The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Computing Perfect Bayesian Equilibria in Sequential Auctions with Verification

Vinzenz Thoma[1,3], Vitor Bosshard[2], Sven Seuken[2,3]
[1]ETH Zurich,
[2]University of Zurich,
[3]ETH AI Center
vinzenz.thoma@ai.ethz.ch, {bosshard, seuken}@ifi.uzh.ch

Abstract

We present an algorithm for computing pure-strategy ε-perfect Bayesian equilibria in sequential auctions with continuous action and value spaces. Importantly, our algorithm includes a *verification phase* that computes an upper bound on the utility loss of the found strategies. Prior work on equilibrium computation in auctions with verification has focussed on the single-round case, but the methods do not work for sequential auctions because of two main challenges: (1) there are infinitely many subgames, and (2) the setting has no optimal substructure as bidders' beliefs and best response strategies depend on the strategies of previous rounds. We make two contributions. First, we introduce a tailor-made *game abstraction* that discretizes the auction and augments the state space with the *public beliefs*, such that an approximate equilibrium can be computed via dynamic programming. Second, we prove a *decomposition theorem* to upper bound the utility loss of the computed equilibrium. This is essential because it is neither guaranteed that the auction has an equilibrium nor that any algorithm converges to it. We validate our algorithm on multiple settings with known equilibria and apply it to a new multi-round combinatorial auction.

1 Introduction

In sequential auctions, multiple items are sold or bought over multiple rounds. Such formats are widely used, for example for selling art, fish, timber, mineral rights, broadcast licences, and government debt (Gale and Stegeman 2001). Sequential auctions are easy to run in practice and are particularly useful in settings where not all bidders or items are present in all rounds (Parkes and Singh 2003).

Equilibrium computation in auctions

As many auctions used in practice are not strategyproof, it is important to study their equilibrium properties. This is necessary to predict how bidders are likely to bid, for providing strategic recommendations to bidders, and to assess the efficiency and revenue properties of new auction designs (Bünz, Lubin, and Seuken 2022). Consider the famous example of the 2000 Swiss 3G sequential auction that only produced a fraction of the expected revenue due to a poor design and a failure to anticipate the equilibrium outcome (Klemperer 2002).

Few analytical results are known for sequential auctions, which motivates computational methods. However, designing scalable or converging algorithms is challenging, given the existing hardness results for computing equilibria (Daskalakis, Goldberg, and Papadimitriou 2009; Cai and Papadimitriou 2014). Prior work has used iterative best-response algorithms to compute Bayes-Nash equilibria in single-round auctions (Vorobeychik and Wellman 2008; Reeves and Wellman 2004; Rabinovich et al. 2013a; Bosshard et al. 2020). This does not extend to sequential auctions, due to two complications: (1) the continuous bidding space leads to infinitely many subgames, in each of which the bidders need to play an equilibrium strategy and form beliefs by Bayesian updating, and (2) the auction does not have an optimal substructure, i.e. the equilibrium of a subgame depends on bidding strategies in previous rounds. Therefore, the auction cannot be solved by backward induction.

Our Contribution

We study the problem of computing perfect Bayesian equilibria (PBEs) with verification in sequential auctions. We make two technical contributions. First, we introduce a bespoke game abstraction that reduces the state space to finite cardinality. It further endows the game with an optimal substructure, such that we can compute PBEs by backward induction (Section 4). We achieve this by (1) introducing public belief states, on which the bidders condition their strategy, (2) restricting them to bid piecewise constant, and (3) choosing how bidders update their beliefs off-equilibrium. To find a PBE candidate in the abstraction, we recursively compute Bayes-Nash equilibria (BNEs) in each round of the auction via an iterated best-response algorithm (Section 4).

Second, we prove a bound on the utility loss of our computed PBEs in the unabstracted game (Section 5). Our proof exploits the optimal substructure and the piecewise convexity of bidders' utilities to bound the utility loss over the infinite set of possible types and histories. To approximate the bound for a given set of strategies, we provide an easy-to-implement verification procedure (Section 5).

Our algorithm reproduces existing analytical results in multiple settings, including sequential sales with and without reserve prices, as well as procurement auctions. It further finds interesting strategic behavior in a novel sequential auction with combinatorial values (Section 6).

2 Related Work

Analytical Equilibria in Sequential Auctions

The seminal work of Milgrom and Weber (2000) derived equilibria for sequential auctions with single-minded bidders. This was later extended to multi-unit demand (Katzman 1999; Rodriguez 2009; Gale and Stegeman 2001) and round-dependent reserve prices (Gong, Tan, and Xing 2013; Landi and Menicucci 2018). Another line of work has focused on equilibria of sequential procurement auctions (Gale, Hausch, and Stegeman 2000; McAfee and Vincent 1993; Ashenfelter 1989; Kokott, Bichler, and Paulsen 2019).

Equilibrium Computation

There exists a vast literature on equilibrium computation in extensive-form games. The main approaches are: computing best-responses (Hendon, Jacobsen, and Sloth 1996; McMahan and Gordon 2007; Ganzfried and Sandholm 2009; Heinrich, Lanctot, and Silver 2015), counterfactual regret minimization (Zinkevich et al. 2007; Moravčík et al. 2017; Brown and Sandholm 2019), and reinforcement learning (RL) (Silver et al. 2016; Brown et al. 2020).

These methods generally assume finite actions and often restrict to two-player zero-sum games. In contrast, auctions are general-sum and have a continuous type and action space. This gap has motivated auction-specific equilibrium solvers. Ideas include using deep learning (Martin and Sandholm 2023; Bichler et al. 2021), gradient-based optimization (Bichler, Fichtl, and Oberlechner 2023; Kohring, Pieroth, and Bichler 2023), and iterated best response computation (Vorobeychik and Wellman 2008; Reeves and Wellman 2004; Rabinovich et al. 2013b; Li and Wellman 2021). Generally, these approaches compute equilibria in an abstracted version of the game, without guarantees on the *utility loss* in the original auction. Bosshard et al. (2017; 2018; 2020) provide a verification procedure to bound this loss. However, all the above approaches for equilibrium computation are limited to single-round auctions.

There are few results on equilibrium computation in sequential auctions. The concurrent work of Pieroth, Kohring, and Bichler (2023) and an earlier work by Greenwald, Li, and Sodomka (2012) both iteratively use RL to compute *ex-ante* Nash equilibria, such that bidders's strategies are best responses to each other *in expectation* before the auction starts and bidders observe their own type. In contrast, we focus on *ex-interim* perfect Bayesian equilibria, requiring that bidders best respond in each subgame after having observed their type and updating their beliefs. d'Eon, Newman, and Leyton-Brown (2024) use RL to compute equilibria in auctions with demand instead of value queries, and Thoma et al. (2023) use it to study the exploitability of truthful bidding.

