_2^2.$$

In the following, we will simply refer to this metric as the Wasserstein distance. We consider the squared Wasserstein-2 distance between the buyers' marginal bid distributions and the equilibrium as a suitable objective for our analysis: This metric captures deviations in the probability mass assigned to different type–action pairs and weighs these based on the distance to the canonical equilibrium β^*. Furthermore, it provides a natural upper bound on the size of equilibrium polytopes ($\mathcal{P}$):

$$\varnothing_{\mathcal{E}}(\mathcal{P}) := \sup_{\sigma_1,\sigma_2\in\mathcal{P}} \sqrt{W_2^2(\beta_1, \beta_2)} \leq 2 \sup_{\sigma\in\mathcal{P}} \sqrt{W_2^2(\sigma, \beta^*)}.$$

As shown by Ahunbay & Bichler (2024), the BCCE in first-price auctions is unique whenever the prior cumulative distribution function (CDF) is strictly concave. As $\mathrm{BSCCE} \subseteq \mathrm{BCCE}$, their contraction result immediately carries over to BSCCE. Thus, the theoretically interesting case for our analysis is when the conditions guaranteeing BCCE contraction are violated.

Their numerical experiments indicate that even under a uniform prior, which does not satisfy strict concavity, the BCCE that maximizes Wasserstein distance to the canonical equilibrium approaches that equilibrium as the discretization is refined. At the same time, they provide counterexamples in which this distance increases when the prior is no longer concave. In particular, such behavior arises for $\mathrm{Beta}(\alpha, 1)$ priors, whose CDFs are strictly convex for $\alpha > 1$.

These findings motivate a closer investigation of whether BSCCE provides a sharper prediction in precisely these regimes. Specifically, we examine whether the Wasserstein distance between BSCCEs and the canonical equilibrium decreases as the discretization becomes finer even when the prior CDF is convex and BCCE need not concentrate.

Discretization. We discretize the type and action space uniformly as $\{0, 1/\xi, \dots, 1\}$, with a size of $\xi + 1$. The probability of a discrete type $\theta_i = \frac{k}{\xi}$ is $F(\frac{k+1}{\xi+1}) - F(\frac{k}{\xi+1})$. For a $\mathrm{Beta}(\alpha, 1)$ prior, the resulting type probabilities become increasingly imbalanced as α grows. In particular, the prior mass concentrates near the highest valuation, leading to $\pi_i(1) \gg \pi_i(0)$ for large α. Similarly, the probability mass concentrates around $\frac{1}{2}$ for a $\mathrm{Beta}(\alpha, \alpha)$ prior. This imbalance induces numerical instabilities in equilibrium computation and makes experiments for larger discretizations challenging.

Results. As the discretization refines, the BSCCE polytope contracts. In particular, its diameter, bounded by the maximal Wasserstein distance to the canonical equilibrium, converges to zero, indicating that all BSCCEs approach the unique continuous BNE. This contraction implies that any equilibrium selected by gradient-based learning dynamics in discretized auctions must lie close to the canonical equilibrium, providing strong numerical evidence for convergence in the continuous limit. The decrease in squared Wasserstein distance is illustrated in Figure 2 for a $\mathrm{Beta}(2, 1)$ prior. In this case, the Wasserstein distance decreases roughly as $\frac{1}{10\xi}$.

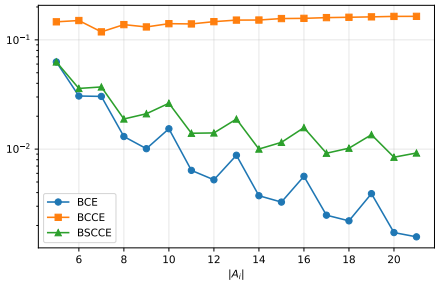


Figure 2: Distance to the BNE depending on ξ for a $\mathrm{Beta}(2, 1)$ prior.

Figure 1 from our introduction visualizes the maximal conditional probability assigned to each type–action pair under different solution concepts for this prior. Specifically, each tile (θ_i, a_i) displays $\max_{\sigma \in \mathcal{P}} \sigma(a_i \mid \theta_i)$. The figures illustrate the inclusion relationship $BNE \subseteq BCE \subseteq BSCCE \subseteq BCCE$. While all solution concepts are centered around the canonical equilibrium $\beta_i^*(\theta_i) = \frac{2}{3}\theta_i$, both BSCCE and BCE are substantially more concentrated around β_i^* than BCCE.

The speed of the concentration depends on the priors (see Appendix C). When the game is close to a complete-information game, or the CDF is too convex, the size of the BSCCE polytope shrinks slowly. We observe that the BSCCE is more robust to the non-concavity of the CDF than the BCCE, but for large values of α (where the convexity strengthens), it is not able to track the BCE convergence.

