Computing Distributional Bayes Nash Equilibria in Auction Games via Gradient Dynamics

Maximilian Fichtl, Matthias Oberlechner, Martin Bichler

Department of Informatics, Technical University of Munich
bichler@in.tum.de

Abstract

Auctions are modeled as Bayesian games with continuous type and action spaces. Computing equilibria in auction games is computationally hard in general and no exact solution theory is known. We introduce algorithms computing distributional strategies on a discretized version of the game via convex online optimization. One advantage is that the expected utility of agents is linear in distributional strategies. It follows that if our regularized optimization algorithms converge to a pure strategy, then they converge to a ε-equilibrium of the discretized game. We also show that the ε-equilibrium of the discretized game approximates an equilibrium in the continuous game. In a number of experiments, we show that the method approximates the analytical (pure) Bayes Nash equilibrium closely in a wide variety of auction games. This is remarkable, because in many games learning dynamics do not converge or are even chaotic. When agents have a low number of strategies or they are symmetric, we find equilibria in seconds. The method allows for interdependent valuations and different types of utility functions and provides a foundation for broadly applicable equilibrium solvers that can push the boundaries of equilibrium analysis in auctions and beyond.

Introduction

Auction games are arguably some of the most important applications of game theory and they are modeled as continuous-type, continuous-action Bayesian games. A bidder's valuation or type in such an auction game is drawn from some continuous distribution and he can choose from a continuous range of possible actions (or bids). Early on, Vickrey (1961) showed how to derive a Bayes-Nash equilibrium (BNE) strategy in a single-object first-price auction in the independent-private values (IPV) model with symmetric bidders and quasi-linear utility functions. The first-order conditions together with the assumption of symmetric bidding behavior lead to an ordinary differential equation, which has a closed-form solution for the BNE bidding strategy.

It turns out that deviations from this benchmark model lead to challenges in the equilibrium analysis (McAfee and McMillan 1987). For example, when the valuations of potential bidders are interdependent, then the system of first-order partial differential equations that characterizes a BNE often becomes intractable (Campo, Perrigne, and Vuong 2003). Computing Nash equilibria (NE) in complete-information finite games is already known to be PPAD-hard. However, computing exact Bayesian Nash equilibria (BNE) can even be PP-hard, a complexity class that is clearly intractable Cai and Papadimitriou (2014). Consequently, the analytical derivation of BNE strategies has been elusive for all but very simple auction games. For example, no BNE strategies have been derived for first-price sealed-bid combinatorial auctions, even though these auctions are widely used in procurement or also for spectrum sales. Even existence of BNE has only been shown for relatively simple models (Jackson and Swinkels 2005).

Inspite of this, there have been a number of approaches to develop numerical techniques for specific environments. For example, Armantier, Florens, and Richard (2008) introduced a BNE-computation method that is based on expressing the Bayesian game as the limit of a sequence of complete-information games. Rabinovich et al. (2013) study best-response dynamics on mixed strategies in auctions with finite action spaces, while Bosshard et al. (2020) contribute an iterated best-response algorithm with an elaborate verification method. More recently, Bichler et al. (2021) introduced a technique to compute approximate Bayes-Nash equilibria (BNE) using neural networks and self-play. Rather than iterated best-response, the Neural Pseudogradient Ascent (NPGA) implements simultaneous gradient ascent to learn pure BNE in auction games. Depending on the prior distribution, the number of items and bidders and their utility function, close approximations could be found within a few minutes or hours in a wide range of auction games. While NPGA implements gradient dynamics, the use of neural networks and evolutionary strategies leads to a relatively complex algorithm. However, the convergence to approximate BNE in a wide variety of auction environments is remarkable given a number of recent results on matrix games, where gradient dynamics either circle, diverge, or are even chaotic (Sanders, Farmer, and Galla 2018).

We introduce *Simultaneous Online Dual Averaging* (SODA), a form of gradient dynamics that learns *distributional strategies* (Milgrom and Weber 1985), which are a form of mixed strategies for Bayesian games. The distributional strategies allow us to derive gradients and implement gradient dynamics without relying on neural networks with self-play. SODA is based on a discretization of the type and action space and allows for interdependent types and different utility functions (e.g., risk aversion), which makes it very simple and generic algorithm compared to existing approaches. Learning distributional strategies rather than pure strategies has two main advantages: First, distributional strategies are known to exist in a larger variety of environments compared to what is known about pure strategy BNE (Jackson and Swinkels 2005; Athey 2001; Reny 2011). Second, the expected utility is linear in the distributional strategies. It can be shown that if regularized convex optimization algorithms converge to a pure strategy, then they converge to a BNE of the discretized game. Understanding gradient dynamics in specific types of games turns out to be very challenging even in simple matrix games. Actually, it is akin to studying dynamical systems and characterizing environments where gradient dynamics converge to a Nash equilibrium (if one exists) can be arbitrarily complex (Andrade, Frongillo, and Piliouras 2021). However, for standard auction models we can quickly check if SODA converges to a pure strategy for a specific game and thus is an equilibrium. This is an advantage over prior numerical methods, which rely on estimates of the utility loss only to certify an approximate equilibrium. Importantly, we also show that the distributional ε-BNE found in the discretized single-object auctions approximates a continuous equilibrium, if one exists.[1]

Computational complexity is always a concern in equilibrium computation, and we know that approximating BNE in multi-item auctions can be NP-hard (Cai and Papadimitriou 2014). The main drivers of complexity for our algorithm are the number of bidders, the number of available strategies (typically driven by the number of items for sale) to each bidder and the level of discretization of the type and action space. If the bidder has an exponential set of strategies in the set of items as it can happen in a combinatorial auction, also our algorithm cannot be polynomial. Yet, in most auction models in the literature the number of strategies for each bidder is fixed and small or the bidders are symmetric, which allows for an effective computation of ε-BNE with SODA.