From an algorithmic perspective, these RL approaches iterate over bidders and approximately compute their best responses for the whole auction by solving single-agent Bellman equations. In contrast, we first introduce equilibrium Bellman equations. These decompose the auction into simpler single-round auctions, whose equilibria we subsequently find via iterated best responses.

3 Preliminaries

Formal Model

A set N of n bidders compete over T rounds for a set K of k items. Each bidder i has a type $\theta_i \in \Theta_i \subseteq \mathbb{R}^{d_i}$ with $d_i \in \mathbb{N}_+$. θ_i is a continuous random variable with density f_i. We use the subscript $-i$ to denote all bidders except i and no subscript to denote all bidders, e.g. θ_{-i} denotes all but i's types and θ all types. In every round t, bidders can submit an XOR bid $b_{i,t} \in \mathbb{R}^q_{\geq 0}$, where $q \leq 2^K$.[1] Given $b_t = (b_{1,t}, \dots, b_{n,t})$, the auctioneer employs an allocation rule $X_t(b_t|x_{<t})$, where $x_{<t}$ denotes past allocations. X_t deterministically assigns a subset $K_t \subseteq K \setminus \bigcup_{\tau<t} K_\tau$ of the remaining items to the bidders, who get charged according to a deterministic payment rule $P_t(b_t|x_{<t}) \in \mathbb{R}^n$. The allocation x_t and a deterministic announcement $a_t = A_t(b_t|x_{<t})$ is made public after each round. Examples for A_t are announcing the winning bid, all bids, or the winner's payment.

Bidders can decide to *drop out* of the auction, e.g. if they already won what they are interested in. $N_t(x_{<t}) \subseteq N$ denotes the set of bidders remaining at round t, given $x_{<t}$.

We illustrate the model in the following example, which we elaborate on in Appendix E.

Example 1 (Two-round FPSB auction). *Consider a two-round first-price sealed-bid (FPSB) auction with uniform types where the highest bid wins. In this case,* $\theta_i \in [0,1]$, $f_i(\theta_i) = 1$. *The bidding space is* $\mathbb{R}_{\geq 0}$. *The allocation rule assigns the good to the highest bidder, i.e.* $X_t = \arg\max_i b_{i,t}$ *for* $t \in \{1,2\}$. *The auctioneer announces the winning bid, such that* $A_t(b_t|h_t) = P_t(b_t|h_t) = \max_i b_{i,t}$. *As bidders are single-minded in the first round, we have* $\forall i : v_i(\theta_i, \{i\}|\emptyset) = \theta_i$, *where* v_i *is the value function of bidder* i.

A sequential auction $\mathcal{A}$ can be modeled as a Bayesian extensive-form game. The state space is $S = \Theta \times H$, where $\Theta = \{\Theta_1, \dots, \Theta_n\}$ and H is the set of *histories*. A history $h_t = \{x_1, a_1, \dots, x_{t-1}, a_{t-1}\}$ consists of past allocations and announcements. Bidders act simultaneously. The resulting allocation and announcement define the transition to the next state. The possible successors of state s_t are denoted by $N^+(s_t) = \{s|\theta^s = \theta^{s_t} \wedge \exists x_t, a_t : h^s = h^{s_t} \cup \{x_t, a_t\}\}$, where θ^s, h^s denote the types and history of s. N^+ endows S with a tree structure, which we call *game tree*. The game starts with nature choosing θ and entering the state $s_0 = (\theta, \emptyset)$. A *subgame* at s is an auction that starts from s instead of s_0.

Each bidder i does not fully observe s but only their respective *information set* $I_i(s) = (\theta_i^s, h^s)$. A (pure) bidding strategy is thus a function $\sigma : \mathcal{I}_i \rightarrow \mathbb{R}^q_{\geq 0}, (\theta_i, h) \mapsto \sigma(\theta_i|h)$, where $\mathcal{I}_i$ is the set of i's information sets. In

[1] We use the XOR language for ease of exposition, but our results apply without loss of generality to other bidding languages.

Section 4 we work with *piecewise-constant* strategies. Following Bosshard et al. (2020), we define a *hyperrectangle* $[y, z) = \prod_{k=1}^{d} [y_k, z_k)$ in $\mathbb{R}^d$ as a Cartesian product of half open intervals and a *partition* $\mathcal{P}_i$ of Θ_i as a set of disjoint hyperrectangles $[y, z)_j$ covering Θ_i. A strategy σ_i is piecewise constant if there exists a partition $\mathcal{P}_i$ such that $\forall [y, z)_j \in \mathcal{P}_i, \forall \theta_i \in [y, z)_j : \sigma_i(\theta_i|h) = \sigma_i(y|h)$.

As bidders do not know each other's types, they form *beliefs*. Initially, the public belief about θ_i, denoted by $\mu_i(\cdot|\emptyset)$, is given by f_i. As the game progresses, h_t reveals information about the other bidders's types—referred to as *signaling*—and is used to update the beliefs using Bayes rule. We denote the common belief about bidder i after history h_t by $\mu_i(\cdot|h_t)$ and definde $\mu = (\mu_1, \ldots, \mu_n)$. To ensure beliefs are independent and common information, we make two assumptions. First, the announcements cannot correlate bidder types. Moreover, any bidder remaining in the auction has no informational advantage from his privately observed payment $p_{i,t}$ over the public observation (x_t, a_t). Both assumptions are satisfied for most standard auctions.[2]

Utility and Equilibrium

Bidder i has a value function $v_{i,t}(\theta_i, x_t|x_{<t})$, linear in θ_i, denoting the value of assignment x_t in round t, given $x_{<t}$. The corresponding utility is defined as $u_i(\theta_i, x_t, p_t|x_{<t}) = v_{i,t}(\theta_i, x_t|x_{<t}) - p_{i,t}$. Given h_t and bidding strategies σ, i's expected utility $\overline{u}_i(\theta_i|\sigma, h_t)$ (over all remaining rounds) is defined as:

$$\mathop{\mathbb{E}}_{\theta_{-i} \sim \mu_{-i}} \left[\sum_{\tau=t}^{T} u_i(\theta_i, X_\tau(\sigma(s_\tau)|x_{<\tau}), P_\tau(\sigma(s_\tau)|x_{<\tau})) \right], \tag{1}$$

where $\sigma(s_\tau) = (\sigma_1(\theta_1^{s_\tau}|h_\tau^{s_\tau}), \ldots, \sigma_n(\theta_n^{s_\tau}|h_\tau^{s_\tau}))$ and the expectation is taken with respect to $\mu_{-i}(\cdot|h_t)$.