Complexity. An implementation of the general BSCCE linear program for a discretization of $1/\xi$ involves $O(\xi^4)$ variables and $O(\xi^3)$ constraints, with a relatively dense constraint matrix. Several model reductions are available when players are symmetric, and additional reductions apply to first-price auctions, where rational players never bid above their type. This allows us to reduce the number of non-zeros in the LP by approximately an order of magnitude. Even after these optimizations and despite the polynomial scaling, the LPs become challenging to solve for very fine discretizations.

For a discretization of $\xi = 20$, the BSCCE LP contains approximately 45k variables, 8k constraints, one million non-zeros, and requires about 20 seconds to compute the Wasserstein-distance-maximizing BSCCE using the Gurobi Optimizer version 12.1. By comparison, without model reductions, the same instance requires 30 minutes. For $\xi = 30$, the BSCCE LP grows to roughly 180k variables, 25k constraints, and 7 million non-zeros, with a solution time of about 11 minutes.

While we can ensure a theoretically well-behaved complexity bound, this does not always imply practical tractability. Due to the increasingly numerical instability and increasing computational burden, we did not pursue even finer discretizations. We still believe that our study here is informative, as it points towards a theoretically justified convergence of gradient-based algorithms.

6 Limitations

Our results have some limitations. First, while BSCCE is defined for general finite Bayesian games, our numerical analysis focuses on discretized symmetric first-price auctions. In particular, the contraction of BSCCE toward the continuous-game BNE outside regimes already covered by existing BCCE results is supported numerically rather than established theoretically. When existing results

imply that BCCE contracts to the BNE, the inclusion $\mathrm{BCE} \subseteq \mathrm{BSCCE} \subseteq \mathrm{BCCE}$ immediately yields the same conclusion for BSCCE. Beyond these regimes, however, deriving conditions under which BSCCE contracts to the BNE remains an open theoretical question.

Second, our formal learning result concerns time averages of deterministic Bayesian projected gradient ascent using expected-utility gradients. It therefore does not establish convergence guarantees for stochastic gradient updates based on sampled auction outcomes. While stochastic-approximation arguments motivate a connection between these settings, a formal treatment of stochastic BPGA requires additional assumptions and is beyond the scope of this paper.

Third, our numerical evidence is necessarily limited by the computational cost of optimizing over the BSCCE polytope. This constrains the discretization sizes that can be studied. Moreover, uniform discretizations can generate grid-alignment effects because the continuous BNE may lie closer to the action grid for some grid sizes than for others. We therefore interpret individual finite-grid results with caution and focus on the behavior across increasingly fine discretizations.

Finally, the auction experiments considered here use symmetric priors. Extending the analysis to asymmetric priors is nontrivial because the corresponding equilibrium is asymmetric and generally does not admit the same closed-form benchmark. Similarly, although the definition of BSCCE and its connection to BPGA apply to finite Bayesian games more generally, whether BSCCE provides a comparably sharp equilibrium prediction in other Bayesian games remains an open empirical and theoretical question.

7 Conclusion

This paper revisits a recurring empirical observation in auction learning: gradient-based algorithms applied to discretized Bayesian auctions reliably converge to outcomes that are close to the unique Bayes–Nash equilibrium (BNE) of the underlying continuous game, even in settings where standard theoretical guarantees fail. While existing no-regret arguments based on Bayesian coarse correlated equilibrium (BCCE) explain convergence in all-pay auctions, they are insufficient for first-price auctions, where the BCCE set can remain large and include outcomes far from the BNE.

To address this gap, we introduced the *Bayesian semicoarse correlated equilibrium* (BSCCE), a refinement of BCCE that restricts admissible deviations to tangential, probability-preserving transformations consistent with projected gradient dynamics. This equilibrium concept admits a linear programming characterization that enables systematic numerical analysis. While the focus of this paper is first-price auctions with independent private values, the solution concept BSCCE is defined for all Bayesian games.

Our main finding is that, in symmetric first-price auctions, the BSCCE set contracts as the discretization is refined. Across a range of priors, including cases where BCCE fails to concentrate, the diameter of the BSCCE polytope shrinks to zero, indicating convergence to the unique continuous BNE. This shrinkage property provides a stability certificate for gradient-based learning dynamics. However, shrinkage is limited when the priors are degenerate, i.e., when the game is close to a complete-information game, or when the CDF is highly convex. While the proposed equilibrium concept, BSCCE, sharpens predictions about the outcome of gradient dynamics, it remains open whether there is an even smaller, tractable solution concept that captures that outcome better.

Overall, our results suggest that equilibrium concepts tailored to the structure of learning dynamics can offer sharper and more robust predictions than generic no-regret notions in Bayesian games. The BSCCE framework provides a principled bridge between gradient-based optimization and equilibrium analysis, and it is applicable beyond auctions to other multi-agent learning problems with private information and high-dimensional strategy spaces.

Acknowledgments and Disclosure of Funding

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).

This project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - GRK 2201/2 - Project Number 277991500.

The work of Mete Şeref Ahunbay was funded by a Walter Benjamin Fellowship from the German Research Foundation (Deutsche Forschungsgemeinschaft), project nr. 558988238.

This project has received funding from the French National Research Agency (ANR) through the PEPR IA FOUNDRY project (ANR-23-PEIA-0003).