We provide extensive experimental results where we approximate the analytical pure BNE closely in a wide variety of auction games. We could actually compute close approximations of the BNE with only a few bidders in seconds even for complex core-selecting combinatorial auctions. If we restrict ourselves to independent private values, we can solve large instances with dozens of bidders in seconds. This allows for a quick exploration of auction models with different priors or different utility functions.

Numerical methods have shown to be important for the engineering sciences and are widely used. In auction theory and market design, numerical methods have not received a similar level of attention. SODA allows for the development of numerical tools for a wide range for Bayesian games with continuous type and action spaces and perform comparative statics with respect to distributional assumptions or the utility functions of the agents. This can be the foundation for widely applicable equilibrium solvers. Importantly, the paper shows that important cases of equilibrium computation problems in auctions are tractable and we can find approximate equilibria quickly in spite of discouraging general complexity results.

Model and Algorithm

We will first introduce the necessary notation and then describe the algorithm more generally.

Notation

An incomplete-information or *Bayesian game* is given by a sextuplet $G = (\mathcal{I}, \mathcal{V}, \mathcal{O}, \mathcal{A}, f, u)$. Here $\mathcal{I} = \{1, \ldots, n\}$ denotes the set of *agents* participating in the game. The joint probability density function $f : \mathcal{O} \times \mathcal{V} \to \mathbb{R}_{\geq 0}$ describes an atomless *prior* distribution over agents' types, given by tuples (o_i, v_i) of observations and valuations. We make no further restrictions on f, thus allowing for arbitrary correlations. f is assumed to be common knowledge and we will denote its marginals by f_{v_i}, f_{o_i}, etc.; its conditionals by $f_{v_i|o_i}$, etc.; and its associated probability measure by F. Agent i's private *observation* is then given as a realization $o_i \in \mathcal{O}_i$, with $\mathcal{O} = \mathcal{O}_1 \times \cdots \times \mathcal{O}_n$ being the set of possible observation profiles. Similarly, $\mathcal{V}$ denotes the set of "true" but possibly unobserved valuations. Crucially, we make this distinction to model interdependencies in settings beyond purely private values or purely common values. Based on the observation o_i, the agent chooses an action, or *bid*, $b_i \in \mathcal{A}_i$, and the set of possible action profiles is given by $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$.

For each possible action and valuation profile, the vector $u = (u_1, \ldots, u_n)$ of F-integrable, individual *(ex-post) utility* functions $u_i : \mathcal{A} \times \mathcal{V}_i \to \mathbb{R}$ assigns the game outcome to each player. *Ex-ante*, before the game, agents neither have observations nor valuations, only knowledge about f. In the *interim* stage, agents additionally observe o_i providing (possibly partial or noisy) information about their own valuations v_i. Full access to the outcomes $u(v, b)$ is given only after taking actions (*ex-post*). In our formulation, we do not assume explicit ex-post access to any values (e.g., v_i, v_{-i}, b_{-i}) beyond the outcome u itself. An index $-i$ denotes a partial profile of all agents but agent i.

Taking an ex-ante view, players are tasked with finding strategies that link observations and bids. Instead of pure strategies, which are measurable functions $\beta_i : \mathcal{O}_i \to \mathcal{A}_i$ that map observations to bids, we are interested in distributional strategies that induce a probability measure on the space of observations and actions (Milgrom and Weber 1985).

Definition 1. *In the private values model, a distributional strategy for player i is probability measure σ on $\mathcal{O}_i \times \mathcal{A}_i$ for which the marginal distribution on $\mathcal{O}_i$ is f_{o_i}. Formally, the marginal condition can be written as $\sigma(O \times \mathcal{A}_i) = F_{o_i}(O)$*

[1] Jackson and Swinkels (2005) discuss examples where there are equilibria in the discretized game, but the continuous game does not have an equilibrium.

for all measurable sets $O \subset \mathcal{O}_i$. When players adopt distributional strategies $(\sigma_1, ..., \sigma_n)$ the expected utility is given by

$$\tilde{u}_i(\sigma) = \int u_i(b, o_i)\sigma_1(db_1|o_i)...\sigma_n(db_n|o_n)F(do) \quad (1)$$

The primary Bayesian games we'll consider are *sealed-bid auctions* on m indivisible items. In general combinatorial auctions we thus have a set $\mathcal{K}$ of possible *bundles* of items and the valuation- and action-spaces are therefore of dimension $|\mathcal{K}| = 2^m$. In the *private values* setting, we always have $o_i = v_i$; in the *common values* setting, there is some unobserved constant $v_c = v_1 = \cdots = v_n$ and the o_i can be considered noisy measurements of v_c. Mixed settings are likewise possible. In any case, based on bid profile b, an *auction mechanism* will determine two things: An allocation $x = x(b) = (x_1, \dots x_n)$ which constitutes a partition of the m items, where bidder i is allocated the bundle x_i; and a price vector $p(b) \in \mathbb{R}^n$, where p_i is the monetary amount bidder i has to pay in order to receive x_i. Formally, one may consider the individual allocations to be one-hot-encoded vectors $x_i \in \{0,1\}^{|\mathcal{K}|}$. In the standard risk-neutral model the utilities u_i are then described by *quasilinear* payoff functions $u_i^{QL}(v_i, b) = (x_i(b) \cdot v_i - p_i(b))$, i. e. by how much a player values her allocated bundle minus the price she has to pay.