Given fixed opponent strategies σ_{-i}, we define the *best response utility* as $\overline{u}_i^{BR}(\theta_i|\sigma_{-i}, h) = \sup_{\sigma_i} \overline{u}_i(\theta_i|\sigma_i, \sigma_{-i}, h)$. If the utility-maximizing strategy exists, we refer to it as the *best response* σ_i^{BR}. The related *utility loss* of agent i is denoted as $l_i^{BR}(\theta_i|\sigma, h) = \overline{u}_i^{BR}(\theta_i|\sigma_{-i}, h) - \overline{u}_i(\theta_i|\sigma, h)$. The latter is used to characterize equilibrium strategies.

Definition 1 (ε-PBE)**.** *A tuple (σ, μ) of strategies and beliefs constitutes an ε-perfect Bayesian equilibrium if it fulfills the following conditions:*

1. ***Sequential rationality:*** $\forall i, \forall \theta_i, \forall h : l_i^{BR}(\theta_i|\sigma, h) \leq \varepsilon$.
2. ***Correct initial beliefs:*** $\forall i : \mu_i(\theta_i|h) = f_i(\theta_i)$.
3. ***Bayesian updating:*** *Whenever possible, beliefs are obtained by Bayesian updating.*

4 Equilibrium Search

We focus on finding pure-strategy PBEs. First, pure strategies are more intuitive for bidders to follow and simplify the search space. Second, there always exists an ε-PBE in pure strategies for some ε, which we can upper bound with our verification procedure introduced in Section 5.[3] Third, most auction mechanisms do not disclose all bids, such that bidders have fewer incentives to mix strategies in order to conceal their types. Fourth, results in finite games link pure-strategy equilibria to better convergence of learning dynamics (Mertikopoulos and Sandholm 2016).

Our approach involves (1) creating a belief-augmented representation, such that finding a PBE can be decomposed into iteratively finding BNEs of the individual auction rounds, and (2) creating a finite abstraction, such that the resulting dynamic program can be approximately solved.

Belief-Augmented Game Representation

In perfect-information extensive-form games, an equilibrium can be written as a dynamic program and computed by backward induction (Laraki, Renault, and Sorin 2019). This is not true for sequential auctions, as the best response in a subgame depends on the beliefs induced by previous-round strategies. In comparable cases, these issues were resolved by augmenting the state space with the *public beliefs* (also referred to as *common information*) (Brown et al. 2020; Nayyar et al. 2014; Zhang and Sandholm 2021). As the *belief-augmented* representation gives agents all necessary information to make their decision, an equilibrium can be computed via backward induction. In particular, Ouyang, Tavafoghi, and Teneketzis (2017) propose a dynamic program to compute PBEs, where each timestep consists of finding the BNE of a simpler stage game. We extend the approach to sequential auctions, where—in contrast to their setting—actions are not public, bidders can leave and have continuous type and action spaces.

So far, a strategy depended on θ_i and $h_t = \{x_{<t}, a_{<t}\}$. Next, we augment h_t by μ. In fact, as the main purpose of announcements $a_{<t}$ is signaling, we replace them with μ.

Definition 2 (PBS)**.** *A Public Belief State (PBS) is a tuple $(t, x_{<t}, \mu)$, where t is the round, $x_{<t}$ the past allocations and μ the set of current public beliefs.*

We denote a single PBS by β. Moreover, $\mathcal{B}_t$ is the set of all possible PBS in round t and $\mathcal{B} = \bigcup_{t=1}^{T} \mathcal{B}_t$. To ensure correct initial beliefs, we set $\mathcal{B}_1 = \{(1, \emptyset, f)\}$. We define a PBS strategy $\sigma_i(\theta_i|\beta)$ as bidding $\sigma_i(\theta_i|\beta)$ in every history h that has the same allocations $x_{<t}^{\beta}$, and induces the same public belief μ^{β} as β.[4] As beliefs are part of a PBS, the belief updates and state transitions merge.

Definition 3 (PBS Transition Function)**.** *A PBS transition function ρ maps a tuple of the current PBS, allocation, and announcement (β, x_t, a_t) to a new PBS β'.*

[2]We defer the details to Appendix D. However, we briefly note that piecewise-constant strategies—as we use later—in combination with random tiebreaking can cause types to become interdependent, which is why we focus on deterministic mechanisms.

[3]While auctions without pure-strategy PBEs exist (Gong, Tan, and Xing 2013), they are not guaranteed to have a mixed-strategy equilibrium either (consider a first-price auction with one good and one bidder with bidding space $(0, 1]$).

[4]Mapping histories to their induced beliefs is not injective. By focussing on PBS strategies we are restricting the strategy space. This is a mild restriction, as past announcements have no direct influence on the expected utility (the mechanism only depends on $x_{<t}$). Moreover, in Proposition 3 we show that the best response to PBS strategies is a PBS strategy again.

Definition 4 (Consistent Transition)**.** *Given PBS strategies σ, a transition function ρ is* consistent *with σ if for all β it holds that if (x_t, a_t) has a nonzero probability under $\sigma(\cdot|\beta)$, then $\rho(\beta, x_t, a_t)$ maps β to $(x_{<t} \cup \{x_t\}, \mu')$, where μ' is obtained by Bayesian updating.*

To compute PBEs by backward induction, we introduce the concept of a *stage auction*.

Definition 5 (Stage Auction)**.** *Given a sequential auction $\mathcal{A}$, round t, a function $U_{t+1} : \Theta_i \times \mathcal{B}_{t+1} \to \mathbb{R}$, and a PBS transition ρ_t, we define the stage auction $\mathcal{A}_t(U_{t+1}, \rho_t, \beta)$ as the following Bayesian game. There is a set of strategies σ, assumed to be common knowledge. Bidder i observes θ_i and β, bids $b_{i,t} = \sigma_i(\theta_i|\beta)$ and receives the utility*

$$\begin{aligned} u_i^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\theta_i|b_t) =& v_i(\theta_i, X_t(b_t)|x^{\beta}_{<t}) - P_{i,t}(b_t) \\ &+ U_{t+1}(\theta_i, \rho_t(\beta, X_t(b_t), A_t(b_t))), \end{aligned}$$

where v_i, X_t, A_t, P_t are as described by $\mathcal{A}$.

Bidder i's expected utility in the stage auction is given by

$$\overline{u}_i^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\theta_i|b_{i,t}) = \mathop{\mathbb{E}}_{\substack{\theta_{-i}\sim\mu_{-i},\\ b_{-i}\sim\sigma_{-i}}} \left[u_i^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\theta_i|b_t)\right].$$

Strategies $\sigma(\cdot|\beta)$ form a BNE of the stage auction $\mathcal{A}_t(U_{t+1}, \rho_t, \beta)$ if for all i and θ_i it holds that

$$\sup_{b_{i,t}} \overline{u}_i^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\theta_i|b_{i,t}) \leq \overline{u}^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\theta_i|\sigma_i(\theta_i|\beta)).$$

To formulate PBEs as a dynamic program, we need a Bellman equation for equilibria.