References

Ahunbay, M. and Bichler, M. On the uniqueness of bayesian coarse correlated equilibria in standard first-price and all-pay auctions, 2024. URL https://arxiv.org/abs/2401.01185.

Ahunbay, M. and Bichler, M. Semicoarse correlated equilibria and lp-based guarantees for gradient dynamics in normal-form games, 2025. URL https://arxiv.org/abs/2502.20466.

Ahunbay, Mete. First-order (coarse) correlated equilibria in non-concave games, 2025.

Benveniste, A., Metivier, M., and Priouret, P. *Adaptive Algorithms and Stochastic Approximations*. 2012.

Bergemann, D., Brooks, B., and Morris, S. First-Price Auctions With General Information Structures: Implications for Bidding and Revenue. *Econometrica*, 85(1):107–143, 2017. doi: 10.3982/ECTA13958.

Bichler, M., Fichtl, M., Heidekrüger, S., Kohring, N., and Sutterer, P. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Machine Intelligence*, 3(8):687–695, 2021. doi: 10.1038/s42256-021-00365-4.

Bichler, M., Fichtl, M., and Oberlechner, M. Computing bayes–nash equilibrium strategies in auction games via simultaneous online dual averaging. *Operations Research*, 2023.

Bichler, M., Lunowa, S. B., Oberlechner, M., Pieroth, F. R., and Wohlmuth, B. On the Convergence of Learning Algorithms in Bayesian Auction Games, 2024.

Bichler, M., Lunowa, S. B., Oberlechner, M., Pieroth, F. R., and Wohlmuth, B. Beyond monotonicity: On the convergence of learning algorithms in standard auction games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 13649–13657, 2025.

Cai, Y., Farina, G., Grand-Clément, J., Kroer, C., Lee, C.-W., Luo, H., and Zheng, W. Fast Last-Iterate Convergence of Learning in Games Requires Forgetful Algorithms. *Advances in Neural Information Processing Systems*, 37:23406–23434, 2024.

Chawla, S. and Hartline, J. D. Auctions with unique equilibria. EC '13, 2013. URL https://doi.org/10.1145/2482540.2483188.

Daskalakis, C., Farina, G., Fishelson, M., Pipis, C., and Schneider, J. Efficient learning and computation of linear correlated equilibrium in general convex games. STOC '25, 2025. URL https://doi.org/10.1145/3717823.3718307.

Forges, F. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35(3):277–310, 1993. doi: 10.1007/BF01075202.

Gordon, G. J., Greenwald, A., and Marks, C. No-regret learning in convex games. In *Proceedings of the 25th international conference on Machine learning*, pp. 360–367. Association for Computing Machinery, 2008. doi: 10.1145/1390156.1390202.

Greenwald, A. and Jafari, A. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Proceedings of the 16th Annual Conference on Learning Theory (COLT)*, pp. 2–12, 2003.

Hartline, J., Syrgkanis, V., and Tardos, E. No-regret learning in bayesian games. *Advances in Neural Information Processing Systems*, 28, 2015.

Lugosi, G. and Stoltz, G. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1):187–208, 2007. doi: 10.1016/j.geb.2006.06.001.

Mertikopoulos, P. and Sandholm, W. H. Learning in Games via Reinforcement and Regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016. doi: 10.1287/moor.2016.0778.

Mertikopoulos, P. and Sandholm, W. H. Riemannian game dynamics. *Journal of Economic Theory*, 177: 315–364, 2018. doi: 10.1016/j.jet.2018.06.002.

Mertikopoulos, P. and Zhou, Z. Learning in games with continuous action sets and unknown payoff functions, 2018.

Mertikopoulos, P., Hsieh, Y.-P., and Cevher, V. A unified stochastic approximation framework for learning in games. *Mathematical Programming*, 203(1):559–609, 2024. doi: 10.1007/s10107-023-02001-y.

Rasooly, I. and Gavidia-Calderon, C. The importance of being discrete: on the inaccuracy of continuous approximations in auction theory, 2022. URL `https://arxiv.org/abs/2006.03016`.

Zhang, D. and Nagurney, A. On the stability of projected dynamical systems. *Journal of Optimization Theory and Applications*, 85(1):97–124, 1995. doi: 10.1007/BF02192301.

A LP Formulations

A.1 BSCCE LP

We state BSCCEs as a linear program over the joint distribution $\sigma \in \Delta(\Theta \times \mathcal{A})$ where the constraints enforce the conditions in Definition 3.2 (prior consistency, conditional independence, incentive compatibility). Note that $\pi(\theta)$ represents the probability density of the type distribution.