An extension to this basic setting includes *risk-aversion*. Here, we model risk-aversion via utilities $u^{RA} = \left(u^{QL}\right)^\rho$ where $\rho \in (0,1]$ is the risk attitude; $\rho = 1$ describes risk-neutrality, smaller values lead to strictly concave, risk-averse transformations of u^{QL}. Risk aversion is an established way to explain why in field studies of single-object first-price sealed-bid (FPSB) auctions, bidders bid higher than their risk-neutral counterparts in analytical BNE (Bichler, Guler, and Mayer 2015). However, different types of utility functions are possible.

Model and Algorithm

The algorithm is based on gradient dynamics applied to the set of discrete versions of the distributional strategies. As mentioned in the example, these are constructed by restricting ourselves to finite subsets of the observation, valuation and action sets and considering finitely atomic measures as a counterpart to the distributional strategies in the continuous setting. Formally speaking, we construct a discrete version $G^d = (\mathcal{I}, \mathcal{V}^d, \mathcal{O}^d, \mathcal{A}^d, f^d, u)$ of the incomplete-information game G. This is done by defining a set of discrete observations $\mathcal{O}^d = \mathcal{O}_1^d \times ... \times O_n^d$ where $\mathcal{O}_i^d := \{o_1^i, ..., o_N^i\} \subset \mathcal{O}_i$. Similarly we define $\mathcal{A}_i^d := \{b_1^i, ..., b_M^i\} \subset \mathcal{A}_i$ and $\mathcal{V}_i^d := \{v_1^i, ..., v_L^i\} \subset \mathcal{V}_i$. We further replace the joint probability density function f by a discrete version f^d over $\mathcal{V}^d \times \mathcal{O}^d$. The marginal distribution of f^d over $\mathcal{O}_i^d$ can be written as $f_{o_i}^d = \sum_{\nu_i=1}^N (f_{o_i}^d)_{\nu_i} \delta_{o_\nu^i}$. The discrete version s_i of a distributional strategy σ_i for bidder i is now an atomic measure over $O_i^d \times A_i^d$ and can be written as $s_i = \sum_{\nu_i,\mu_i=1}^{N,M} (s_i)_{\nu_i\mu_i} \delta_{o_k^i} \otimes \delta_{b_l^i}$ with $o_\nu^i \in \mathcal{O}_i^d$ and $b_{\mu^i} \in \mathcal{A}_i^d$. Since the discrete points are fixed, the probability measures are completely determined by their coefficients. For simplicity in the notation, we will focus from now on them. In that case, the marginal condition translates to $\sum_{\mu_i} (s_i)_{\nu_i\mu_i} = (f_{o_i}^d)_{\nu_i}$ for all $\nu_i = 1, ..., N$. Therefore the set of all possible discrete distributional strategies for bidder i can be identified by the following set of coefficients:

$$\mathcal{S}_i^d := \{s_i \in \mathbb{R}^{N \times M} : (s_i)_{\nu_i\mu_i} \geq 0 \ \forall \nu_i, \mu_i, \sum_{\mu_i} s_{\nu_i\mu_i} = (f_{o_i}^d)_{\nu_i} \ \forall \nu_i\} \quad (2)$$

For a given strategy profile $(s_1, ..., s_n) \in \mathcal{S}_1^d \times ... \times \mathcal{S}_n^d$ we can compute the expected utility. This corresponds to equation (1) but in the discrete setting.

$$\tilde{u}_i(s_1, ..., s_n) = \sum_{\lambda,\nu,\mu} \big(u_i(b_\mu, v_{\lambda_i}) \prod_{j=1}^n (s_j)_{\nu_j\mu_j} \frac{(f^d)_{\lambda,\nu}}{(f_{o_1}^d)_{\nu_1} \cdots (f_{o_n}^d)_{\nu_n}} \big) \quad (3)$$

Note that $\mu = (\mu_1, ..., \mu_n)$ is a multi-index and $b_\mu = (b_{\mu_1}^1, ..., b_{\mu_n}^n)$ the action profile of all bidders (same for v and o respectively). We can split the sum in equation (3) in two parts. First we sum over ν_i, μ_i and multiply $(s_i)_{\nu_i\mu_i}$ with the remaining terms. Then we denote the second part as $(c_i)_{\nu_i,\mu_i}$ and write $\tilde{u}_i(s_1, ..., s_n) = \langle s_i, c_i \rangle_F$. Since the second sum, i. e. c_i, does not depend on s_i, the expected utility function for bidder i is linear in the bidder's own strategy. Instead of considering the discretized game G^d, we can use the expected utility $\tilde{u}_i$ and the sets of discrete distributional strategies S_i^d to define a complete-information game.

Definition 2. *Given the Bayesian game $G = (\mathcal{I}, \mathcal{V}, \mathcal{O}, \mathcal{A}, f, u)$, we construct a discrete version $G^d = (\mathcal{I}, \mathcal{V}^d, \mathcal{O}^d, \mathcal{A}^d, f^d, u)$ of the game by discretizing the respective spaces and probability distributions. The resulting sets of discrete distributional strategies $\mathcal{S}_i^d$ and the expected utility $\tilde{u}_i$ define a complete-information game $\Gamma = (\mathcal{I}, \mathcal{S}^d, \tilde{u})$, which we call the approximation game of G.*

Besides the linear utility function $\tilde{u}_i$, the continuous action sets, namely the sets of discrete distributional strategies $\mathcal{S}_i^d$, are also compact and convex. This structure allows us to use algorithms from online convex optimization. We will focus on Dual Averaging (DA) (Nesterov 2009), which was analyzed in the context of games with continuous action sets by Mertikopoulos and Zhou (2019). The method is closely related to other no-regret learners such as Online Mirror Descent (OMD). An overview and comparison of these algorithms in a more general framework can be found for instance in McMahan (2017). Dual Averaging is based on two steps: (1) Given a current state (strategy profile) $s_t \in \mathcal{S}^d$, one computes the individual gradients and uses them to update a score variable θ_t in the dual space. (2) The updated score variable θ_{t+1} is then linked or mirrored back to the feasible set in the primal space using a link function g, resulting in the new state s_{t+1}. Formally speaking, for all $i \in \mathcal{I}$ we have

$$\begin{aligned} &(1) \quad \theta_{i,t+1} = \theta_{i,t} + \eta_t \nabla_{s_i} \tilde{u}_i(s_{i,t}, s_{-i,t}) \\ &(2) \quad s_{i,t+1} = g(\theta_{i,t+1}). \end{aligned} \quad \text{(DA)}$$

For the update step of the score variable we use a nonincreasing sequence of step sizes $\{\eta_t\}$ of the form $\eta_t = \eta_0/t^\beta$ for some $\beta \in (0, 1]$ as proposed in Mertikopoulos and Zhou (2019). For the second part we use the link function

$$g(\theta_{i,t+1}) := \arg\max_{s_i \in \mathcal{S}_i^d} \langle \theta_{i,t+1}, s_i \rangle - H(s_i) \tag{4}$$

with a strongly convex regularizer H. If we ignore H for a moment, the link function returns a strategy which is most closely aligned with the aggregated gradients. Since the utility function in our application is linear, we can interpret the output as the strategy, which performs best over an aggregated version of all previous adversarial strategies. This makes the connection to the basic idea of no-regret learning even clearer. Including the regularizer again, there is a natural choice in the context of our problem. The entropic regularization or negative Gibbs entropy $H(x) = \sum_{l=1}^d x_l \log x_l$ is 1-strongly convex with respect to the 1-norm over the probability simplex $\{x \in \mathbb{R}^d : x \geq 0, \|x\|_1 = 1\}$. Since the set of feasible discrete distributional strategies $\mathcal{S}_i^d$ is a subset of the probability simplex in $\mathbb{R}^{NM}$ we use this regularization term and get

$$H(s_i) = \sum_{\nu_i, \mu_i} (s_i)_{\nu_i, \mu_i} \log(s_i)_{\nu_i, \mu_i}. \tag{5}$$

Another advantage of the entropic regularizer is that the recursive rule of (DA) can be written into a single update step, also known as normalized exponentiated gradient:

$$(s_{i,t+1})_{\nu_i \mu_i} = (f_{o_i}^d)_{\nu_i} \frac{(s_{i,t})_{\nu_i \mu_i} \exp\left(\eta_t (c_{i,t})_{\nu_i \mu_i}\right)}{\sum_{\mu_i'} (s_{i,t})_{\nu_i \mu_i'} \exp(\eta_t (c_{i,t})_{\nu_i \mu_i'})} \tag{6}$$

for all ν_i, μ_i. Note that all gradients $c_{i,t} = \nabla_{s_i} \tilde{u}_i(s_{i,t}, s_{-i,t})$ in (DA) are computed simultanously, i. e., they depend only on the current strategy profile s_t. Therefore we refer to the algorithm as simultaneous online dual averaging (SODA). If the algorithm is implemented with a fixed step size $\eta_t = \eta$, it is similar to online mirror descent (OMD) as described in Shalev-Shwartz (2012).

Convergence of the Discretization

Next, we show that approximate BNEs of the discrete game G^d naturally induce approximate BNEs of the continous game G, where the quality of the approximation depends on the coarseness of the discretization. Thus, if our algorithm finds a good solution to the discretized setting, this also induces a good solution for the continuous setting, where the quality depends on the coarseness of the discretization. For the sake of brevity, we only consider the case of single-object first-price sealed-bid auctions here. Apart from that, we do not postulate any strong assumptions, like symmetry or independence.

Proposition 3. *Let $s \in \mathcal{S}^d$ be an ε-BNE of the discrete game G^d of a first-price sealed-bid single-object auction. Let $\sigma \in \mathcal{S}$ be the strategy profile, where each σ_i is the strategy induced by s_i. Then σ is an $\varepsilon + \mathcal{O}(M\delta + L\delta^2)$-BNE of the continuous game G.*

Here M and L are constants that are introduced in the appendix and are independent of the actual discretization of the game. δ denotes the coarseness of the discretization. The central message of the proposition is that if we found an approximate BNE for the discrete game, we also found an approximate BNE for the continuous game with an additional error term decreasing linearly with the coarseness of the discretization.

The idea of the proof is as follows. Given an arbitrary strategy profile $s \in \mathcal{S}^d$ of the discrete game, we show that s naturally induces a feasible strategy profile $\sigma \in \mathcal{S}$ of the continuous game and that the difference of utilities of these two solutions is in $\mathcal{O}(\delta)$. Conversely, we can construct a feasible discrete strategy profile s from a given continuous strategy profile σ. Our central argument is that if we start with a continuous strategy profile $\sigma \in \mathcal{S}$, consider the induced discrete strategy profile $s \in \mathcal{S}^d$, which in turn induces a continuous strategy $\tilde{\sigma}_i \in \mathcal{S}_i$ for each agent i, the loss of utility is in $\mathcal{O}(\delta)$. Now suppose we found an ε-BNE $s^* \in \mathcal{S}^d$ of the discrete game and consider the induced continuous strategy profile σ^*. Let σ_i be a best response to σ^*_{-i}. Then the discrete strategy s_i induced by σ_i cannot be much better than s_i^*, since s^* is an ε-BNE. But by the result mentioned above, the utility of the continuous strategy $\tilde{\sigma}_i$ neither differs by much from s_i, nor from σ_i^*. Thus, the gain of utility from switching to σ_i is in $\mathcal{O}(\delta)$.

Mertikopoulos and Zhou (2019) proved in their Theorem 4.1 that if a complete-information game with finite-dimensional continuous action space $\mathcal{A} \subseteq \mathbb{R}^d$ is pseudo-concave and the sequence of pure strategy profiles $(a_i^t)_{t \in T}$ resulting from dual averaging converges to $a_i^* \in \mathcal{A}_i$ for all $i \in \mathcal{I}$ with positive probability, then a^* is a Nash equilibrium. A consequence of the distributional strategies that we learn is that the expected utility $\tilde{u}(s_1, \cdots, s_n)$ is linear in the bidder's own strategy. Consequently, if SODA converges to a pure strategy, it also converges to a Nash equilibrium.

Whether the gradient dynamics converge or they don't depends on the very game or the type of equilibrium in a game. In our experiments the SODA always converged in games with only a few players. Convergence can be checked quickly, for example, by looking at the distance to the last iterate, i.e., the computed solution.

There are also some conditions, when convergence is known a priori. This is the case when the utility gradients are variationally stable (Mertikopoulos and Zhou 2019), or when a game has strict Nash equilibria. As a matter of fact, variational stability and strictness of the Nash equilibrium coincide. In contrast, in games with only mixed Nash equilibria the broader class of no-regret learners cannot be expected to converge (Vlatakis-Gkaragkounis et al. 2020).

Definition 4. *Let $s^* \in \mathcal{S}$ be a pure Nash equilibrium of the approximation Γ, which is characterized by*

$$\tilde{u}_i(s^*) \geq \tilde{u}_i(s_i, s^*_{-i}) \quad \forall s_i \in \mathcal{S}_i^d, \quad \forall i \in \mathcal{I} \tag{7}$$

If the equation holds as a strict inequality for all $s_i \neq s_i^*$ for all $i \in \mathcal{I}$, the equilibrium is said to be strict.

For example, the continuous game G of a first-price sealed-bid auction in the standard independent private values model has a strict equilibrium (Krishna 2009). This means, each equilibrium strategy has a unique best response. The result by Vlatakis-Gkaragkounis et al. (2020) is for finite games with a finite set of actions, while our actions are continuous probabilities $s_i(o, b)$. However, the argument applies to our environment as well:

Proposition 5. *Let $s^* \in \mathcal{S}^d$ be a strict Nash equilibrium in the approximation game $\Gamma = (\mathcal{I}, \mathcal{S}^d, \tilde{u})$. Then s^* is variationally stable in a neighborhood U of s^*.*

The proof follows the lines of argument in Proposition A.6 by Giannou, Vlatakis-Gkaragkounis, and Mertikopoulos (2021)

Scaleability

Drivers for complexity of SODA are the number of players, the number of items or bundles (which drives the number of strategies), and the level of discretization. If the number of strategies is exponential in the number of items (as in a combinatorial auction with general valuations), then gradient-based optimization like in SODA explores all exponentially-many strategies. As a result, an algorithm learning even only approximate ε-BNE cannot be polynomial in the number of items. Cai and Papadimitriou (2014) showed with a similar argument that computing approximate ε-BNE in combinatorial auctions is NP-hard.

In most auction-theoretical models, the number of items or strategies per agent is small. Examples include single-minded bidders in combinatorial auctions or split-award auctions with two or three items only. Apart from this, a standard assumption in auction theory is that of symmetric priors and symmetric equilibrium strategies, which leads to the fact that we only need to explore the strategies of a single and not of multiple players. For example, if we further assume that the bidders are independent, the computational effort can be further reduced. In such a first-price sealed-bid auction, the expected utility can be written as

$$\tilde{u}_i(s_1, ..., s_n) = \sum_{\nu,\mu} \big((s_i)_{\nu_i\mu_i}(o_{\nu_i} - b_{\mu_i}) \\ \mathbb{P}(b_{\mu_i} \text{ is highest bid}; s_{-i})\big). \tag{8}$$

Compared to the very general formulation (3), where we sum over all combination of bids which grows exponentially in the number of bidders n, we compute the first order statistic. This way the complexity depends linearly on n, which allows us to analyze much larger settings.

So, while we know that the complexity of finding ε-BNE in general is NP-hard, computation is not necessarily a limiting factor in most of the models analyzed in economic theory and we can compute approximate BNE in due time with an appropriate level discretization.

Experimental Evaluation

Let us first describe the auction games that we analyze, before we discuss the evaluation criteria and the results.

Auction Games

We illustrate the versatility of our method in the context of single-object auctions and combinatorial auctions. For single-object and combinatorial auctions with only a few bidders, we can compute BNE within a few minutes or seconds. We compare our results to those in (Bichler et al. 2021) in order to illustrate the performance increase we get for these environments. Let us briefly revisit the auction models analyzed in their paper and in the following.

Single-Object Auctions We start with interdependencies in single-item auctions. The most well-known examples of interdependencies are the common value model (with independent observations o) and the affiliated value model for single-item auctions (Krishna 2009). We explore the second-price auction in an environment where there is one pure common value that is the same among all bidders. Three bidders $i \in \{1, 2, 3\}$ share a common $\mathcal{U}(0, 1)$-distributed value for the item of interest. Conditioned on this value, the observation o_i of bidder i is uniformly—and independently from the other observations—distributed on the interval from zero to two-times the common value. Formally, we can define the joint prior probability density function f with a four-dimensional uniformly distributed random variable $\Omega = [0, 1]^4$. For a draw $\omega \sim \mathcal{U}(\Omega)$ we set each player's type to $v_i(\omega) = \omega_4$ and each observation to be $o_i(\omega) = 2 \cdot \omega_i \cdot \omega_4$. Notice, all agents have the same value (or type), but they only learn their value if they win the auction. In this model, the symmetric BNE strategy profile can be stated in closed form as

$$\beta_i^*(o_i) = \frac{2o_i}{2 + o_i}. \tag{9}$$

For this setting, all functions required for the calculation of the utility loss from (15) can be derived analytically.

In the *affiliated values model* the individual observations are correlated. In a model with two bidders (see also (Krishna 2009, Example 6.2)), we can set $\Omega = [0, 1]^3$ and bidder $i \in \{1, 2\}$ then makes the observation

$$o_i(\omega) = \omega_i + \omega_3 \tag{10}$$

and both have a common value of $v(\omega) = \frac{1}{2}(\omega_1 + \omega_2) + \omega_3$. The symmetric BNE-strategy for both agents under a second-price payment rule is to bid truthfully and for a first-price payment rule to bid according to $\beta_i^*(o_i) = \frac{2}{3}o_i$.

The Local-Local-Global Model The LLG model consists of two objects $\{1, 2\}$, two local bidders $i \in \{1, 2\}$ and one global bidder $i = 3$, each being only interested in one specific bundle (of the single object i (locals) or both objects (global)), and we denote the valuation of each bidder's single bundle by $v_i \in \mathbb{R}$. We consider a private values (but not *independent* private values) setting with $o_i = v_i$ which allows for correlation. It was shown that with independent private values and risk-neutral bidders, core-selecting payment rules lead to significant inefficiencies in equilibrium (Goeree and Lien 2016) in combinatorial auctions. Essentially, the two local bidders attempt to free-ride on each other. Depending on the prior value distributions, it can happen that both local bidders bid too low in total and they fail to outbid the global bidder, even if their combined valuations are higher than the global bidder's. This results in an inefficient outcome and it has been used as an argument against core-selecting combinatorial auctions (Bichler and Goeree 2017). Now, it is interesting to understand equilibria with different assumptions. For example, it is reasonable to believe that bidder valuations in spectrum auctions are correlated, because telecoms face the same downstream market.

Ausubel and Baranov (2019) investigate two models of correlation among local bidders' private values and derive analytical BNE, which we will use as a baseline in our experiments besides the results of NPGA. Let's define the joint prior f to be the five-dimensional uniform distribution of a latent random variable $\omega \sim \mathcal{U}[0,1]^5$. Then let $v_3 = 2\omega_3$ be the valuation of the global bidder and

$$\begin{aligned} v_1(\omega) &= w\omega_4 + (1-w)\omega_1 \\ v_2(\omega) &= w\omega_4 + (1-w)\omega_2 \end{aligned} \tag{11}$$

be the valuations of the local bidders where the *weight* w is a random variable depending on ω_5 only. The valuations of the local bidders can be thought of a as a linear combination of an individual component ω_i and a common component ω_4. Now given an exogenous correlation parameter $\gamma \in [0,1]$, Ausubel and Baranov (2019) propose (among other ways) to choose w such that $\text{corr}(v_1, v_2) = \gamma$ via the *Bernoulli weights* model:

$$w(\omega) = \begin{cases} 1 & \text{if } \omega_5 < \gamma, \\ 0 & \text{else}, \end{cases} \tag{12}$$

The authors analytically derive the unique symmetric BNE strategies for multiple bidder-optimal core-selecting payment rules including the nearest-zero (NZ), nearest-VCG (NVCG), and nearest-bid (NB) rule in the Bernoulli weights model. These rules all choose the efficient allocation x (according to the submitted bids) but select different price vectors p from the set of core-stable outcomes. For example, the nearest-VCG rule picks the point in the core that minimizes the Euclidean distance to the (unique) Vickrey-Clarke-Groves payments. Similarly, the nearest-zero point takes the origin of the coordinate system as a reference point, while the nearest-bid rule minimizes the distance to the vector of submitted bids b.

Evaluation Criteria

The experiments are evaluated using two metrics. First we use the relative utility loss with respect to the best response ℓ to decide whether we are close to a equilibrium within the approximation game Γ. For a given strategy profile $(s_1, ..., s_n)$, the best response s_i^{br} of bidder i given the opponents strategies s_{-i} can be computed by solving the respective LP. The relative utility loss is then given by

$$\ell(s_i, s_i^{br}; s_{-i}) = 1 - \frac{\tilde{u}_i(s_i, s_{-i})}{\tilde{u}_i(s_i^{br}, s_{-i})} \tag{13}$$

After computing a discrete distributional strategy with sufficiently small relative utility loss ℓ, we want to evaluate the solution within the initial continuous setting of the auction game G. To compare our results with NPGA from Bichler et al. (2021), we choose the same approach and estimate the ex-ante utility using the sample-mean of the ex-post utilities:

$$\hat{u}_i(\cdot, \beta_{-i}) := \frac{1}{H} \sum_h u_i(\cdot, \beta_{-i}((o_{-i})_h)) \tag{14}$$

We then compare the outcome of a player bidding according to the computed strategy s_i versus bidding according to the known equilibrium strategy β_i, while all opponents j play the equilibrium strategy β_j.

$$\mathcal{L}(s_i; \beta) = 1 - \frac{\hat{u}_i(s_i, \beta_{-i})}{\hat{u}_i(\beta_i, \beta_{-i})} \tag{15}$$

Similar to NPGA, we sample $H = 2^{22}$ observation/valuation profiles and choose the actions according to the given strategies. For a given equilibrium strategy $\beta_i : \mathcal{O}_i \to \mathcal{A}_i$ the bid for an observed o_i is simply $b_i = \beta_i(o_i)$. In the case of a discrete distributional strategy s_i it is not that obvious to get a corresponding action in the continuous setting. The idea is that we identify the corresponding interval or rather discrete value o_{ν_i} to the signal o_i and sample a bid according the induced behavioral strategy $((s_i)_{\nu_i \mu_i})_{\mu_i=1}^M$. Instead of sampling only the discretized values b_{μ_i}, we extend the behavioral strategy to $\mathcal{A}_i$ by using the corresponding piecewise constant probability densities. This induces a piecewise linear cumulative distribution function from which we can sample using inverse transform sampling.

In these experiments the discretizations was done by using equidistant points. The discrete prior was then computed by evaluating the density function at these points and scaling the vector accordingly (basically midpoint rule). Other discretizations, that take the specific priors into account might lead to better results but weren't considered here.

To visualize discrete distributional strategies, i. e. probability measures over the discretized spaces, we plot for each fixed observation the 0.5% and 99.5% quantiles of the induced probability measure over the discretized action space. This way, we get areas where 99% of the bids for all valuations are contained. This allows for a visual comparison to the pure equilibrium strategies.

Our test computer contains an Intel Core i7-8565U CPU @ 1.80 Ghz and 16GB of RAM. The implementation of the algorithm uses Python 3.8.5.

Results

Single-Item Auctions In the common values and the affiliated values model all spaces are discretized using $N = M = 64$ ($L = 64$) discretization points. We assume that bidders are symmetric and therefore compute only one strategy that is played by all bidders. In both cases the algorithm starts with a random inintial strategy, an initial step size η_0 and is stopped after the relative utility loss ℓ is less than some threshold. For the affiliated values model we choose $\eta_0 = 0.1$ and $\ell < 0.1\%$. In all ten runs SODA stopped after less than 5 seconds. In the commin value model we have $\eta_0 = 0.7$ and $\ell < 0.5\%$ which leads to a running time of 1-2 minutes. In contrast, NPGA was run for 15 minutes to achieve a comparable utility loss. The resulting approximated utility loss $\mathcal{L}$ for NPGA and SODA is reported in Table 1.

Table 1: Results for single-item auctions with interdependencies.
Mean and standard deviation of $\mathcal{L}$ over 10 runs.

Auction Game	NPGA	SODA
Common Value	0.000 (0.000)	0.003 (0.002)
Affiliated Values	0.002 (0.001)	0.003 (0.001)

Figure 1: Computed strategies for single-item auctions with interdependencies

Combinatorial Auctions Next, we report our results for core-selecting combinatorial auctions in the LLG model. In this experiment, the action space and the observation space are both discretized using $N = M = 64$ points. We assume that the local bidders are symmetric and compute one strategy for both. We focus the Bernoulli weights model $\gamma \in \{0.1, 0.5, 0.9\}$ and the three core-selecting payment rules: nearest-zero (NZ), nearest-VCG (NVCG), and nearest-bid (NB). Starting with a random initial strategy, we used a step size $\eta_0 = 5$ and stop with a relative utility loss ℓ of less than 10^{-4}. The algorithm converged for the nearest-zero rule in less than 1-3 min, and for the other rules in less than 1 min. In comparison to NPGA, the results (see Table 2) show an equally low utility loss for both methods. However NPGA was again run for 15 minutes. Figure 2 illustrates the resulting BNE strategies for the local bidders assuming different Bernoulli weights.

Table 2: Relative utility loss $\mathcal{L}$ of NPGA and SODA in the LLG model with Bernoulli weight $\gamma = 0.5$. Mean and standard deviation of experiments over 10 runs.

Auction Game	NPGA	SODA
LLG NZ	-0.000 (0.000)	0.000 (0.001)
LLG NVCG	0.000 (0.000)	0.002 (0.002)
LLG NB	0.001 (0.000)	0.000 (0.001)

All strategies where computed with a relative utility loss ℓ less than 0.01%.

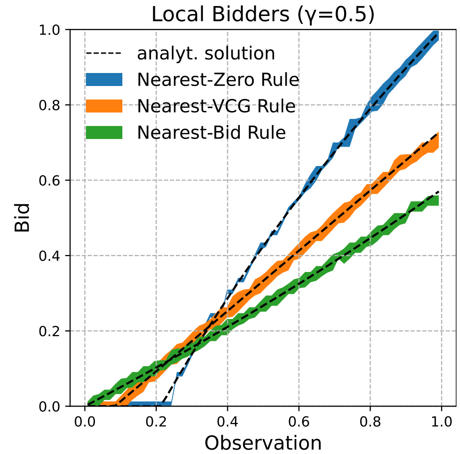

Figure 2: Computed strategies for the local bidders in the LLG model

Conclusions

Equilibrium learning has almost exclusively focused on finite normal-form games. Computing Bayesian Nash equilibria for continuous-type and -action auction games was considered intractable and only recently numerical techniques have addressed this problem. SODA is a new technique that relies on distributional strategies and a discretization of the type and action spaces. The method is very fast for auction models with a small number of bidders or with symmetric bidders. In the standard independent private values environment SODA computes approximate equilibria also for large numbers of bidders very quickly, which makes SODA a convenient numerical tool for auction theorists and a fast and simple alternative to other equilibrium computation methods. We analyzed a wide variety of auction environments and SODA converged in all of them. This suggests that although equilibrium computation in general is computationally very hard, for many relevant models an approximate equilibrium can even be found in minutes or seconds. We also know that if the gradient dynamics in SODA converge, then they converge to a Nash equilibrium.

The fact that SODA converges in such a wide variety of auction games is remarkable. Not only do we know that approximating Bayesian Nash equilibrium in multi-item auction games can be NP-hard (Cai and Papadimitriou 2014), also the analysis of adaptive learning in games is notoriously difficult and experiments suggest that non-equilibrium behavior, exemplified by chaos, may be the norm for complicated games with many players (Sanders, Farmer, and Galla 2018). Unfortunately, the dynamics generated by gradient-based learning can be very intricate and hard to interpret. Andrade, Frongillo, and Piliouras (2021) leave little hope for a general understanding of the behaviors arising from optimization-driven dynamics in games. As such, the reasons why SODA converges to an equilibrium in so many auction environments is a challenging question and one that we leave for future research.

References

Andrade, G. P.; Frongillo, R.; and Piliouras, G. 2021. Learning in Matrix Games can be Arbitrarily Complex. *arXiv preprint arXiv:2103.03405*.

Armantier, O.; Florens, J.-P.; and Richard, J.-F. 2008. Approximation of Nash equilibria in Bayesian games. *Journal of Applied Econometrics*, 23(7): 965–981.

Athey, S. 2001. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4): 861–889.

Ausubel, L. M.; and Baranov, O. 2019. Core-selecting auctions with incomplete information. *International Journal of Game Theory*.

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

Bichler, M.; and Goeree, J. K. 2017. *Handbook of spectrum auction design*. Cambridge University Press.

Bichler, M.; Guler, K.; and Mayer, S. 2015. Split-award procurement auctions—Can Bayesian equilibrium strategies predict human bidding behavior in multi-object auctions? *Production and Operations Management*, 24(6): 1012–1027.

Bosshard, V.; Bünz, B.; Lubin, B.; and Seuken, S. 2020. Computing Bayes-Nash equilibria in combinatorial auctions with verification. *Journal of Artificial Intelligence Research*, (69): 531–570.

Cai, Y.; and Papadimitriou, C. 2014. Simultaneous Bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 895–910.

Campo, S.; Perrigne, I.; and Vuong, Q. 2003. Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics*, 18(2): 179–207.

Giannou, A.; Vlatakis-Gkaragkounis, E.-V.; and Mertikopoulos, P. 2021. Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. *arXiv:2101.04667 [cs, math]*. ArXiv: 2101.04667.

Goeree, J. K.; and Lien, Y. 2016. On the impossibility of core-selecting auctions. *Theoretical Economics*, 11(1): 41–52.

Jackson, M. O.; and Swinkels, J. M. 2005. Existence of equilibrium in single and double private value auctions 1. *Econometrica*, 73(1): 93–139.

Krishna, V. 2009. *Auction Theory*. Academic press.

McAfee, R. P.; and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature*, 25(2): 699–738.

McMahan, H. B. 2017. A survey of algorithms and analysis for adaptive online learning. *Journal of Machine Learning Research*, 18(1): 3117–3166.

Mertikopoulos, P.; and Zhou, Z. 2019. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2): 465–507.

Milgrom, P. R.; and Weber, R. J. 1985. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4): 619–632.

Nesterov, Y. 2009. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1): 221–259.

Rabinovich, Z.; Naroditskiy, V.; Gerding, E. H.; and Jennings, N. R. 2013. Computing pure Bayesian-Nash equilibria in games with finite actions and continuous types. *Artificial Intelligence*, 195: 106–139.

Reny, P. J. 2011. On the existence of monotone pure-strategy equilibria in Bayesian games. *Econometrica*, 79(2): 499–553.

Sanders, J. B.; Farmer, J. D.; and Galla, T. 2018. The prevalence of chaotic dynamics in games with many players. *Scientific Reports*, 8(1): 1–13.

Shalev-Shwartz, S. 2012. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2): 107–194.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1): 8–37.

Vlatakis-Gkaragkounis, E.-V.; Flokas, L.; Mertikopoulos, P.; and Piliouras, G. 2020. No-regret learning and mixed Nash equilibria: They do not mix. In *Annual Conference on Neural Information Processing Systems*.