Definition 6 (Bellman Utility Update)**.** *The Bellman utility update is given by*

$$\mathcal{T}(t, i, U_{t+1}, \sigma_t, \rho_t)(\sigma_i, \beta_t) = \overline{u}_i^{\mathcal{A}_t(U_{t+1},\rho_t,\beta)}(\cdot|\sigma_i(\theta_i|\beta)).$$

Next, we extend the PBE dynamic program of Ouyang, Tavafoghi, and Teneketzis (2017) to sequential auctions.

Proposition 1. *A set of strategies σ and a transition functions ρ form a PBE of a sequential auction $\mathcal{A}$, if they solve the following dynamic program for $t \in \{1, \dots, T\}$:*

1. $\forall i : U^i_{T+1} = 0$ and $\forall i \notin N^t(\beta_t) : U^i_t(\cdot|\beta_t) = 0$.
2. $\forall i \in N, \forall t : U^i_t = \mathcal{T}(t, i, U_{t+1}, \sigma_t, \rho_t)$.
3. $\forall t, \forall \beta_t \in \mathcal{B}_t : \sigma_t(\cdot|\beta_t)$ is a BNE of $\mathcal{A}_t(U_{t+1}, \rho_t, \beta_t)$.
4. $\forall t, \forall \beta_t \in \mathcal{B}_t : \rho_t(\beta_t, \cdot, \cdot)$ is consistent with $\sigma_t(\cdot|\beta_t)$.

The proof works by backward induction and is deferred to Appendix C.

Game Abstraction

The dynamic program defined in Proposition 1 is not yet computationally solvable. First, $\mathcal{B}_t$ is infinite for all $t > 1$, such that there are infinitely many stage auctions to solve. Second, computing the BNE of a single stage auction is hard, especially due to the continuous type space. For this reason, we create a finite abstraction of $\mathcal{A}$, compute a PBE of the abstraction and then map it back to $\mathcal{A}$. This is a common approach for solving games with prohibitively large strategy spaces (Sandholm et al. 2012; 2015). Below, we seperately describe our approach to create a finite set of PBS $\mathcal{B}_t$ (to solve only finitely many stage auctions) and to discretize Θ (to compute approximate BNEs).

Type Space Abstraction Given $\mathcal{A}_t(U_{t+1}, \rho_t, \beta)$ and strategies σ_{-i}, the continuity of Θ_i renders it infeasible to compute a best response for each θ_i. However, Bosshard et al. (2020) have presented an approach for computing BNEs using piecewise-constant strategies that we extend to stage auctions.

By restricting PBS strategies to be piecewise constant with respect to a partition $\mathcal{P}$, we are creating an abstracted game, where each bidder i has a finite typespace $\hat{\Theta}_i = \mathcal{P}_i$, with types $\hat{\theta}_i$ corresponding to hyperrectangles $[y, z)_j$.[5] In this case, the initial distributions become probability mass functions $\hat{F}_i(\hat{\theta}_i) = \int_{\hat{\theta}_i} f_i(\theta_i)d\theta_i$, which simplifies Bayesian updating (cf. Appendix E). The utility of a bidder in the abstraction with type $[y, z)_j$ is the utility of the corresponding bidder in $\mathcal{A}$ with type y (the vertex nearest to the origin). This causes an approximation error, which we bound in Section 5.

Belief Space Abstraction Given partition $\mathcal{P}$ of Θ and initial distribution f, we denote by $\mathcal{B}^{f,\mathcal{P}}_t \subset \mathcal{B}_t$ the set of PBS in round t that can be reached by Bayesian updating from f. We further define $\mathcal{B}^{f,\mathcal{P}} = \bigcup_{t=1}^{T} \mathcal{B}^{f,\mathcal{P}}_t$. $\mathcal{B}^{f,\mathcal{P}}$ is finite, which is crucial to solving only finitely many stage auctions.

Proposition 2. *Given a partition $\mathcal{P}$ of Θ and initial type distribution f, it holds that $|\mathcal{B}^{f,\mathcal{P}}| < \infty$.*

This follows from the finiteness of the $\mathcal{P}$ and the auction being deterministic. The full proof is deferred to Appendix C.

If all bidders follow a piecewise-constant strategy $\hat{\sigma}$, the number of bids and thus of payments, announcements and histories with positive probability is finite. In particular, all PBS with positive probability are in $\mathcal{B}^{f,\mathcal{P}}$. However, for a PBE we need counterfactual reasoning. In order for $\hat{\sigma}_i$ to be a best response to $\hat{\sigma}_{-i}$ in a given stage auction, we need to know what utility i would achieve by deviating. There are infinitely many such deviations that have probability zero under $\hat{\sigma}$ and thus infinitely many possible PBS and stage auctions to solve. However, recalling Definition 1, in a PBE we are free to choose beliefs on histories with zero probability under $\hat{\sigma}$. In particular, we can choose a consistent transition function ρ that projects all those outcomes onto the finite set $\mathcal{B}^{f,\mathcal{P}}$. We call a transition function ρ consistent with a piecewise-constant strategy $\hat{\sigma}$ and *finite* if it performs Bayesian updating on all histories with positive probability under $\hat{\sigma}$ and projects all other histories onto $\mathcal{B}^{f,\mathcal{P}}$. There are several ways to implement a finite and consistent transition function. We discuss this in Appendix E. Next, we formalize the auction abstraction.

Definition 7 (Abstracted Belief-Based Sequential Auction)**.** *A partition $\mathcal{P}$, strategies $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$, and transition $\hat{\rho}$ induce a finite belief-augmented abstraction $\hat{A}$ of a sequential auction $\mathcal{A}$ if the following hold:*

- *For all i, $\hat{\sigma}_i : \mathcal{P}_i \times \mathcal{B}^{f,\mathcal{P}} \to \mathbb{R}^q$ is a well-defined piecewise-constant (with respect to $\mathcal{P}$) PBS strategy.*
- *ρ is finite and consistent with $\hat{\sigma}$.*

[5] The definition of piecewise-constant strategies canonically extends to PBS strategies.

A Recursive Best Repsonse PBE Algorithm

Backward induction In Algorithm 1, we combine the dynamic program from Proposition 1 with the abstraction from Definition 7 to compute a PBE of $\hat{\mathcal{A}}$. This is done by recursively computing the BNEs of its finitely many stage auctions, starting in the final round T.

Algorithm 1: Recursive PBE Computation

```
Input: Auction A, Partition P
Compute B_t^{f,P} for all t
Set U_{T+1} = 0 and U_t^i(·|β) = 0 for all i ∉ N^t(β)
for t = T to 1 do
    for all β_t ∈ B_t^{f,P} do
        σ̂_t(·|β_t), ρ̂_t(·,·|β_t) ← FindBNE(A_t(U_{t+1}, ·, β_t))
        U_t(·|β_t) ← T(t, i, U_{t+1}, σ̂_t, ρ̂_t)
    end for
end for
return σ̂, ρ̂
```

Computing BNEs of Stage Auctions To compute the BNEs of the abstracted stage auctions, we build upon the work of Bosshard et al. (2020) for single-round auctions. We describe the main ideas of how we adapt their algorithm below and defer the details to Appendix A.