For the conditional independence, we write it explicitly in terms of the joint distribution and already assume prior consistency

$$\sigma(\theta_{-i}|\theta_i, a_i) = \sigma(\theta_{-i}|\theta_i) \Leftrightarrow \frac{\sigma(\theta, a_i)}{\sigma(\theta_i, a_i)} = \frac{\sigma(\theta)}{\sigma(\theta_i)} = \frac{\pi(\theta)}{\pi(\theta_i)} \Leftrightarrow \pi(\theta_i)\sigma(\theta, a_i) = \pi(\theta)\sigma(\theta_i, a_i)$$

Besides Prior Consistency

$$\sum_a \sigma(\theta, a) = \pi(\theta)$$

it must hold for each player $i \in \mathcal{N}$ and all corresponding types $\theta_i \in \Theta_i$ that

$$\pi(\theta_i) \sum_{a_{-i}} \sigma(\theta, a_i, a_{-i}) = \pi(\theta) \sum_{\theta'_{-i}, a_{-i}} \sigma(\theta_i, \theta'_{-i}, a_i, a_{-i}) \quad \forall\, a_i, \theta_{-i},$$

(Conditional Independence)

$$\sum_{\theta_{-i}, a} \sigma(\theta, a) \left[\sum_{a'_i \in \mathcal{A}_i} d_i(\theta_i, a'_i, a_i) u_i(a'_i, a_{-i}|\theta_i) \right] \leq 0 \qquad \text{(Semicoarse incentive compatibility)}$$

$$\forall\, d_i(\theta_i, a'_i, a_i) \in \mathcal{D}_i.$$

The suitable BSCCE deviations $d_i(\theta_i, a'_i, a_i) = Q_i(\theta_i, a'_i, a_i) + q_i(\theta_i, a'_i) \in \mathcal{D}_i$ are defined as follows:
for all players $i \in \mathcal{N}$ and for all corresponding types $\theta_i \in \Theta_i$ it must hold

$$\begin{aligned}
\sum_{a'_i \in \mathcal{A}_i} d_i(\theta_i, a'_i, a_i) = 0 \qquad & \forall a_i, && \text{(probability preservation, } \mu_i(\theta_i, a_i)) \\
d_i(\theta_i, a'_i, a_i) \geq 0 \qquad & \forall a_i \neq a'_i, && \text{(non-negativity, } \gamma_i(a_i, a'_i)) \\
Q_i(\theta_i, a'_i, a_i) - Q_i(\theta_i, a_i, a'_i) = 0 \qquad & \forall a_i, a'_i. && \text{(symmetry, } \rho_i(a_i, a'_i))
\end{aligned}$$

A.2 Tractable BSCCE LP

As there are infinitely many 'suitable' deviations, the primal before is numerically unsolvable. We use the approach from Theorem 3.4 in Ahunbay & Bichler (2025) to get a tractable version. The idea is to consider the 'worst-case' deviations, which correspond to the extreme rays of the deviation polyhedral cone, and incorporate them into the objective. This results in an inner and outer problem. Then we consider the dual of the inner maximization problem, with the dual variables μ, ρ and $\gamma \geq 0$ as introduced before. The dual is a feasibility problem as the right-hand sides are 0, and the constraints can be incorporated into the outer problem. As $\mu_i(\theta_i, a_i)$ will only appear in one equality constraint, it can be eliminated from the resulting LP by fixing its value suitably. The Semicoarse incentive compatibility constraint is then replaced by the following two types of tractable constraints (for each player and type)

$$\begin{aligned}
w_i(\theta_i, a_i, a'_i) - w_i(\theta_i, a_i, a_i) + \gamma_i(\theta_i, a_i, a'_i) + \rho_i(\theta_i, a_i, a'_i) - \rho_i(\theta_i, a'_i, a_i) = 0 \quad & \forall a_i \neq a'_i \\
\sum_{a_i} [w_i(\theta_i, a_i, a'_i) - w_i(\theta_i, a_i, a_i)] + \sum_{a_i \neq a'_i} \gamma_i(\theta_i, a_i, a'_i) = 0 \quad & \forall a'_i,
\end{aligned}$$

where

$$w_i(\theta_i, a_i, a'_i) = \sum_{\theta_{-i}, a_{-i}} \sigma(\theta, a) u_i(a'_i, a_{-i}|\theta_i).$$

Note that if $\rho = 0$, then the constraints reduce to the BCE incentive constraints.

If all players have identical type and action spaces, i.e. $|\Theta_i| = |\Theta_j|$ and $|\mathcal{A}_i| = |\mathcal{A}_j|$ for all $i \neq j \in \mathcal{N}$, the resulting LP has $O\big((|\Theta_i|\,|\mathcal{A}_i|)^{|\mathcal{N}|}\big)$ variables and $O\big(|\mathcal{N}|\,|\Theta_i|^{|\mathcal{N}|}\,|\mathcal{A}_i| + |\mathcal{N}|\,|\Theta_i|\,|\mathcal{A}_i|^2\big)$ constraints, where the first term comes from the prior-consistency and conditional-independence constraints, and the second term from the tractable semicoarse incentive-compatibility constraints.

B Proof of Corollary 4.1

Proof. We show that the distribution

$$\sigma_T(\theta, a) := \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j)$$

satisfies the constraints in Definition 3.2.