In brief, the approach relies on iteratively computing the best response of each bidder in the stage auction and updating their strategies accordingly. While such *best response dynamics* are not guaranteed to converge, when they do, the obtained strategies form a BNE by definition.

As the expected utility is not differentiable, we cannot use gradient-based methods to compute best responses.[6] Instead, we use pattern search—a zeroth order method proposed by Hooke and Jeeves (1961). To evaluate the expected utility, we use Monte Carlo (MC) integration, sampling types from the belief distribution μ_{-i}^{β}. Having computed a best response $\hat{\sigma}_i^{BR}$ to $\hat{\sigma}_i$, we do not directly use $\hat{\sigma}_i^{BR}$ in the next iteration but instead update according to $\sigma_i \leftarrow (1-\gamma)\sigma_i + \gamma\sigma_i^{BR}$, where γ decreases over time to avoid oscillations. Additionally, we design a setting-specific rule on how to choose a finite and consistent transition function for the stage auction. The pseudocode is given in Appendix A (Algorithm 2), along with further details.

5 Equilibrium Verification

There are two issues with using Algorithm 1 as a standalone method. First, even if a pure-strategy PBE exists, the algorithm might not converge. The lack of convergence guarantees for equilibrium solvers in auctions has been well-documented (Bichler et al. 2023). Second, even if Algorithm 1 converges in the abstraction $\hat{\mathcal{A}}$, the utility loss in $\mathcal{A}$ might be much larger, such that the found strategies only form an ε-PBE in $\mathcal{A}$ for impractically large ε. We refer to the problem of bounding the utility loss in $\mathcal{A}$ as *verification*.

[6]Indeed a small change in the bid could mean the difference between winning or loosing.

The difficulty with verification is that we need to argue over the set of all possible strategies and cannot simply restrict the search space as in Section 4.[7] To address this, we show in Theorem 2 that the utility loss in $\mathcal{A}$ can be bound by only computing best responses at finitely many vertices and in finitely many PBS. Based on this result, we introduce a verification procedure to run after Algorithm 1.

Decomposing the Utility Loss

All strategies $\hat{\sigma}$ found by Algorithm 1 satisfy Bayesian updating and correct initial beliefs. To verify them, we only have to bound the utility loss, i.e. find ε, such that $\forall \theta_i, \forall h$: $l_i^{BR}(\theta_i|\hat{\sigma}, h) \leq \varepsilon$. This revives the issue that H and Θ_i are continuous. To tackle the continuity of H, we show that the best response to a PBS strategy is again a PBS strategy.

Proposition 3 (Closedness of PBS strategies)**.** *Given a set of PBS strategies $\hat{\sigma}$, a transition function ρ, and a history h_t. Let $[h_t]$ denote the equivalence class of all histories that induce the same PBS according to ρ. Then $\forall \theta_i \in \Theta_i, \forall h' \in [h_t], \forall b_{i,t} \in \mathbb{R}_{\geq 0}^q$, it holds that*

$$\overline{u}_i(\theta_i|\sigma_{-i}, b_{i,t}, \sigma_{i,>t}, h_t) = \overline{u}_i(\theta_i|\sigma_{-i}, b_{i,t}, \sigma_{i,>t}, h').$$

In particular, for any strategy, there exists a corresponding PBS strategy with the same expected utility.

The proof is deferred to Appendix C. Intuitively, the result holds because $\overline{u}_i$ is independent of $a_{<t}$. Proposition 3 shows if a best response to PBS strategies $\hat{\sigma}_{-i}$ exists, there is an equivalent PBS best response. In particular, the utility loss is the same for all h inducing the same PBS. Therefore, we can restrict to bounding the utility loss in the finitely many PBS in $\mathcal{B}^{f,\mathcal{P}}$. We denote by $l_i^{BR}(\theta_i|\hat{\sigma}, \beta) = \sup_{\tilde{\sigma}_i} \overline{u}_i(\theta_i|\hat{\sigma}_{-i}, \tilde{\sigma}_i, \beta) - \overline{u}_i(\theta_i|\hat{\sigma}, \beta)$ $i's$ utility loss at β from not playing a PBS best response. Similarly, we define the *immediate utility loss* $l_i^{IBR}(\theta_i|\hat{\sigma}, \beta) = u_i^{IBR}(\theta_i|\hat{\sigma}, \beta) - u_i(\theta_i|\hat{\sigma}, \beta)$, where $u_i^{IBR}(\theta_i|\hat{\sigma}, \beta)$ is the highest utility i can reach by changing bids in the stage auction $\mathcal{A}_t(U_{t+1}, \rho, \beta)$, with ρ, U_{t+1} given by Algorithm 1. Below, we show one of our main contributions: that the utility loss decomposes into a finite sum of immediate utility losses.

Theorem 1 (Decomposition)**.** *Given piecewise-constant PBS strategies $\hat{\sigma}$, it holds for all i, β_t, θ_i that*

$$l_i^{BR}(\theta_i|\hat{\sigma}, \beta_t) \leq l_i^{IBR}(\theta_i|\hat{\sigma}, \beta_t) + \max_{\beta_{t+1} \in \mathcal{B}_{t+1}^{f,\mathcal{P}}(\beta_t, \rho)} l_i^{BR}(\theta_i|\hat{\sigma}, \beta_{t+1}),$$

where $\mathcal{B}_{t+1}^{f,\mathcal{P}}(\beta_t, \rho) \subseteq \mathcal{B}_{t+1}^{f,\mathcal{P}}$ is the set of possible successors of β_t under ρ.

The proof works by backward induction and is deferred to Appendix C. Proposition 3 and Theorem 1 let us tackle the continuity of H. Fixing θ_i, we can bound the utility loss by calculating the immediate utility loss at the finitely many possible PBS. However, to verify $\hat{\sigma}$, we also need to upper bound the utility loss over the continuous type space. To do so, we extend an idea by Bosshard et al. (2020) for

[7]We can restrict to all pure strategies. Fixing all other bidders, i faces an MDP with a pure optimal strategy (Puterman 2005).

verifying BNEs in single-round auctions. We can show that $\overline{u}_i^{IBR}$ is convex in each hyperrectangle of $\mathcal{P}_i$ and thus upper bound by a linear function. $\overline{u}_i$ is linear as well, so that the immediate utility loss can be upper bound by the difference of two linear functions evaluated at the hyperrectangle's finitely many vertices.[8]

Proposition 4. *Given a piecewise-constant (on $\mathcal{P}$) PBS strategy $\hat{\sigma}$, the following holds for all i, β:*

$$\sup_{\theta_i \in \Theta_i} l_i^{IBR}(\theta_i|\hat{\sigma}, \beta) \leq \max_{[y,z)_j \in \mathcal{P}_i} \max_{w \in V([y,z)_j)} \overline{u}_i^{IBR}(w|\hat{\sigma}_i(y), \hat{\sigma}_{-i}, \beta) - \overline{u}_i(w|\hat{\sigma}_i(y), \hat{\sigma}_{-i}, \beta),$$

where the last term is i's expected utility with type w bidding $\hat{\sigma}_i(y)$, and $V([y,z)_j)$ denotes the vertex set of $[y,z)_j$.