Prior consistency. We must show that the marginal of σ_T over actions equals the prior π: $\sigma_T(\theta) = \pi(\theta)$. It holds that,

$$\begin{aligned}\sigma_T(\theta) = \sum_{a\in A} \sigma_T(\theta, a) &= \sum_{a\in A} \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j) \\ &= \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \sum_{a\in A} \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j) = \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) = \pi(\theta).\end{aligned}$$

The second-to last step follows, as $x_{j,t}(\cdot \mid \theta_j)$ are a probability distribution over A, hence the inner sum equals 1.

Conditional independence. We first show that $\sigma_T(a_i \mid \theta_i) = \frac{1}{T}\sum_{t=1}^{T} x_{i,t}(a_i \mid \theta_i) \quad (*)$.
As a first step, we use Bayes' rule and the definition of σ_T. Afterwards, we rearrange terms to obtain

$$\begin{aligned}\sigma_T(a_i \mid \theta_i) = \sum_{a_{-i},\theta_{-i}} \frac{\sigma_T(a,\theta)}{\sigma_T(\theta_i)} &= \frac{1}{\sigma_T(\theta_i)} \cdot \sum_{a_{-i},\theta_{-i}} \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j) \\ &= \underbrace{\frac{\pi(\theta_i)}{\sigma_T(\theta_i)}}_{=1} \cdot \frac{1}{T}\sum_{t=1}^{T} x_{i,t}(a_i \mid \theta_i) \sum_{\theta_{-i}} \pi(\theta_{-i} \mid \theta_i) \sum_{a_{-i},} \prod_{j\neq i} x_{j,t}(a_j \mid \theta_j) \quad = \frac{1}{T}\sum_{t=1}^{T} x_{i,t}(a_i \mid \theta_i).\end{aligned}$$

The last step is a result of prior consistency and the insight that $\pi(\theta_{-i} \mid \theta_i)$ is a distribution over Θ_{-i} and $\prod_{j\neq i} x_j(\cdot \mid \theta_j)$ is a distribution over $\mathcal{A}_{-i}$, so they sum up to one.

We now apply this to show conditional independence.

$$\begin{aligned}\sigma_T(a_i, \theta) = \sum_{a_{-i}} \sigma_T(a, \theta) \quad &= \sum_{a_{-i}} \frac{1}{T}\sum_{t=1}^{T} \pi(\theta) \prod_{j\in\mathcal{N}} x_{j,t}(a_j \mid \theta_j) \\ &= \pi(\theta) \cdot \frac{1}{T}\sum_{t=1}^{T} x_i(a_i \mid \theta_i) \sum_{a_{-i}} \prod_{j\neq i} x_{j,t}(a_j \mid \theta_j) \\ &= \pi(\theta) \cdot \frac{1}{T}\sum_{t=1}^{T} x_i(a_i \mid \theta_i) \overset{(*)}{=} \pi(\theta) \cdot \sigma_T(a_i \mid \theta_i)\end{aligned}$$

Incentive compatibility. We need to show that for every player i, every type $\theta_i \in \Theta_i$, and every deviation $d^{\theta_i} \in \mathcal{D}$

$$\lim_{T\to\infty} \mathop{\mathbb{E}}_{(\theta,a)\sim\sigma_T} \left[\sum_{a_i'\in\mathcal{A}_i} d^{\theta_i}(a_i', a_i)\, u_i(a_i', a_{-i} \mid \theta_i) \middle| \theta_i \right] \leq 0.$$

Remember that the players play the strategy game in expected-utility formulation as described in Section 4. In particular, the ex-ante utility for playing a vector $x_i \in X_i := \prod_{\theta_i\in\Theta_i} \Delta(\mathcal{A}_i)$ is

$$u_i(x_i, x_{-i}) = \mathbb{E}_{\theta\sim\pi, a\sim x(\cdot|\theta)}[u_i(a_i, a_{-i})] = \sum_{\theta,a} \left(\pi(\theta) \cdot \prod_{j\in\mathcal{N}} x_j(a_j \mid \theta_j) \right) u_i(a_i, a_{-i} \mid \theta_i).$$

This is essentially a complete-information game, and we can apply Proposition 2.1 from Ahunbay & Bichler (2025). The proposition states that, for any tangential $h : X \to \mathbb{R}$ with bounded, Lipschitz continuous gradients, the following inequality holds, assuming all players implement projected gradient ascent with suitably chosen step sizes.

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{N}}\langle\nabla_i h(x_t), \nabla_i u_i(x_t)\rangle \leq \frac{1}{T}\cdot\left(\frac{1}{\eta_T}+\frac{1}{\eta_1}+\sum_{t=1}^{T}\eta_t\right)\cdot const \quad \overset{T\to\infty}{\longrightarrow} 0$$

The constant depends on the game and the function h. We choose

$$h(x) = \sum_i \sum_{\theta_i} \frac{1}{\pi(\theta_i)}\left(\frac{1}{2}x_i[\theta_i]^T Q_i^{\theta_i} x_i[\theta_i] + x_i[\theta_i]^T q_i^{\theta_i}\right)$$

with pairs $(Q_i^{\theta_i}, q_i^{\theta_i})$ that satisfy the tangency, non-negativity, and symmetry conditions. $x_i[\theta_i] \in \Delta(\mathcal{A}_i)$ is the probability vector that player i assigns upon observing θ_i. With this h, we now show that BSCCE incentive compatibility follows. We can assume that $\pi(\theta_i) > 0$ without loss of generality because if this were not the case, we could remove the type from the game.

Let us first describe the gradients $\nabla_i h(x)$ and $\nabla_i u_i(x)$. They both are vectors in $\mathbb{R}^{|\Theta_i|\cdot|\mathcal{A}_i|}$, and, like the strategy vectors x_i, they have a block-wise structure: There are $|\Theta_i|$ blocks with $|\mathcal{A}_i|$ entries, corresponding to the gradients with respect to $x_i(a_i \mid \theta_i)$. In particular, we find

$$\begin{aligned}
\nabla_i h(x)[\theta_i, a_i'] &= \frac{1}{\pi(\theta_i)}\left(\sum_{a_i} Q_i^{\theta_i}(a_i', a_i) x_i[\theta_i, a_i]\right) + \frac{1}{\pi(\theta_i)} q_i^{\theta_i}(a_i') \\
\nabla_i u_i(x)[\theta_i, a_i'] &= u_i(a_i', x_{-i} \mid \theta_i)\cdot\pi(\theta_i),
\end{aligned}$$

where the indices refer to the position in the vector corresponding to the respective type and action. The utility $u_i(a_i', x_{-i} \mid \theta_i)$ is the interim utility of a player that observes type θ_i and plays action a_i' while the opponents commit to playing x_{-i}:

$$\begin{aligned}
u_i(a_i', x_{-i} \mid \theta_i) &= \mathbb{E}_{\theta_{-i}\sim\pi(\cdot|\theta_i)}\left[\mathbb{E}_{a_{-i}\sim x_{-i}(\cdot|\theta_{-i})}[u_i(a_i', a_{-i} \mid \theta_i)]\right] \\
&= \sum_{a_{-i},\theta_{-i}}\left(\pi(\theta_{-i} \mid \theta_i)\prod_{j\neq i} x_j(a_j \mid \theta_j)\right) u_i(a_i', a_{-i} \mid \theta_i).
\end{aligned}$$

In what follows, we derive the equality

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i\in\mathcal{N}}\langle\nabla_i h(x), \nabla_i u_i(x)\rangle = \sum_{i\in\mathcal{N}}\sum_{\theta_i\in\Theta_i}\mathbb{E}_{(\theta,a)\sim\sigma_T}\left[\sum_{a_i'} d^{\theta_i}(a_i', a_i) u_i(a_i', a_{-i} \mid \theta_i) \mid \theta_i\right],$$

which allows us to apply the bound from the proposition to our distribution σ_T.

We start with the scalar product of $\nabla_i u_i$ and $\nabla_i h$:

$$\begin{aligned}
\langle\nabla_i u_i(x_t), \nabla_i h(x_t)\rangle &= \sum_{a_i', \theta_i} \nabla_i u_i(x_t)[\theta_i, a_i'] \cdot \nabla_i h(x_t)[\theta_i, a_i'] \\
&= \sum_{a_i', \theta_i} u_i(a_i', x_{-i,t} \mid \theta_i)\cdot\left(\left(\sum_{a_i} Q_i^{\theta_i}(a_i', a_i)\cdot x_{i,t}(a_i \mid \theta_i)\right) + q_i^{\theta_i}(a_i')\right) \\
&= \sum_{a_i', \theta_i} u_i(a_i', x_{-i,t} \mid \theta_i)\sum_{a_i}\left(Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i')\right)\cdot x_{i,t}(a_i \mid \theta_i)
\end{aligned}$$

In the second step, we used our gradients and note that the $\pi(\theta_i)$ terms cancel. We also moved from the vector-based notation for the players' strategies to a conditional-probability notation. In the second step, we applied the identity $\sum_{a_i} x_{i,t}[\theta_i, a_i] = 1$.

We expand the interim utility terms to obtain

$$= \sum_{a_i', \theta_i} \underbrace{\sum_{a_{-i}, \theta_{-i}} \left(\pi(\theta_{-i} \mid \theta_i) \prod_{j \neq i} x_{j,t}(a_j \mid \theta_j) \right) \cdot u_i(a_i', a_{-i} \mid \theta_i)}_{u_i(a_i', x_{-i,t} \mid \theta_i)} \cdot \sum_{a_i} \left(Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i') \right) \cdot x_{i,t}(a_i \mid \theta_i)$$

$$= \sum_{a_i', \theta, a} \left(\pi(\theta_{-i} \mid \theta_i) \prod_{j \in \mathcal{N}} x_{j,t}(a_j \mid \theta_j) \right) \cdot u_i(a_i', a_{-i} \mid \theta_i) \cdot \left(Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i') \right),$$

where we moved the order of summation and regrouped the player strategies $x_{j,t}$ to one product.

When we average over time, we can identify $\frac{1}{T} \sum_{t=1}^{T} \pi(\theta_{-i} \mid \theta_i) \prod_{j \in \mathcal{N}} x_{j,t}(a_j \mid \theta_j)$ with $\sigma_T(a, \theta_{-i} \mid \theta_i)$:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T} \langle \nabla_i u_i(x_t), \nabla_i h(x_t), \rangle &= \sum_{a_i', a, \theta} \sigma_T(a, \theta_{-i} \mid \theta_i) \cdot u_i(a_i', a_{-i} \mid \theta_i) \cdot \left(Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i') \right) \\ &= \sum_{a_i', \theta_i} \mathbb{E}_{(a,\theta) \sim \sigma_T} \left[u_i(a_i', a_{-i} \mid \theta_i) \cdot (Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i')) \mid \theta_i \right]. \end{aligned}$$

By summing over all players, this establishes the equivalence, so we get

$$\limsup_{T \to \infty} \sum_{i \in \mathcal{N}} \sum_{a_i', \theta_i} \mathbb{E}_{(a,\theta) \sim \sigma_T} \left[u_i(a_i', a_{-i} \mid \theta_i) \cdot (Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i') \mid \theta_i \right] \leq 0.$$

Note that this holds for any choice of matrices $Q_i^{\theta_i}$ and vectors $q_i^{\theta_i}$ ($i \in \mathcal{N}$, $\theta_i \in \Theta_i$) that satisfy symmetry, tangency, and non-negativity. We can thus set all of them to zero aside for one player i and one of its types θ_i. This implies

$$\limsup_{T \to \infty} \mathbb{E}_{(a,\theta) \sim \sigma_T} \left[u_i(a_i', a_{-i} \mid \theta_i) \cdot (Q_i^{\theta_i}(a_i', a_i) + q_i^{\theta_i}(a_i') \mid \theta_i \right] \leq 0, \qquad \forall i \in \mathcal{N}, \theta_i \in \Theta_i,$$

which in turn concludes the proof.

□

C Numerical experiments

C.1 Hardware Information

We ran our experiments on a small-scale server with an Intel(R) Xeon(R) Silver 4116 CPU with 48 cores (2.10 GHz) and 192 GiB of RAM. Our experiments are based on Python 3.12 and Gurobi 12.1. We left the CPU usage to Gurobi. Each experiment took between several seconds and several minutes (see "Complexity" in Section 5).

C.2 Wasserstein Distances in the Incomplete-Information game

The convergence curves for the squared Wasserstein-2 distance exhibit a periodic "zig-zag" pattern for the uniform prior (Figure 3a). This is an artifact of the uniform discretization grid aligning perfectly with the canonical continuous BNE actions ($1/2\theta_i$). Thus, the oscillation is a predictable effect rather than an algorithmic or numerical instability.