The full proof is deferred to Appendix C. Combining Proposition 3, Proposition 4, and Theorem 1 yields the main result of this section: that we can verify an ε-PBE in the full auction by computing the immediate utility loss at finitely many vertices in finitely many abstracted stage auctions.

Theorem 2 (ε-PBE Verification)**.** *A set of piecewise-constant (on partition $\mathcal{P}$) PBS strategies $\hat{\sigma}$ and a finite and consistent transition ρ form an ε-PBE in a sequential auction $\mathcal{A}$, with ε being upper bound by*

$$\varepsilon \leq \max_{i \in N} \max_{\beta_{t:T}} \sum_{\beta \in \beta_{t:T}} \max_{[y,z) \in \mathcal{P}_i} \max_{w \in V([y,z))} \overline{u}_i^{IBR}(w|\hat{\sigma}, \beta) - \overline{u}_i(w|\hat{\sigma}_i(y), \hat{\sigma}_{-i}, \beta),$$

where $\beta_{t:T}$ denotes a vector of successor PBS under ρ going from t to T.[9]

The proof is deferred to Appendix C.

Verification Procedure

Theorem 2 gives a strict upper bound on the utility loss of a set of piecewise-constant PBS strategies $\hat{\sigma}$, expressed in terms of the immediate utility losses evaluated at finitely many vertices of $\mathcal{P}$ and finitely many PBS. The result directly motivates an implementable verification procedure for PBEs. We have already developed an algorithm to compute best responses (and hence the corresponding immediate utility loss) in stage auctions in Section 4. Analogously, we can use pattern search to approximate the immediate utility loss at the finitely many vertices of $\mathcal{P}$ and then find the largest possible utility loss for any $\beta \in \mathcal{B}^{f,\mathcal{P}}$ by backward induction over the rounds. The detailed pseudocode is given in Appendix A (Algorithm 4).

6 Experiments

We evaluate our algorithm in a range of settings with and without known analytical results.[10] Unless otherwise specified, we perform 10 runs per setting, using a partition with uniform intervals and a grid size of 100 for our piecewise-constant strategies, 100 iterations of iterated best responses to compute BNEs, and up to 100,000 MC samples for search and 200,000 for verification. Our algorithm is embarrassingly parallelizable, as inherently all β_t in round t can be solved independently of each other. Therefore, the wall clock time for solving each setting is a small fraction of the reported core hours. Below we describe the different settings and our corresponding results.

Env.	**Pay**	n	**|K|**	$L_2^{t=1}$	$L_2^{t=2}$	$L_2^{t=3}$	$L_2^{t=4}$	**Hrs.**
Seq.	1st	3	2	0.008	0.010			22
Seq.	1st	4	3	0.007	0.007	0.012		110
Seq.	1st	5	4	0.010	0.005	0.007	0.012	176
Seq.	2nd	3	2	0.008	0.006			4
Seq.	2nd	4	3	0.012	0.008	0.006		76
Seq.	2nd	5	4	0.014	0.009	0.008	0.006	154
Res.	1st	3	2	0.004	0.005			28
Res.	1st	4	2	0.008	0.007			47
Res.	2nd	3	2	0.005	0.006			3
Res.	2nd	4	2	0.006	0.006			5
Split	1st	3	2	0.003	0.003			85
Split	1st	4	2	0.006	0.005			139

Table 1: Average L_2 distances (with standard error below 10^{-5}) to the known equilibrium and the runtime (in core hours). Seq = Sequential, Res = Reserve price, Split = Split-award. 1st = First-price, 2nd = Second-price. n = number of bidders, $|K|$ = number of goods.

Sequential Sales

An auctioneer is selling K identical goods to n single-minded bidders with values uniformly distributed on $[0, 1]$. There are K rounds. Each round, the auctioneer runs a first-price or second-price sealed-bid auction for one item, breaking ties by a predefined order. For the first-price auction, the payments are made public (i.e. $a_t = p_t$) and for the second-price auction, the winning bid (i.e. $a_t = \max_i b_{i,t}$). The winning bidder leaves the auction. The equilibrium strategy in round k is given by $\sigma_i^*(\theta_i) = (n-K)\theta_i/(n-k+1)$ and $\sigma_i^*(\theta_i) = (n-K)\theta_i/(n-k)$, for first-price and second-price auctions, respectively (Krishna 2002).

In Table 1, we report L_2 distances to the known analytical equilibrium for $K \in \{2, 3, 4\}$ with $n = K + 1$. We plot the found strategies in Figures 3–20 in Appendix G. Our algorithm converges to the known PBEs, achieving small L_2 distances across all settings.

Sequential Sales with Ascending Reserve Price

Gong, Tan, and Xing (2013) present a complication of the sequential sales setting, where the auctioneer further

[8] The argument in the single-round case of Bosshard et al. (2020) is simpler, as there are no subsequent rounds, and the best response utility is convex instead of piecewise convex.

[9] The set of $\beta_{t:T}$ is finite, since $\mathcal{B}^{f,\mathcal{P}}$ is finite.

[10] We provide our source code in the supplementary material. We ran all experiments on a Debian compute cluster with 20 nodes, each node having 128 GB RAM and two Intel E5-2680 2.80GHz processors for a total of 40 cores. To speed up computation, we enforced monotonicity in the strategies after each iteration, as this makes the number of possible beliefs linear instead of exponential in the grid size.