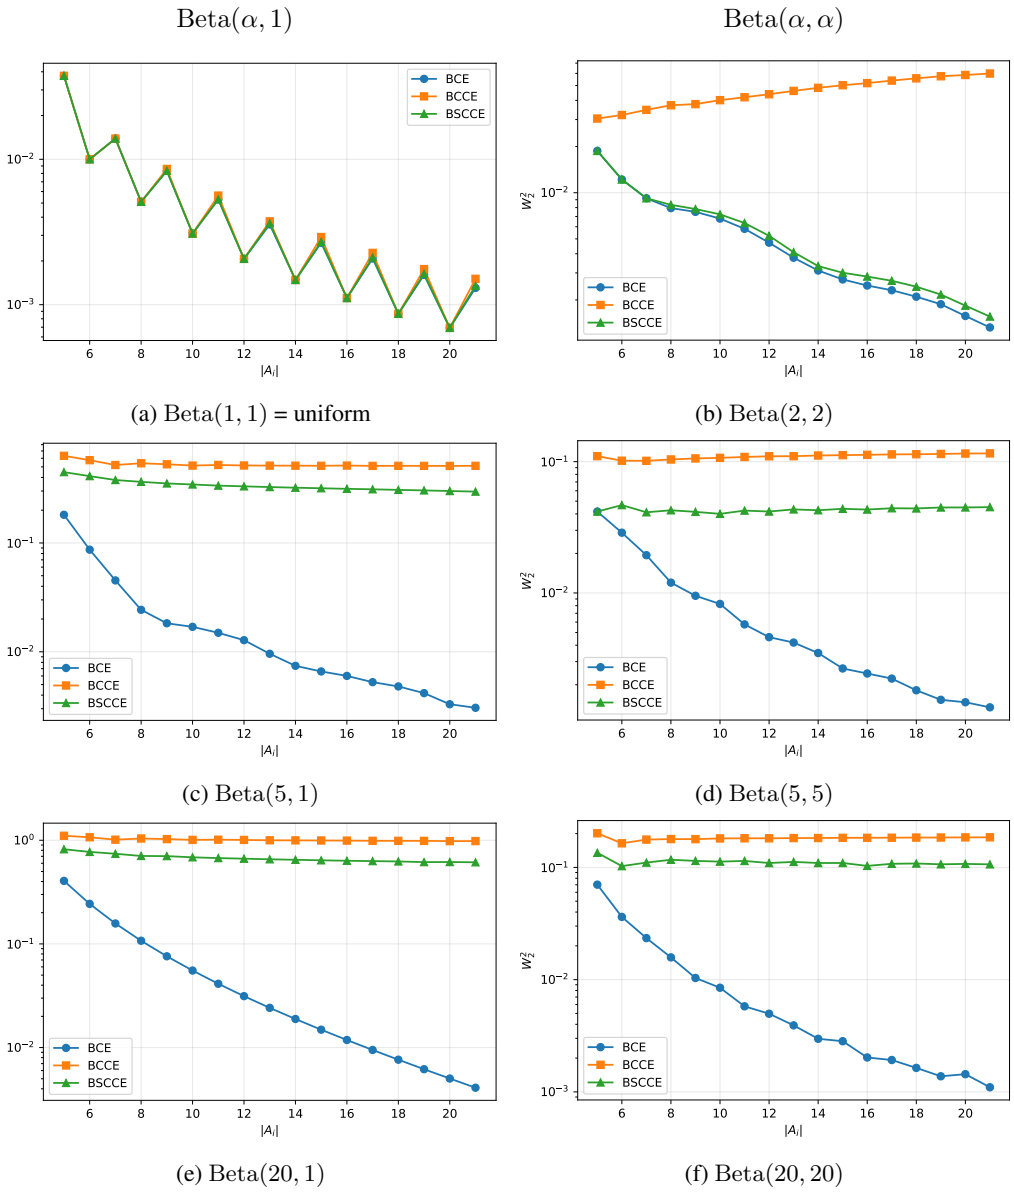
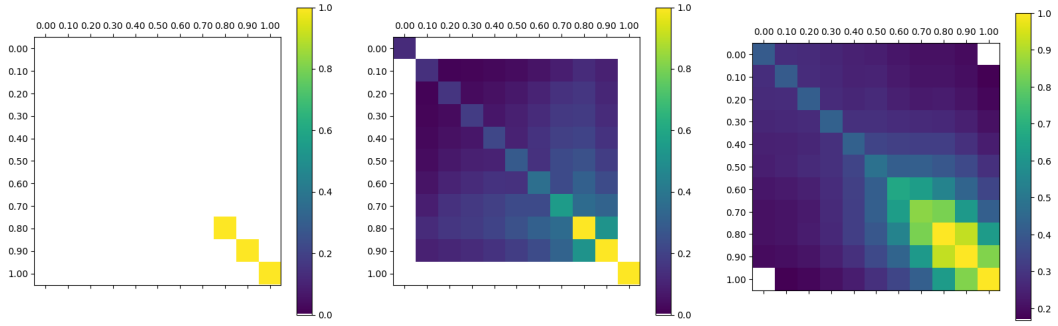


Figure 3: Dependence of the distance to the canonical BNE on the fineness of uniform discretization for various Beta priors.

C.3 Equilibrium Polytopes in the Complete-Information Case

As the previous results suggest that the tightness of the BSCCE concept decreases as priors become degenerate, we investigate the limiting complete-information case. Figure 4 illustrates the different equilibrium polytopes for a first-price auction with two players, who both have type $\theta_i = 1$. In this case, the only possible CEs are that both players perfectly coordinate to play the exact same action out of the top three, whereas the SCCE and CCE sets are much larger. This implies that, although the SCCE set is smaller than the CCE set, an even stronger equilibrium refinement might be necessary.



(a) CE polytope (b) SCCE polytope (c) CCE polytope

Figure 4: Comparison of equilibrium polytopes for a first price auction where both players have type $\theta_i = 1$.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We provide a theoretical proof for the relation to gradient-based learners and empirical results that justify the convergence claims.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We discuss computational complexity and limited convergence for degenerate priors.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution

is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.

- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: We clearly lay out the setting used for Corollary 4.1, provide a proof in the appendix, and reference prior work.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We describe the first-price sealed-bid auction and how we discretize it. The parameterizations of the prior distributions are provided in the respective figures. In Appendix Appendix A.2, we detail how the BSCCE problem can be made tractable.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The corresponding Code repository will be made publicly available. Currently, it is not included to preserve authors' anonymity.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The experiments are based on auction games. We describe in the main part, including how we discretized and parameterized the auctions. No external datasets are used.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: The experiments are deterministic: our results are based on solutions of linear programs.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Our manuscript includes a paragraph on complexity which also includes runtimes. Details about the hardware are reported in Appendix C.1.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: We have reviewed the NeurIPS Code of Ethics. Our paper does not violate the principles described.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: Our paper adds to an understanding of learning dynamics in games; there is no immediate societal impact.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: We do not publish any sensitive data or models.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: We do not build on any previous data, code, or model.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: We do not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not include experiments with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not include experiments with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: Our method does not use LLMs.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.