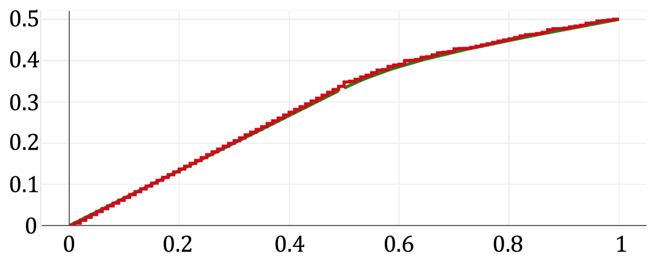

Figure 1: Bidding strategies for the first round of the first-price sequential sales auction with 3 bidders, 2 goods, and reserve prices $r_1 = 0, r_2 = 0.5$. Green denotes the analytical solution and red our found strategies. Types are plotted on the x-axis, bids on the y-axis.

chooses ascending reserve prices $r_1 \leq \cdots \leq r_K$ for each round. The equilibrium strategies are explained in Appendix F. We consider the cases where $n \in \{3, 4\}$, $K = 2$ and $(r_1, r_2) = (0, 0.5)$. The L_2-distances are given in Table 1 and the first-round bidding strategy for the three bidder case is plotted in Figure 1 (see Figures 21–32 in Appendix G for the other plots). Across all settings, the L_2 distance to the known PBE is very small; in most settings, the non-zero distance is due to the finite gridsize of the piecewise-constant strategies.

Two-Round Split-Award Auction

Our algorithm can also handle reverse auctions. Consider an auctioneer that can either procure 100% of a business from one bidder or 50% each from two bidders. Their type θ_i is the cost of providing 100%. All bidders share a cost parameter $C \in (0, 1)$ determining their cost $C\theta_i$ for providing the first 50%. For $C > 0.5$ ($C < 0.5$), the setting has economies (diseconomies) of scale. The auctioneer, unaware of C, cannot decide a priori whether to buy 100% from one or 50% each from two bidders and instead performs a combinatorial auction. Each bidder i submits both a split bid $b_i^{(sp)}$ for 50% and a sole bid $b_i^{(so)}$ for 100%. The first round has two possible outcomes:

- **Split award:** If $2 \min_i b_i^{(sp)} \leq \min_i b_i^{(so)}$, the auctioneer procures the first half of the business from the bidder with the lowest split bid. Next, the auctioneer runs a first-price sealed-bid auction for the remaining split. In the second round the first-round winner w has cost $(1 - C)\theta_w$.
- **Sole award:** If $2 \min_i b_i^{(sp)} > \min_i b_i^{(so)}$, The auctioneer procures the whole business from the most competitive sole bidder and the auction ends.

We study the setting for $n \in \{3, 4\}$. The bidders' costs are distributed uniformly on $[1, 2]$. We choose $C = 0.2$; this ensures *strong diseconomies of scale* for which Kokott, Bichler, and Paulsen (2019) derive the PBE that can be found in Appendix F. In Table 1, we provide the L_2 distances. Figures 33–36 in Appendix G show the plotted strategies. As before, our algorithm closely recovers the known equilibria.

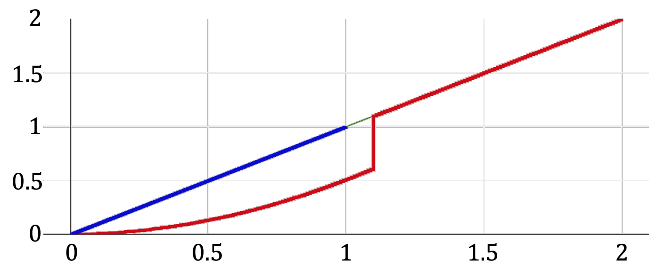

Figure 2: Bidding strategies in the first round of the sequential LLG auction. Green denotes truthful bidding, blue the strategy of the local bidder and red the strategy of the global bidder. Types are plotted on the x-axis, bids on the y-axis.

Two-Round Local-Local-Global Auction

We study a sequential version of the Local-Local-Global (LLG) auction (Ausubel and Milgrom 2006), without previous analytical results. There are two goods A and B sold sequentially via second-price sealed-bid auctions to three bidders. The two local bidders are single-minded. Bidder 1 wants A and bidder 2 wants B. Their values are uniformly distributed on $[0, 1]$. Bidder 3 is the global bidder whose value for the bundle $\{A, B\}$ is distributed uniformly on $[0, 2]$. In the first round, only bidders 1 and 3 bid; in the second, only bidders 2 and 3 bid.

We run our algorithm with a grid size of 200 and 150 best response iterations in each round. Figure 2 shows the computed first-round strategies (see Figure 38 in Appendix G for the second-round strategies). Our verification procedure computes a small ε of $1.47 \cdot 10^{-4}$, indicating near-perfect convergence. In the second round, both bidders bid truthfully. In the first round, the global bidder shades the bids for types less than 1.1 due to the uncertainty of winning B in the second round. In contrast, for types larger than 1.1 the global bidder—likely to win B as well—bids truthfully. This interesting, non-continuous behavior shows the potential of our approach for studying less-understood auction formats.

7 Conclusion

In this work, we have introduced a new algorithm for computing pure-strategy ε-perfect Bayesian equilibria in sequential auctions. We have shown how to abstract and decompose a sequential auction such that its PBEs can be computed via backward induction, and how the output can be verified.

Our method complements concurrent works using RL to compute equilibria in auctions. Instead of first fixing all bidders except i and using deep RL to compute i's best response (solving the Bellman equation and abstracting the game using neural networks), we start with an equilibrium Bellman equation and abstraction that simplify our problem to computing a finite number of BNEs. Only then do we iteratively compute best responses—in single-round games rather than multi-round ones. By reversing this order, we obtain a simple verification procedure and a stronger equilibrium concept (ex-interim PBE vs ex-ante Nash).

Our approach opens up exciting future directions. Combining our decomposition and verification with other methods (such as double oracle or deep RL) could lead to scalable, yet verifiable approaches. On the theory side, an interesting open question is how to address interdependent types and mixed strategies and whether there is an auction class for which these algorithms provably converge.

Acknowledgments

V. Thoma is supported by an ETH AI Center Doctoral Fellowship and acknowledges funding from the Swiss National Science Foundation (SNSF) Project Funding No. 200021-207343. This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant agreement No. 805542).

References

Ashenfelter, O. 1989. How Auctions Work for Wine and Art. *Journal of Economic Perspectives*, 3(3): 23–36.

Ausubel, L.; and Milgrom, P. 2006. The Lovely but Lonely Vickrey Auction. In Cramton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial Auctions*. MIT Press.

Bichler, M.; Fichtl, M.; Heidekrüger, S.; Kohring, N.; and Sutterer, P. 2021. Learning equilibria in symmetric auction games using artificial neural networks. *Nature Machine Intelligence*, 3(8): 687–695.

Bichler, M.; Fichtl, M.; and Oberlechner, M. 2023. Computing Bayes Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, 294. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701047.

Bichler, M.; Lunowa, S. B.; Oberlechner, M.; Pieroth, F. R.; and Wohlmuth, B. 2023. On the Convergence of Learning Algorithms in Bayesian Auction Games.

Bosshard, V.; Bünz, B.; Lubin, B.; and Seuken, S. 2017. Computing Bayes-Nash Equilibria in Combinatorial Auctions with Continuous Value and Action Spaces. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. Melbourne, Australia.

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.

Bosshard, V.; Wang, Y.; and Seuken, S. 2018. Non-decreasing Payment Rules in Combinatorial Auctions. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. Stockholm, Sweden.

Brown, N.; Bakhtin, A.; Lerer, A.; and Gong, Q. 2020. Combining Deep Reinforcement Learning and Search for Imperfect-Information Games. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 17057–17069. Curran Associates, Inc.

Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890.

Bünz, B.; Lubin, B.; and Seuken, S. 2022. Designing Core-Selecting Payment Rules: A Computational Search Approach. *Information Systems Research*, 33(4): 1157–1173.

Cai, Y.; and Papadimitriou, C. 2014. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 895–910. ACM.

Daskalakis, C.; Goldberg, P. W.; and Papadimitriou, C. H. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2): 89–97.

d'Eon, G.; Newman, N.; and Leyton-Brown, K. 2024. Understanding Iterative Combinatorial Auction Designs via Multi-Agent Reinforcement Learning. arXiv:2402.19420.

Gale, I.; and Stegeman, M. 2001. Sequential Auctions of Endogenously Valued Objects. *Games and Economic Behavior*, 36: 74–103.

Gale, I. L.; Hausch, D. B.; and Stegeman, M. 2000. Sequential Procurement with Subcontracting. *International Economic Review*, 41(4): 989–1020.

Ganzfried, S.; and Sandholm, T. 2009. Computing Equilibria in Multiplayer Stochastic Games of Imperfect Information. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI'09, 140–146. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Gong, Q.; Tan, X.; and Xing, Y. 2013. Ordering sellers in sequential auctions. *Review of Economic Design*, 18(1): 11–35.

Greenwald, A.; Li, J.; and Sodomka, E. 2012. Approximating Equilibria in Sequential Auctions with Incomplete Information and Multi-Unit Demand. In Pereira, F.; Burges, C.; Bottou, L.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.

Heinrich, J.; Lanctot, M.; and Silver, D. 2015. Fictitious Self-Play in Extensive-Form Games. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 805–813. Lille, France: PMLR.

Hendon, E.; Jacobsen, H. J.; and Sloth, B. 1996. Fictitious Play in Extensive Form Games. *Games and Economic Behavior*, 15(2): 177–202.

Hooke, R.; and Jeeves, T. A. 1961. "Direct Search" Solution of Numerical and Statistical Problems. *Journal of the ACM (JACM)*, 8(2): 212–229.

Katzman, B. 1999. A Two Stage Sequential Auction with Multi-Unit Demands. *Journal of Economic Theory*, 86(1): 77–99.

Klemperer, P. 2002. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46(4–5): 829–845.

Kohring, N.; Pieroth, F. R.; and Bichler, M. 2023. Enabling First-Order Gradient-Based Learning for Equilibrium Computation in Markets. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 17327–17342. PMLR.

Kokott, G.-M.; Bichler, M.; and Paulsen, P. 2019. The beauty of Dutch: Ex-post split-award auctions in procurement markets with diseconomies of scale. *European Journal of Operational Research*, 278(1): 202–210.

Krishna, V. 2002. *Auction Theory*. Academic Press.

Landi, M.; and Menicucci, D. 2018. Sequential Auctions with Decreasing Reserve Prices. *The B.E. Journal of Theoretical Economics*, 19(1).

Laraki, R.; Renault, J.; and Sorin, S. 2019. *Mathematical Foundations of Game Theory*. Springer.

Li, Z.; and Wellman, M. P. 2021. Evolution Strategies for Approximate Solution of Bayesian Games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6): 5531–5540.

Martin, C.; and Sandholm, T. 2023. Finding Mixed-Strategy Equilibria of Continuous-Action Games without Gradients Using Randomized Policy Networks. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 2844–2852. International Joint Conferences on Artificial Intelligence Organization. Main Track.

McAfee, R.; and Vincent, D. 1993. The Declining Price Anomaly. *Journal of Economic Theory*, 60(1): 191–212.

McMahan, H. B.; and Gordon, G. J. 2007. A Fast Bundle-based Anytime Algorithm for Poker and other Convex Games. In Meila, M.; and Shen, X., eds., *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, 323–330. San Juan, Puerto Rico: PMLR.

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

Milgrom, P.; and Weber, R. 2000. *A Theory of Auctions and Competitive Bidding II*. Edward Elgar Publishing Ltd.

Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337): 508–513.

Nayyar, A.; Gupta, A.; Langbort, C.; and Başar, T. 2014. Common Information Based Markov Perfect Equilibria for Stochastic Games With Asymmetric Information: Finite Games. *IEEE Transactions on Automatic Control*, 59(3): 555–570.

Ouyang, Y.; Tavafoghi, H.; and Teneketzis, D. 2017. Dynamic Games With Asymmetric Information: Common Information Based Perfect Bayesian Equilibria and Sequential Decomposition. *IEEE Transactions on Automatic Control*, 62(1): 222–237.

Parkes, D. C.; and Singh, S. 2003. An MDP-Based Approach to Online Mechanism Design. In Thrun, S.; Saul, L.; and Schölkopf, B., eds., *Advances in Neural Information Processing Systems*, volume 16. MIT Press.

Pieroth, F. R.; Kohring, N.; and Bichler, M. 2023. Equilibrium Computation in Multi-Stage Auctions and Contests. *ArXiv*, abs/2312.11751.

Puterman, M. L. 2005. *Markov decision processes : discrete stochastic dynamic programming*. Wiley series in probability and statistics. Hoboken, N.J: Wiley, 2nd ed. edition. ISBN 0471727822.

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

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

Reeves, D. M.; and Wellman, M. P. 2004. Computing Best-Response Strategies in Infinite Games of Incomplete Information. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, 470–478. Arlington, Virginia, USA: AUAI Press. ISBN 0974903906.

Rodriguez, G. E. 2009. Sequential Auctions with Multi-Unit Demands. *The B.E. Journal of Theoretical Economics*, 9(1).

Sandholm, T. 2015. Abstraction for Solving Large Incomplete-Information Games. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*. Austin, TX, USA.

Sandholm, T.; and Singh, S. 2012. Lossy stochastic game abstraction with bounds. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. Valencia, Spain.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.

Thoma, V.; Curry, M.; He, N.; and Seuken, S. 2023. Learning Best Response Policies in Dynamic Auctions via Deep Reinforcement Learning. arXiv:2312.13232.

Vorobeychik, Y.; and Wellman, M. P. 2008. Stochastic Search Methods for Nash Equilibrium Approximation in Simulation-Based Games. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '08, 1055–1062. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9780981738116.

Zhang, B.; and Sandholm, T. 2021. Subgame solving without common knowledge. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 23993–24004. Curran Associates, Inc.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007. Regret Minimization in Games with Incomplete Information. In Platt, J.; Koller, D.; Singer, Y.; and Roweis, S., eds., *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc.