



Computing pure Bayesian-Nash equilibria in games with finite actions and continuous types

Zinovi Rabinovich^{a,b}, Victor Naroditskiy^b, Enrico H. Gerding^{b,*}, Nicholas R. Jennings^{b,c}

^a Department of Computer Science, Bar-Ilan University, Ramat Gan, 52900, Israel

^b Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, United Kingdom

^c Department of Computing and Information Technology, King Abdulaziz University, Saudi Arabia

ARTICLE INFO

Article history:

Received 31 May 2011

Received in revised form 17 August 2012

Accepted 15 September 2012

Available online 18 September 2012

Keywords:

Algorithmic game theory

Bayes–Nash equilibrium

Epsilon-Nash equilibrium

Fictitious play

Simultaneous auctions

ABSTRACT

We extend the well-known fictitious play (FP) algorithm to compute pure-strategy Bayesian-Nash equilibria in private-value games of incomplete information with finite actions and continuous types (G-FACTs). We prove that, if the frequency distribution of actions (fictitious play beliefs) converges, then there exists a pure-strategy equilibrium strategy that is consistent with it. We furthermore develop an algorithm to convert the converged distribution of actions into an equilibrium strategy for a wide class of games where utility functions are linear in type. This algorithm can also be used to compute pure ϵ -Nash equilibria when distributions are not fully converged. We then apply our algorithm to find equilibria in an important and previously unsolved game: simultaneous sealed-bid, second-price auctions where various types of items (e.g., substitutes or complements) are sold. Finally, we provide an analytical characterisation of equilibria in games with linear utilities. Specifically, we show how equilibria can be found by solving a system of polynomial equations. For a special case of simultaneous auctions, we also solve the equations confirming the results obtained numerically.

Crown Copyright © 2012 Published by Elsevier B.V. All rights reserved.

1. Introduction

We study the problem of finding a symmetric pure Bayesian-Nash equilibrium in static games (i.e., where decisions are made simultaneously by all players) of incomplete information with independent private values (where the utility of a player depends only on the actions performed by others and not on their type), continuous type spaces and finite action spaces. Existing analytical results for such games mostly focus on auctions, a special case of incomplete information games. However, despite extensive research in this area, the developed theory has little to offer in terms of equilibrium derivation beyond the simplest models such as a single auction selling one or multiple homogeneous items (for an overview of the results, see [1]). On the computational side, solvers have been designed primarily for games of complete information (e.g., [2–4]), and can be applied to games of incomplete information with only a small number of actions and types. The main contribution of this paper is an algorithmic technique for computing Bayesian-Nash equilibria in games of incomplete information. We show its efficacy in simultaneous auctions, an important game of which only special cases were solved before. On the analytical side, we provide a novel characterisation of equilibria in a large class of games. This characterisation allows us to derive all equilibria for small simultaneous auction games confirming computational findings.

* Corresponding author.

E-mail addresses: zr@zinovi.net (Z. Rabinovich), vn@ecs.soton.ac.uk (V. Naroditskiy), eg@ecs.soton.ac.uk (E.H. Gerding), nrj@ecs.soton.ac.uk (N.R. Jennings).

In more detail, our computational technique is an extension of the fictitious play (FP) algorithm [5,6] to games of incomplete information with continuous types. Fictitious play was initially proposed as an iterative method for computing equilibria in zero-sum games of complete information. In each iteration, the algorithm chooses a best response to the frequency distribution of actions from previous iterations. If this frequency distribution, known as FP *beliefs*, converges, the converged distribution yields a *mixed* strategy Nash equilibrium of the game (see, e.g., [7]). Building on this, we develop an algorithm that generalises fictitious play to a wide class of games of incomplete information. Unlike regular fictitious play, if our algorithm converges, a *pure*-strategy equilibrium is produced.

Following much of the game-theoretic literature (see, e.g., [1,8]), we focus on *symmetric* games (where all players have the same type-dependent utility function, action space, type space, and type distribution) with single-dimensional types.¹ Our goal is to find a pure symmetric equilibrium, which is known to exist in this class of games under very mild assumptions (see Section 3 for details). The class of games we consider includes a wide range of commonly studied static games of incomplete information. Examples include single-sided auctions, double auctions, Cournot/Bertrand duopoly with asymmetric information and negotiation with incomplete information. Whereas we assume a continuous type space, our algorithm requires the space of *actions* to be finite. In fact, in many cases, such as auctions with discrete bids (consider the auctioneer *stepping* up the price in an English auction), finite action spaces are inherent to the problem, yet more difficult to analyse theoretically. Furthermore, the combination of finite actions and continuous type distributions guarantees (see Section 3) existence of a pure equilibrium, and also simplifies the representation of distributions over actions, i.e. FP beliefs.

While the steps of the fictitious play algorithm are the same in games of complete and incomplete information, novel challenges, such as recovering a pure equilibrium strategy from the converged beliefs, and computing a best-response action distribution, arise in the latter class of games. Unlike games of complete information, the converged frequency distribution in incomplete information games does not correspond to a single mixed strategy.² Moreover, it is not very useful to study mixed equilibria (as opposed to pure ones) for the types of games we consider. This is because, in games of incomplete information with continuous types, if a mixed equilibrium exists and under mild conditions on the type distributions, there always exists a corresponding pure-strategy equilibrium (resulting in the same action distribution), the latter being a more practical and desirable solution concept.³

To this end, we start by proving that, for converged FP beliefs, there exists a pure-strategy equilibrium generating that distribution. This theoretical result applies to converged beliefs which may only be observed asymptotically. In practice, we can only run a finite number of iterations of a FP procedure, never reaching an exact convergence. Therefore, we need a way to compute equilibria from FP beliefs that have not completely converged. In this case, we turn to approximate equilibria: after each iteration of our algorithm, we check whether we can produce an ϵ -equilibrium strategy from current beliefs. For this, we need an algorithm that converts FP beliefs into a strategy, such that the action distribution resulting from this strategy is the same as the beliefs. We design such an algorithm, which we call *BeliefsToStrategy*, for games where the agent's utility is linear in single-dimensional type (see Section 3 for formal definitions). Linearity in type is a standard assumption in most commonly studied single-parameter games including all forms of single-item auctions where an agent's type denotes the value for receiving the item (Appendix A discusses how our technique can be applied to domains with multi-dimensional types such as multi-unit or combinatorial auctions). When applied to converged beliefs, *BeliefsToStrategy* produces a pure-strategy equilibrium. Furthermore, if a sequence of beliefs is converging, *BeliefsToStrategy* yields an ϵ -equilibrium for any ϵ after a finite number of iterations.

We illustrate the power of our approach by finding equilibria in an important and previously unsolved problem: simultaneous sealed-bid auctions. In particular, we study a complete spectrum of combinatorial preferences, from perfect substitutes to perfect complements. We choose simultaneous auctions as it is a well-known fundamental model that has received attention in the literature before. However, previously both analytical and numerical results were obtained only for special cases [10,11].

Finally, in order to benchmark our numerical results, we provide an analytical characterisation of the equilibrium for games with linear utility functions in terms of a system of polynomial equations. Using this characterisation, we show how, for the simultaneous auctions setting with two players, two bids, and two auctions, the system can be solved analytically, providing exact equilibrium and uniqueness results. We then show that this derived equilibrium matches the results obtained numerically.

Against this background, our contributions to the state-of-the-art are as follows.

- We extend fictitious play to games of incomplete information with finite actions and single-dimensional continuous types. We prove that, whenever fictitious play beliefs converge, there exists a pure-strategy Bayesian-Nash equilibrium consistent with the converged beliefs.

¹ In Appendix A.1 we relax the symmetry assumption. The assumption of single-dimensional types is further discussed in Appendix A.2.

² This is because, in the case of incomplete information, a strategy is a mapping from type to actions, and there is a continuum of mappings that result in the same distribution of actions.

³ Pure equilibrium is a preferred solution concept as it is conceptually simpler and it does not rely on the ability of a player to randomise (see, e.g., discussion on mixed equilibria implementation in [9]).

- We provide an algorithm that converts converged beliefs into a pure equilibrium strategy for games with linear utility functions. We also show that, using this algorithm, if beliefs converge asymptotically, we can obtain an ϵ -equilibrium for any $\epsilon > 0$ within finite time.
- We find equilibria in a prominent, yet previously unsolved, auction game — simultaneous auctions for items ranging from perfect substitutes to perfect complements.
- We show that an equilibrium in a wide class of games with finite actions and continuous types can be found by solving a system of polynomial equations. Using this characterisation, we derive an equilibrium for a special case of simultaneous auctions.

The remainder of the paper is structured as follows. We begin with a review of related work in Section 2. Our model of games of incomplete information is formally stated in Section 3. A generalised fictitious play algorithm for these games is presented in Section 4. In Section 5 we provide a best-response algorithm and a procedure for converting FP beliefs to a strategy for games with utility functions linear in type. Section 6 applies our approach to a simultaneous auctions model. Finally, an analytical characterisation along with an exact derivation for the special case of two auctions, two players, and two bid levels appears in Section 7. Section 8 concludes.

2. Related work

This section provides an overview of fictitious play literature as well as other methods for finding Nash equilibria. The key distinction between ours and extant work is that our technique applies to games of *incomplete* information with *continuous* types. Note that games of incomplete information with discrete types can be viewed as games of complete information with a separate player for each possible type (see, e.g., Definition 26.1 in [12]). However, this representation is exponential in the number of types, making techniques for complete information games applicable to only very small game instances. A few techniques, which we review below, have been developed specifically for incomplete information games with discrete types. However, they also become intractable as the number of types increases.

In more detail, fictitious play was initially proposed as an iterative method for computing equilibria in static zero-sum games of complete information. It was subsequently shown to converge in several restricted classes of games, such as potential games [13] and bi-matrix $2 \times N$ games [14]. For instance, the work of Monderer and Shapley [15] shows FP convergence in a restricted class of complete information games (specifically, games that are *response equivalent to identical payoff games*). A related method is no-regret learning (or regret matching) where a player compares actions based on their average performance in the past (see, e.g., [16,17]). This method has been shown to converge to a Nash equilibrium in the same restricted class of games where FP converges (see, e.g., [18]). However, in general, the frequency distribution of actions produced by this method converges to the set of correlated equilibria, which is a weaker solution concept than Nash.

The literature mentioned above applies to settings with complete information. This setting is well studied, and a number of other general-purpose solvers exist for computing Nash and correlated equilibria [2–4,19–21]). In contrast, there are many fewer solution algorithms for incomplete information games, though some (e.g. [22–24]) can handle (or be adapted to) incomplete information games with discrete finite type sets at the expense of computational feasibility. Notice, however, that they are still inapplicable to the setting of our paper, since we focus on games with continuous types.

To address the issue of scalability, compact representations such as tree games [25] (where the utility structure induces a set of dependencies between players that form a tree), Bayesian Action-Graph Games (BAGGs) [26], and Multi-Agent Influence Diagrams (MAIDs) [27] have been developed to exploit the game structure: e.g., independence of type distributions and symmetry. An additional feature of the latter two approaches is their ability to make the game structure available to general-purpose solution algorithms. In particular, Jiang et al. [26] show how two different algorithms, the global Newton method [28] and the simplicial subdivision method [29], can be used with BAGGs, and demonstrate experimentally that these algorithms can result in exponential speedup. However, both BAGGs and MAIDs rely on the fact that the type spaces of the game they encode are discrete and finite.

Furthermore, unlike the case of BAGGs and MAIDs, most representations and solution algorithms impose strong restrictions on each other, which consequently limits the class of games they can be efficiently (if at all) applied to. For example, Koller et al. [24] had to convert an extensive form game into a sequence form⁴ in order to supply a payoff matrix to the underlying Lemke's algorithm. While linear in the size of the extensive form's tree, the number of action sequences in Koller's conversion can be exponential in the number of information sets of the game, which significantly impacts the scalability of the overall algorithm. In normal form games with infinite strategy spaces, Stein et al. [30] had to either limit the scope to just two players or approximate the solution by discretising the strategy space. Reeves and Wellman [31] restrict attention to games with two players and piecewise-uniform type distribution and apply an iterated best response to search for Bayesian-Nash equilibrium.

Another related area of research has resulted from the international poker competition [32], which has inspired a number of generally applicable algorithms for solving games of incomplete information. For instance, the counterfactual regret

⁴ Sequence form is a game description similar to normal form, where action sequences replace pure strategies. For typed games it assumes discrete and finite type space, hence is inapplicable to our domain.

algorithm was developed by Zinkevich et al. [33], and a method combining fictitious play with value iteration was proposed by Ganzfried and Sandholm [34]. Furthermore, Hawkin et al. [35] focus on transforming a game with a continuum of actions into a smaller game, and develop a new regret-minimisation algorithm to solve this game which builds on the counterfactual regret algorithm from [33]. These papers differ from our approach in that they view poker as a game with a finite discrete type space, and their algorithms rely on this property. Another approach is by Ganzfried and Sandholm [36], who formulate the problem as a mixed integer linear feasibility program. Their algorithm requires the set of types to be finite (and the number of constraints increases linearly with the number of types), but the authors then discuss how the approach can be extended to deal with continuous types. However, this extension requires the type distributions to be piecewise linear, and additional constraints are needed for each segment. By contrast, our algorithm is specifically designed for settings with continuous type spaces, and does not rely on assumptions about the shape of the distribution. Furthermore, different from our approach, their obtained equilibrium is a mixed one (whereas our algorithm always produces a pure-strategy equilibrium). In addition, their approach relies on having a qualitative model of the domain, which means that the number of intervals that divide the type space, as well as the actions associated with each of these intervals, are known. In contrast, our algorithm assumes no such knowledge.

Closer to the settings considered in this work, Gerding et al. [11] applied a variant of fictitious play called *smoothed fictitious play* to find mixed strategy equilibria in simultaneous auctions selling perfect substitutes when the number of bidder types is small. By contrast, here we consider continuous types and show that FP can be used to find *pure* equilibria. The FP algorithm for finding pure equilibria in games with incomplete information was first introduced in our previous work where the algorithm was applied to simultaneous auctions with perfect substitutes [37]. The current paper builds on [37] and significantly extends that paper. In particular, we introduce, for the first time, the `BeliefsToStrategy` algorithm to recover the pure strategy from the beliefs; we formally prove several properties of this algorithm; we extend the analysis of single-sided simultaneous auctions to a range of combinatorial preferences, from perfect substitutes to perfect complements; finally, we provide an analytical characterisation of the equilibrium strategy for small settings.

Our work also contributes to the literature on analytical derivations of equilibrium bidding strategies in the domain of simultaneous single-sided auctions, and auctions with discrete bids. Simultaneous Vickrey auctions selling complementary goods are studied in [10]. There, a distinction is made between local bidders, who only participate in one given auction, and global bidders who can participate in all auctions. The equilibrium and resultant market efficiency are derived for a model where each auction contains both global and local bidders. The model studied in [10] is further extended to the case of common values in [38]. The model we consider is more general in that it also applies to games other than auctions, and we obtain solutions for a variety of preferences, from perfect complements to perfect substitutes. The case with substitutable goods is studied by [39] in a setting restricted to three sellers and two global bidders and with each bidder having the same value (and thereby knowing the value of other bidders). The space of symmetric mixed equilibrium strategies is derived for this special case. Another setting where bidders face multiple simultaneous sealed-bid auctions is studied in e.g. [40–42]. These papers assume that bidders bid in only a single auction and choose this auction with some probability (where this probability depends on the reserve prices of the auctions). In [43], however, it was shown that choosing a single auction is not optimal. Specifically, if all other bidders choose only one auction, and when their types are sampled from distributions with non-decreasing hazard rates (which includes a wide range of common distribution functions including uniform, normal and exponential), the best response is always to place non-zero bids in all auctions. Our paper differs from [43] since we consider equilibrium behaviour, whereas the analysis in [43] is decision theoretic. Moreover, the analysis is limited to perfect substitutes and empirical evaluation, and relies on discretising the type space. In [44] the authors attempt to find the equilibrium strategies for this setting using iterative best response, but they show that, in fact, the strategies never converge.

Finally, a number of researchers have investigated auctions with discrete bid levels. A first-price auction for a single item is considered in [45]. There, equilibrium is characterised, and revenues are compared for different increments, defining sets of evenly spaced discrete bids. Discrete bids that are not necessarily uniformly spaced are studied in [46] in the context of a second-price auction for a single item. A special case of our characterisation of the best response for linear utilities appears there for a single action/item case. Our analytical characterisation goes beyond the case of single-item auctions, allowing derivation of equilibria in previously unsolved problems such as simultaneous auctions (see Section 7).

3. Games with finite actions and continuous types

We consider symmetric games of incomplete information with a finite number of actions and players with single-dimensional types, where types are sampled from a continuous type space. A game consists of n players, and the set of players is denoted by N . Each player draws his type $\theta \in \Theta \subset \mathbb{R}$ independently from a commonly known continuous distribution over Θ with density f , and a corresponding cumulative distribution F . Without loss of generality, we take the type space to be $\Theta = [0, 1]$. The same finite set of actions $A = \{a_1, \dots, a_m\}$ is available to each player. We adopt a standard independent private value model where the utility of a player is independent of the types of other players, and of the identity of the player performing an action (only the action matters, not who executed it). Therefore, the utility of a player is a function that depends on his type, his action, and the actions of the other players, $u : \Theta \times A^n \rightarrow \mathbb{R}$. For our theoretical results, we furthermore require that the utility is continuous in θ . The tuple $\Gamma = \langle N, A, u(\cdot), \Theta, F(\cdot) \rangle$ then defines a

Bayesian game.⁵ In the following, we refer to this setting as Games with Finite Actions and Continuous Types (G-FACTs). Our algorithm works in the context of G-FACTs as described here, but note that some of these assumptions can be relaxed. In particular, we discuss extensions to asymmetric games and multi-dimensional types in [Appendix A](#). The assumptions of independent private values and finite actions are inherent to the algorithm.

As is common in literature on Bayesian games, we study symmetric Bayesian-Nash equilibria: i.e., equilibria where all players follow the same strategy (see, e.g., Chapters 2, 3, 6, 7 in [\[1\]](#)). A pure strategy $s : \Theta \rightarrow A$ is a function that specifies an action for each player type. We denote by S the set of all strategies $s : \Theta \rightarrow A$. Letting $X = (X_1, \dots, X_{n-1}) \in \Theta^{n-1}$ denote the random variables representing the types of the other $n - 1$ players, the expected utility of a player of type θ playing action a_i when all other players follow the strategy s is $\mathbb{E}_{\{X_j \sim F\}_{j=1}^{n-1}}[u(\theta, a_i, (s(X_1), \dots, s(X_{n-1})))]$.

Instead of expressing the expected utility in terms of the strategies of other players, it is more convenient to use an equivalent representation in terms of the distribution of actions of the other players. The latter representation allows us to take advantage of the finiteness of the action space enabling an efficient best response calculation. The action distribution resulting from a strategy is derived as follows. Let $s^{-1}(a_i) \subseteq \Theta$ denote the set of all types playing action a_i . The probability that an agent's type is from this set is $h_s(a_i) = \int_{s^{-1}(a_i)} f(x) dx$, and $h_s \in \Delta(A)$ is the distribution of actions resulting from an agent playing s . When all other agents follow the same strategy s , the expected utility of an agent of type θ playing action a_i can be written as:

$$\hat{u}(\theta, a_i, h_s) = \mathbb{E}_{\{Y_j \sim h_s\}_{j=1}^{n-1}}[u(\theta, a_i, (Y_1, \dots, Y_{n-1}))]. \quad (1)$$

The expected utility from playing a strategy $s'(\cdot)$ when everyone else plays a strategy $s(\cdot)$ is $\tilde{u}(s', h_s) = \mathbb{E}_\theta[\hat{u}(\theta, s'(\theta), h_s)]$.

Example 1. To illustrate the notation, and to give an example of a game from the G-FACT class, consider a simple, single-item first-price auction with n bidders, each bidder's value uniformly distributed in $\Theta = [0, 1]$ ($F = U(0, 1)$), and 4 discrete bids from 0 to 3 dollars ($A = \{0, 1, 2, 3\}$). Furthermore, the strategy is given by:

$$s(\theta) = \begin{cases} a_1 = 0 & \text{if } 0 \leq \theta < 0.2, \\ a_2 = 1 & \text{if } 0.2 \leq \theta < 0.3, \\ a_3 = 2 & \text{if } 0.3 \leq \theta < 0.65, \\ a_4 = 3 & \text{if } 0.65 \leq \theta \leq 1. \end{cases}$$

Given that θ is uniformly distributed, the action distributions are as follows: $h_s(a_1) = 0.2$, $h_s(a_2) = 0.1$, $h_s(a_3) = 0.35$, and $h_s(a_4) = 0.35$. Suppose that a player with type θ derives a utility of $(3\theta - a_i)$ if she wins the item, and 0 otherwise. Furthermore, considering a fair tie breaking rule, the probability of winning when the agent ties with j other agents is $\frac{1}{j+1}$. Note that a player wins the auction if either all other bids are lower (and thus, $j = 0$), or if $n - j - 1$ bids are lower and the remaining j (excluding his own) bids are equal and she wins the tie. Furthermore, there are $\binom{n-1}{j}$ ways to choose j bidders to tie with. Each such tie occurs with probability $(h_s(a_i))^j (\sum_{a_k < a_i} h_s(a_k))^{n-j-1}$, where the first term ensures that there are j bids equal to the agent's bid, and the second term ensures that all other bids are lower. Then the expected utility of a player with type θ when playing action a_i , given $n - 1$ other bidders is:

$$\hat{u}(\theta, a_i, h_s) = (3\theta - a_i) \sum_{j=0}^{n-1} \binom{n-1}{j} \frac{1}{j+1} (h_s(a_i))^j \left(\sum_{a_k < a_i} h_s(a_k) \right)^{n-j-1}. \quad (2)$$

For example, the probability of winning the auction when placing action a_1 is equal to $\frac{1}{n} h_s(a_1)^{n-1}$. Note that, due to the tie breaking rule, we need to sum over all possible values of $1 \leq j \leq n$, and multiply this by the number of possible occurrences using the binomial coefficient. This calculation becomes more onerous when we consider multiple simultaneous auctions in [Section 6](#).

Now, as mentioned earlier, we are interested in finding equilibrium strategies. We define the necessary terms in the following.

Definition 1. A strategy $s : \Theta \rightarrow A$ is a symmetric pure-strategy equilibrium of a game Γ if:

$$\tilde{u}(s, h_s) \geq \tilde{u}(s', h_s) \quad \forall s' \in S.$$

Some of our results do not produce an exact equilibrium. In those cases, we use an approximate equilibrium defined below.

⁵ Our notation for agents' utility and type exploits the fact that the game is *symmetric*: i.e., each agent has the same type and action spaces and the utility of each agent is independent of his identity. Specifically, each agent's utility is given by the same function u , which is not indexed by i .

Definition 2. A strategy $s : \Theta \rightarrow A$ is a symmetric pure-strategy ϵ -equilibrium of a game Γ if:

$$\tilde{u}(s, h_s) + \epsilon \geq \tilde{u}(s', h_s) \quad \forall s' \in S.$$

It will sometimes be convenient to state this definition in terms of deviations in actions for each type rather than deviations in strategies. Definition 1 can be re-stated using deviations in actions (see Definition 8.E.1 and Proposition 8.E.1 in [8]).

Definition 3. A strategy $s : \Theta \rightarrow A$ is a symmetric pure-strategy equilibrium of a game Γ if for almost⁶ every $\theta \in \Theta$ (w.r.t. F):

$$\hat{u}(\theta, s(\theta), h_s) \geq \hat{u}(\theta, a_i, h_s) \quad \forall a_i \in A.$$

We note that limiting our analysis to symmetric pure-strategy equilibria does not impact the space of games we can solve.

Proposition 1. (From [49,50].) Every G-FACT has a symmetric pure-strategy equilibrium.

Proof. G-FACTs belong to a larger class of games where a pure-strategy symmetric equilibrium is known to exist, if two conditions hold [49,50]: (a) the distribution of types is continuous; (b) $\hat{u}(\theta, s(\theta), h_s)$ is continuous in type θ and in distribution of actions h_s . The first condition holds for G-FACT by definition. The second also follows from the assumption that the utilities $u(\theta, a_1, \dots, a_n)$ are continuous in θ , and in addition from the fact that \hat{u} is an expectation of such functions and the corresponding space of events (i.e. the probability that certain combinations of actions are played) is finite. In more detail, \hat{u} is a (finite) linear combination of continuous functions $u(\theta, \dots)$, therefore \hat{u} is itself continuous in θ . Furthermore, because h_s dictates the coefficients of the linear combination that defines \hat{u} , \hat{u} is also continuous in h_s . As a result, a G-FACT can always be solved in terms of a pure-strategy symmetric equilibrium.

In addition, any mixed equilibrium in such games can be converted into a pure-strategy equilibrium using a purification procedure [50]. Intuitively, since a strategy is a mapping from types to actions, such a purification involves finding a pure-strategy mapping which results in the same action distribution as the original mixed strategy. \square

4. Fictitious play for G-FACTs

In this section we extend fictitious play, which finds mixed equilibria in some complete information games, to search for pure strategy equilibria in G-FACTs. Before doing so, we consider the basic FP algorithm as it applies to normal form games with complete information.

In detail, at each iteration t , the FP algorithm consists of the following two steps:

- Compute best response: given the *belief* that the mixed strategy of the opponent is s'_t , calculate a best response s_t .
- Update beliefs: merge s'_t and s_t into a new mixed strategy s'_{t+1} .

These steps are then repeated until some convergence criterion is satisfied. A standard way to perform the merge in the second step is by averaging all best responses observed thus far. As a result, the influence of any subsequent best-response strategy diminishes with time. However, other approaches are suggested, e.g. using a weighted average, where higher weights are assigned to more recent strategies, or using a sliding window average, where only a small list of recent best responses is kept (see [51] for a discussion on these variations).

The algorithm in Fig. 1 generalises the two steps described above to symmetric games of incomplete information with finite actions and continuous types. An input to the algorithm is initial beliefs, h^0 , about the action distribution. At each iteration t , the best-response strategy is computed (line 3) with respect to the beliefs about the action distribution of an opponent, h^t (since we search for symmetric equilibria, each opponent draws his action from the same distribution). The algorithm for computing a best-response strategy is referred to as `BestResponse`, and the details of the algorithm depend on the specific domain (since the types are continuous and we cannot simply enumerate all possible strategies as with discrete type spaces). In Section 5.1 (see Fig. 4) we provide an instantiation of the algorithm for efficiently computing the best response for the setting with linear utility functions. Formally, s is a best-response strategy (or simply a “best response”) to an action distribution h if:

$$s(\theta) \in \arg \max_{a_i \in A} \hat{u}(\theta, a_i, h) \quad \forall \theta. \quad (3)$$

⁶ “Almost” in this context means that the probability of all types for which the strategy s does not prescribe an optimal action is zero [47,48].

Algorithm FictitiousPlay
Input: game $\Gamma = \langle N, A, u(\cdot), \Theta, F(\cdot) \rangle$, initial beliefs h^0 , update rule κ
Output: if converges, equilibrium strategy

```

1: set iteration count  $t = 0$ 
2: repeat
3:   strategy  $s = \text{BestResponse}(\Gamma, h^t)$ 
4:   compute the corresponding action distribution:
      $\forall a_i \in A: h_s(a_i) = \int_{s^{-1}(a_i)} f(x) dx$ 
5:   update beliefs:
      $\forall a_i \in A: h^{t+1}(a_i) = \kappa(t)h^t(a_i) + (1 - \kappa(t))h_s(a_i)$ 
6:   set  $t = t + 1$ 
7: until converged
8: return BeliefsToStrategy( $h^{t+1}$ )

```

Fig. 1. Fictitious play algorithm for symmetric games of incomplete information.

Once the best response, s , is obtained, its corresponding action distribution, h_s , is calculated (line 4), and the beliefs of the next iteration, h^{t+1} , are generated using an update rule (e.g., a standard update rule, $\kappa(t) = \frac{t}{t+1}$). If beliefs converge (line 7), a pure-strategy equilibrium can be obtained from these beliefs (line 8) as we prove in Theorem 1. We provide an algorithm termed `BeliefsToStrategy` for recovering an equilibrium strategy for the case of linear utilities in Fig. 5.

Convergence in fictitious play occurs if $h^t \rightarrow h^*$ as $t \rightarrow \infty$. This asymptotic convergence is called *convergence in beliefs* (see, e.g., [7]). In games of complete information, h^* is both the frequency distribution of actions and an equilibrium mixed strategy. In the incomplete information games studied here, h^* is just the frequency distribution of actions, and does not explicitly correspond to an equilibrium strategy (we demonstrate this in Section 6.2.3). Nevertheless, a pure-strategy equilibrium can be recovered from the converged beliefs h^* by taking the best response to it as stated below.

Theorem 1. *If fictitious play beliefs converge $h^t \rightarrow h^*$ as $t \rightarrow \infty$, then there is a strategy s^* that is a best response to the converged beliefs h^* and that induces h^* as its action distribution $h_{s^*} = h^*$; i.e., s^* is an equilibrium strategy.*

Proof. To prove the theorem, we first consider the mapping from an action distribution, h , to the set of action distributions produced by all best-response functions to h . This step is also used in Theorem 2 of [49] to prove existence of an equilibrium distribution of actions, i.e. h_{s^*} . We then proceed by assuming the converse: h^* is not a member of distributions produced by best-response strategies to h^* . We show that it contradicts convergence of h^t to h^* . Details of the proof follow.

Given that the action space $A = \{a_1, \dots, a_m\}$ is finite, all distributions over A form a simplex $\Delta(A)$. Denote for $1 \leq i \leq m$, $e_i \in \Delta(A) \subset \mathbb{R}^m$ a vector with i th element set to one, $e_i^i = 1$, and all others to zero, $e_i^j = 0, \forall i \neq j$. Let $E = \{e_i\}_{i=1}^m$. Then for any type θ , and a distribution of opponent actions h , the set of (pure) best responses can be described in terms of the following correspondence:

$$\Phi(\theta, h) = \{e_i \in E \mid u(\theta, a_i, h) \geq u(\theta, a_j, h) \forall a_j \in A\}.$$

Φ is non-empty and closed-valued, and upper hemicontinuous. By integrating over θ we can obtain the set of all action distributions produced by best responses to h . Let $\Psi(h) = \int \Phi(\theta, h) f(\theta) d\theta$, i.e. the set of action distributions generated by different best-response strategies, where:

$$\int \Phi(\theta, h) f(\theta) d\theta = \left\{ \int \phi(\theta, h) f(\theta) d\theta \mid \phi : \Theta \times \Delta(A) \rightarrow E, \phi(\theta, h) \in \Phi(\theta, h) \right\}.$$

Since f is continuous, Ψ is non-empty, compact and convex-valued, and upper hemicontinuous.⁷

Now, assume that $h^* \notin \Psi(h^*)$, in other words no best-response strategy to h^* has the action distribution h^* . In this case, since Ψ is compact and convex-valued, h^* can be separated from $\Psi(h^*)$. Intuitively, it means that the distance from h^* to any distribution generated by a best response to h^* , although small, is not negligible. More formally, there exist two open neighbourhoods, U_1 of h^* and U_2 of $\Psi(h^*)$, so that the following holds:

$$h^* \in U_1, \quad \Psi(h^*) \subset U_2, \quad U_1 \cap U_2 = \emptyset.$$

Furthermore, these neighbourhoods can be chosen so that there exists $\epsilon > 0$ such that U_1 is an open ball of radius ϵ , $U_1 = B_\epsilon(h^*)$, and the distance between U_1 and U_2 , $d(U_1, U_2)$, is at least ϵ . In addition, since Ψ is upper hemicontinuous, we can reduce ϵ to guarantee that $\forall h \in U_1, \Psi(h) \subset U_2$. In other words, best responses to distributions close to h^* have action distributions that are very close to those generated by best responses to h^* itself.

⁷ Relevant theorems, their origin and application to equilibria analysis can be found in the book by Hildenbrand [48]. Specifically, see Theorem 4 on p. 64 and Proposition 7 on p. 73.

Notice that since Ψ is compact and convex-valued, there is a constant $c > 0$ such that for any $h \in U_1, h' \in \Psi(h) \subset U_2$, and $0 < \lambda < 1$ it holds that $d(h, U_2) > d((1 - \lambda)h + \lambda h', U_2) + c\lambda$.

Now, let T be such that $d(h^T, h^*) < \epsilon$; i.e., $h^T \in B_\epsilon(h^*)$. During an FP update, $s(h^T)$ is a best response to h^T , and $h^{T+1} = \frac{T}{T+1}h^T + \frac{1}{T+1}h_s$, where $h_s(a_i) = \int \mathbb{1}(s(\theta) = a_i)f(\theta)d\theta$. By definition $h_s \in \Psi(h^T)$ and $\Psi(h^T) \subset U_2$, thus $h_s \in U_2$. As argued above $d(h^{T+1}, U_2) + c\frac{1}{T+1} < d(h^T, U_2)$. Since $\sum_t \frac{1}{t} = \infty$ and in each FP iteration the distance between h^T and U_2 decreases by $c\frac{1}{T+1}$, there exists $t > T$ so that $h^t \notin B_\epsilon(h^*)$. This contradicts convergence of h^t to h^* . Therefore, $h^* \in \Psi(h^*)$.

Since $h^* \in \Psi(h^*)$, there exists a selection function $s^*: \Theta \rightarrow A$, so that almost everywhere $u(\theta, s^*(\theta), h^*) > u(\theta, a_j, h^*) \forall a_j \in A$, and $h^*(a_i) = h_{s^*}(a_i) = \int \mathbb{1}(s(\theta) = a_i)f(\theta)d\theta$. In other words s^* is the best-response strategy to the action distribution it produces, and hence an equilibrium. \square

Corollary 1. *If the best response s^* to h^* is unique, then it is an equilibrium strategy.*

Equilibrium properties of fictitious play apply only to asymptotically converged beliefs. In numerical simulations, we are limited to a finite number of iterations and have to deal with approximate convergence. Consider a natural measure to establish convergence, which occurs once the (Euclidean) distance between h^t and h^{t+1} falls below some *convergence error*. This, however, is not a reliable convergence measure, as there is no guarantee that the distance in beliefs does not exceed the convergence error in later iterations.

To avoid the problem with identifying convergence in beliefs in a finite number of iterations, we can instead check at each iteration if an ϵ -equilibrium has been reached.⁸ This is done by constructing a strategy from the beliefs and checking if that strategy is an ϵ -equilibrium strategy. We provide a procedure for converting beliefs to strategy in the next section for the case of linear utilities.

5. Applying FP to G-FACTs with linear utilities

In this section we instantiate the `BestResponse` and `BeliefsToStrategy` algorithms for a particular setting where a player's *expected* utility, $\hat{u}(\theta, s(\theta), h_s)$, is linear in his type, θ . Note that our `FictitiousPlay` algorithm from Fig. 1 does not rely on linearity, and other procedures can be developed for non-linear settings. However, linearity is actually common in many games: in particular, it is inherent in all single-parameter models where the type of an agent denotes the value an agent receives in a “winning set of outcomes” (e.g., when the agent wins the item in an auction or when a public project is undertaken).⁹ This can be seen in Example 1 from Section 3 (see Eq. (2)), where the expected utility is linear in θ . This includes not only all one-shot single-item auctions (e.g., first-price, second-price, all-pay, see [1]), but also the simultaneous auctions studied in Section 6.

In the following, we start by making a few observations about the structure of a best response when utilities are linear and provide an algorithm for finding it. We then use these results to construct an algorithm for converting converged beliefs to a pure-strategy equilibrium. Together with a convergence metric described below, these algorithms instantiate our `FictitiousPlay` algorithm for games with linear utilities.

5.1. Best response

When (expected) utilities are linear in θ for a given a_i and h , we refer to the expected utility functions $\hat{u}(\cdot, a_i, h)$ as *utility lines*, and these functions are of the form:

$$\hat{u}(\theta, a_i, h_s) = \theta \cdot \text{slope}(a_i, h_s) + \text{intercept}(a_i, h_s), \quad (4)$$

where the slope and y-intercept are constant for a given action a_i and action distribution h_s . In the following, let $L = \{\hat{u}(\theta, a_i, h_s)\}_{a_i \in A}$ denote the set of all utility lines. Each utility line can be represented by its slope and intercept, and so we will sometimes use $L = \{\sigma_i, \iota_i\}_{i \in \{1, \dots, m\}}$, where σ_i and ι_i are the slope and intercept associated with action a_i .

Now, in general, we can see that the best response corresponds to the actions associated with the *upper envelope* of the utility functions in L . Formally, an upper envelope is given by $u^*(\theta) = \max_{a_i \in A} \hat{u}(\theta, a_i, h)$. In the case of linear utility functions, this upper envelope consists of a piecewise linear function, where each line segment corresponds to a particular utility line (and each utility line corresponds to a particular action). Furthermore, the upper envelope is always convex (to see this, note that, for any two crossing lines, their upper envelope is convex).

Observation 1. *In the case of linear utility functions, the upper envelope is piecewise linear and convex.*

⁸ We discuss the appropriateness of this measure of convergence in Section 6.2.2.

⁹ Note that, since the expected utility (see Eq. (1)) is a linear combination of the individual utilities $u(\theta, a)$, the expected utility $\hat{u}(\theta, s(\theta), h_s)$ is linear in θ if the individual utilities are linear. Therefore, when we say that utilities are linear, this also means that the *expected* utility is linear.




Fig. 2. Utility lines for actions $A = \{a_1, a_2, a_3, a_4\}$ given action distribution h . The bold intervals are the best-response (A', c) : $A' = \{a_1, a_3, a_4\}$, $c = (c_1, c_2)$.

Algorithm UtilityLines
Input: game $\Gamma = \langle N, A, u(\cdot), \Theta, F(\cdot) \rangle$, distribution of actions h
Output: utility lines $L = \{\sigma_i, \iota_i\}_{i \in \{1, \dots, m\}}$

- 1: **for** $i = 1$ to m
- 2: Given game Γ , calculate the slope and intercept:
 $\sigma_i = \text{slope}(a_i, h)$
 $\iota_i = \text{inter}(a_i, h)$
- 3: **return** $\{\sigma_i, \iota_i\}_{i \in \{1, \dots, m\}}$

Fig. 3. Generating utility lines for a game Γ given a distribution of actions h , where a utility line is defined by its slope and intercept.

An example of a best response is shown in Fig. 2. More formally, the upper envelope can be represented as a partition of the type space $[0, 1]$ into intervals, each labelled with its utility line and corresponding action. Let $c \in [0, 1]^{m'-1} \mid 0 < c_1 < c_2 < \dots < c_{m'-1} < 1$ denote a partition into m' intervals $[0, c_1], [c_1, c_2], [c_2, c_3], \dots, [c_{m'-1}, 1]$ and let A' denote the set of actions $\{a'_1, \dots, a'_{m'}\} \subseteq A$, where $a'_j \in A$ is the best-response action on the interval $[c_{j-1}, c_j]$. Note that each action $a'_j \in A'$ maps to an action $a_i \in A$, but the indexing is different. Similarly, let $L' = \{\hat{u}(\theta, a'_i, h_s)\}_{a'_i \in A'} = \{\sigma'_i, \iota'_i\}_{i \in \{1, \dots, m'\}} \subseteq L$ denote the corresponding utility lines. Due to Observation 1, note that $\sigma'_i \geq \sigma'_j$ whenever $i > j$. In Fig. 2, the upper envelope is given by the triple (L', A', c) where $A' = \{a'_1, a'_2, a'_3\} = \{a_1, a_3, a_4\}$ and $c = (c_1, c_2)$. Then, the pair (A', c) describes the corresponding best response.

We are now ready to introduce the algorithms needed to compute the best response. The first step is to compute the utility lines, which depends on the rules of the game $\Gamma = \langle N, A, u(\cdot), \Theta, F(\cdot) \rangle$. We assume an algorithm for doing this is available (see Fig. 3), but cannot provide a specific algorithm since this depends on the details of the problem domain. We will, however, instantiate the algorithm for the simultaneous auctions setting in Section 6 (in which case the expected utility is given by Eq. (5)).

Given an algorithm for computing the utility lines, Fig. 4 presents an algorithm for computing the best response, which proceeds as follows. First, we generate a utility line for each action $a \in A$. Then, all lines are sorted according to their slope (line 2). For ease of exposition, we slightly abuse notation and refer to σ_1 as the lowest slope, followed by σ_2 , etc. Similarly, the action with the lowest corresponding slope is referred to as a_1 followed by a_2 , etc. The best response at $\theta = 0$ is selected in line 3. Note that this simply requires finding the utility line with the highest intercept (since $\hat{u}(0, a_i, h) = \iota_i$). This utility line forms the initial upper envelope (L', A', c) . Now, due to the convexity of the upper envelope (Observation 1), line segments at $\theta > 0$ need to have slopes of at least σ_i , which means that we can disregard any utility lines $j < i$. Hence, the for loop at line 5 starts with $j = i + 1$. In each iteration of the main loop (line 5), we consider whether to include the j th utility line in the current envelope (L', A', c) possibly replacing one or more previously added lines. Now, since the lines are considered in the order of their slope, there are only two possible cases: it can either lie entirely below the current envelope, in which case it has no effect on the upper envelope and can be disregarded; or, it intersects the envelope at a unique point $x \in (0, 1)$. Note that it cannot lie entirely above the envelope, since it has to be below the envelope at point $\theta = 0$ (otherwise, this would mean that $\iota_j > \iota_i$ which contradicts the maximisation in line 3) and that it cannot cross the envelope (L', A', c) at more than one point (its slope is higher than the slopes of L' , and, therefore, once it crosses (L', A', c) , it increases faster than any of the lines L' and does not cross them).

Whenever the j th utility line crosses the current envelope, this envelope is updated as follows. First of all, due to Observation 1, and since we know that σ_j is higher than any existing slope in L' , the line segment necessarily needs to appear at the end of the envelope. Therefore, since the intersection point is x , we can remove any utility lines in the current

Algorithm BestResponse
Input: game $\Gamma = \langle N, A, u(\cdot), \Theta, F(\cdot) \rangle$
distribution of actions h
Output: best response $s: \Theta \rightarrow A$ represented by (A', c)

- 1: generate utility lines $L = \{\sigma_i, \iota_i\}_{i \in \{1, \dots, m\}} = \text{UtilityLines}(\Gamma, h)$
- 2: sort the utility lines in increasing order of slope
let a_i denote the action with the i th lowest slope: $\sigma_1 \leq \sigma_2 \leq \dots$
- 3: find the index of the utility line that maximises utility at $\theta = 0$:
 $i = \arg \max_{j \in \{1, \dots, m\}} \hat{u}(0, a_j, h_s) = \arg \max_{j \in \{1, \dots, m\}} \iota_j$
- 4: let $\sigma'_1 = \sigma_i$, $\iota'_1 = \iota_i$, and $a'_1 = a_i$ and define the initial envelope as:
 $L' = \{\sigma'_1, \iota'_1\}$, $A' = \{a'_1\}$ and $c = ()$
- 5: **for** $j = i + 1$ to m
if $\hat{u}(\theta, a_j, h_s) = \theta \sigma_j + \iota_j$ intersects the envelope (L', A', c) at $x \in \Theta$
let $k \in \{1, \dots, |A'|\}$ denote the index of the intersected utility line in L' :
i.e., where $x \in [c_{k-1}, c_k]$, with $c_0 = 0$ and $c_{|A'|} = 1$
update the envelope (L', A', c) as follows:
5.1: remove all utility lines with index $z > k$ from (L', A', c)
5.2: append utility line with index j at position $k + 1$,
i.e. $\sigma'_{k+1} = \sigma_j$, $\iota'_{k+1} = \iota_j$, $a'_{k+1} = a_j$, and $c_k = x$
- 6: **return** (A', c)

Fig. 4. An algorithm for computing best response when agents' utilities are linear in a single-parameter type $\theta \in [0, 1]$.

envelope which appear after x . These are the lines $z > k$, and they are removed in line 5.1. Then, a new line segment is added to the envelope which intersects with the k th line at point x , and provides the best response for $\theta \in [x, 1]$ (i.e., the new utility line is appended at the end). Note that, since the utility lines are considered in the order of their slope, and new lines are always appended at the end, the resulting upper envelope is convex (as it should be).

Ignoring computation of the utility lines, the runtime of the BestResponse algorithm is dominated by line 5. Note that, to find the intersection point and the corresponding line segment in L' (if any) requires looping through all lines in L' , which has at most $m - 1$ values. Moreover, the for loop at line 5 also has at most $m - 1$ values. Therefore, the worst-case runtime is in the order $O(m^2)$. We note that the upper envelope can be computed in $m \log m$ time (see, e.g., [52]). However, we opt for a simpler implementation since it is efficient in practice (since m' is typically much smaller than m) and (as discussed in Section 6.2.2) the total run-time of one iteration of our algorithm is likely to be dominated by the computation of the individual utility lines.

To apply our method, one needs to be able to compute utility lines. Since there are m actions, the utility lines are represented by m slopes and intercepts. This is independent of other game parameters such as the number of players.¹⁰ The computation of utility lines is specific to the particular domain, and this can become a bottleneck. However, in practice, it is often possible to reduce the computation of a utility line by using a compact representation of a game (such as action-graph games [26]), but we cannot provide any general analysis. We do provide the details of how to efficiently compute utility lines in the domain we study in Section 6. There, computation is dominated by the domain-specific tie-breaking rule (see Appendix B).

5.2. Converting fictitious play beliefs to a pure strategy

Although Theorem 1 shows that an equilibrium strategy can be recovered from the limit of fictitious play beliefs, it does not provide an exact algorithm for doing so, but rather assumes that such a procedure exists. From a theoretical point of view this assumption is valid, since the necessary purification procedures are guaranteed to exist (see, e.g., [50]). However, by itself, existence of a procedure is insufficient to apply the algorithm. For this reason, in this section provide an explicit purification algorithm BeliefsToStrategy (see Fig. 5) for G-FACTs with type-linear utilities. In addition, based on the insight of this procedure, we point out in Appendix A the steps necessary for generalisation of this procedure to non-linear utility functions.

The BeliefsToStrategy algorithm in Fig. 5 constructs a strategy where action a_i is played with probability $h(a_i)$. Actions are sorted in ascending order of slopes of their utility lines and the action with the lowest slope a'_1 is played by the types $\theta \in [0, h(a'_1)]$. The next action a'_2 is played by the types $\theta \in [h(a'_1), h(a'_1) + h(a'_2)]$, etc. The algorithm has two important properties. First, for a converged distribution of actions h^* , BeliefsToStrategy(h^*) is an equilibrium. Second, if h is sufficiently close to a converged distribution, BeliefsToStrategy(h) is an ϵ -equilibrium strategy. These properties are the subject of the following two theorems.

Theorem 2. *If fictitious play beliefs converge $h^t \rightarrow h^*$ as $t \rightarrow \infty$, then a pure-strategy equilibrium can be constructed using the algorithm in Fig. 5.*

¹⁰ This is largely due to symmetry, but even in the asymmetric case (discussed in Appendix A.1), the number of utility lines is $n \cdot m$ and so scales linearly in the number of players.

Algorithm BeliefsToStrategy
Input: game $\Gamma = \langle N, A, u(\cdot), \theta, F(\cdot) \rangle$
distribution of actions h
Output: equilibrium strategy (A', c)

- 1: gather actions played with positive probability $\hat{A} = \{a_i \in A \mid h(a_i) > 0\}$
- 2: generate utility lines for actions \hat{A}
 $L = \{\sigma_i, \iota_i\}_{i \in \{1, \dots, |\hat{A}|\}} = \text{UtilityLines}((N, \hat{A}, u(\cdot), \theta, F(\cdot)), h)$
- 3: sort the utility lines in increasing order of slope
let a_i^h denote the action with the i^{th} lowest slope: $\sigma_1 \leq \sigma_2 \leq \dots$
and define the ordered set $A' = (a_1^h, \dots, a_{|A'|}^h)$
- 4: generate the strategy that produces action distribution h
define $c = (c_1, \dots, c_{|A'|-1}) \in R^{|A'|-1} \mid c_j = \sum_{i=1}^j h(a_i^h) \quad \forall 1 \leq j \leq |A'| - 1$
- 5: **return** (A', c)

Fig. 5. An algorithm for converting distribution of actions h into a pure strategy when agents' utilities are linear in a single-parameter type $\theta \in [0, 1]$.

Proof. As we have shown in Theorem 1, convergence of beliefs means that the distribution h^* is produced by an equilibrium strategy: i.e., there exists a best response s^* to h^* that generates h^* itself. We also know that the latter property holds for $\text{BeliefsToStrategy}(h^*)$, since it produces h^* . Therefore, to conclude that $\text{BeliefsToStrategy}(h^*)$ is an equilibrium, it remains to be shown that $\text{BeliefsToStrategy}(h^*)$ is a best response to h^* . To achieve this, we will compare the $\text{BeliefsToStrategy}(h^*)$ strategy and an arbitrary best response to h^* , s^* . We will show that either $\text{BeliefsToStrategy}(h^*)$ completely coincides with s^* , or has the same utility for all types, and therefore is also a best response to h^* .

Now, recall that in the case of linear utility lines, any best response can be expressed using the interval representation (A', c) . In this representation, the intervals (A', c) are listed in a non-decreasing order of their respective utility line slopes (see Observation 1). Furthermore, the length of interval $[c_{i-1}, c_i]$ is given by $h^*(a_i')$. But this is exactly the strategy produced by $\text{BeliefsToStrategy}(h^*)$. Hence, if the best response to h^* is unique, $\text{BeliefsToStrategy}(h^*)$ is necessarily this best-response strategy.

The only discrepancy between $\text{BeliefsToStrategy}(h^*)$ and a best response s^* to h^* may occur if some set of actions, $A'' \subset A'$, have the same utility line, when evaluated at h^* . In this situation, it is possible that $\text{BeliefsToStrategy}(h^*)$ and s^* choose different actions from A'' for some *disagreement* types. However, all these *disagreement* types will belong wlog. to a single interval K where all actions from A'' constitute a best response. Since, any action $a_i \in A''$ is a best-response for types in K , $\text{BeliefsToStrategy}(h^*)$ is a best response for all types in K .

Furthermore, actions in A'' are in a strict order (with respect to the utility line slope) with actions in $A' \setminus A''$. Hence, $\text{BeliefsToStrategy}(h^*)$ assigns actions in A'' for types in K , and only for types in K , and the same is true for any best response to h^* . If the remaining structure, i.e. the use of actions in $A' \setminus A''$, of the best response s^* is unique for the complement of K , \bar{K} , then this structure coincides with that of $\text{BeliefsToStrategy}(h^*)$ over \bar{K} . As a result, $\text{BeliefsToStrategy}(h^*)$ chooses best-response actions for all types, and is a best-response strategy to h^* for all types. If the best-response structure over \bar{K} using $A' \setminus A''$ is not unique, we repeat the argument regarding *disagreement* types for \bar{K} using $A' \setminus A''$. Since the set of actions is finite, there may be only a finite number of such iterations before the remaining best-response structure becomes unique. We conclude, that even if $\text{BeliefsToStrategy}(h^*)$ differs from s^* in its choices of actions for some types, this does not lead to reduction of utility. Therefore, $\text{BeliefsToStrategy}(h^*)$ is a best response to h^* .

Finally, since $\text{BeliefsToStrategy}(h^*)$ is a best response to h^* and generates h^* , it is the equilibrium strategy. \square

Theorem 3. Let $\{h^t\}_{t=1}^\infty$ be a converging fictitious play sequence, $h^t \rightarrow h^*$. Denote by s^t the best response to h^t calculated during iteration t of fictitious play. Then:

$$|\tilde{u}(\text{BeliefsToStrategy}(h^t), h^t) - \tilde{u}(s^t, h^t)| \rightarrow 0.$$

Proof. Notice that, if there are multiple best responses to h^* , their utility is necessarily the same. Since the utility lines $\hat{u}(\theta, a_i, h)$ are continuous in θ and h , and m is finite, their upper envelope is a continuous function in θ and h . Recall also that we have assumed the set of types to be compact and f continuous, hence $\tilde{u}(s, h) = \mathbb{E}_{\theta \sim f}[\hat{u}(\theta, s(\theta), h)]$ is also uniformly continuous in h for any (not necessarily best-response) strategy s (with a finite interval representation (A, c)). As a result, for any strategy s and for any $\delta > 0$ exists T such that for all $t > T$, $|\tilde{u}(s, h^t) - \tilde{u}(s, h^*)| < \delta$. In particular, we can choose T so that $|\tilde{u}(\text{BeliefsToStrategy}(h^t), h^t) - \tilde{u}(\text{BeliefsToStrategy}(h^t), h^*)| < \delta$. Furthermore, a similar result holds for $\tilde{u}(s[h], h)$, where $s[h]$ is a functional that returns a best-response strategy (e.g., the one from Fig. 4), such that $|\tilde{u}(s^t, h^t) - \tilde{u}(s^*, h^*)| < \delta$. To see this, recall that the utility of the best response is an upper envelope of a finite set of functions continuous in h , hence also continuous in h . As a result, $\tilde{u}(s[h], h)$ is uniformly continuous as a function of h , and the necessary inequality follows.

Now, define the following correspondence:

$$\bar{\Phi}(h) = \{s : \forall \theta, \hat{u}(\theta, s(\theta), h) \geq \hat{u}(\theta, a_j, h) \forall a_j \in A\}.$$

Notice that, similarly to Φ and Ψ from Theorem 1, this correspondence is non-empty, closed- and compact-valued and upper hemicontinuous. We define the distance between two strategies as $d(s_1, s_2) = \int \mathbb{1}(s_1(\theta) \neq s_2(\theta)) f(\theta) d\theta$. Then for any positive δ that is less than the probability of any action that is part of a best response to h' or h^* , i.e. holds that:

$$\delta < \min\left(\min_{a'_i \in A'} h'(a'_i), \min_{a^*_i \in A^*} h^*(a^*_i)\right),$$

there exists T such that for all $t > T$ it holds that for any $s_1 \in \bar{\Phi}(h^t)$ and $s_2 \in \bar{\Phi}(h^*)$, the distance $d(s_1, s_2) < \delta$. In particular, there exists $s^* \in \bar{\Phi}(h^*)$ such that $d(s^*, s^t) < \delta$. Since both strategies are a best response and the distance between the strategies is at most δ , the order of actions in (A', c) must be the same as in (A^*, c) . In particular, this means that $\text{BeliefsToStrategy}(h^t)$ and $\text{BeliefsToStrategy}(h^*)$ use the same order of actions in their interval structure. In fact, they differ only over a set of types of size $d(\text{BeliefsToStrategy}(h^t), \text{BeliefsToStrategy}(h^*)) < \delta$. Since utility is bounded, we have:

$$\begin{aligned} & |\tilde{u}(\text{BeliefsToStrategy}(h^t), h^*) - \tilde{u}(\text{BeliefsToStrategy}(h^*), h^*)| \\ &= |\tilde{u}(\text{BeliefsToStrategy}(h^t), h^*) - \tilde{u}(s^*, h^*)| < c\delta, \end{aligned}$$

where $c > 0$ is some constant. Aggregating all three bounds together we have:

$$\begin{aligned} & |\tilde{u}(\text{BeliefsToStrategy}(h^t), h^t) - \tilde{u}(\text{BeliefsToStrategy}(h^t), h^*)| < \delta, \\ & |\tilde{u}(\text{BeliefsToStrategy}(h^t), h^*) - \tilde{u}(s^*, h^*)| < c\delta, \\ & |\tilde{u}(s^t, h^t) - \tilde{u}(s^*, h^*)| < \delta. \end{aligned}$$

Hence, we obtain that, for any δ , there exists a T so that for all $t > T$, the following holds:

$$|\tilde{u}(\text{BeliefsToStrategy}(h^t), h^t) - \tilde{u}(s^t, h^t)| < c'\delta, \quad \text{for some finite } c' > 0. \quad \square$$

Theorem 3 guarantees that if FP converges, then an ϵ -Nash equilibrium is necessarily obtained at some finite iteration. Furthermore, the proof structure allows another practical simplification. Specifically, before applying BeliefsToStrategy , we can simplify h^t by filtering out all actions that appear with numerically negligible probability (i.e. below the threshold of δ) and renormalising.¹¹

6. Simultaneous auctions

In previous sections we discussed an algorithm for finding equilibria in G-FACTs, and an implementation for linear utility functions. In the current section, we apply the algorithm to a setting where bidders participate in multiple simultaneous, single-sided, sealed-bid auctions.¹² Simultaneous auctions are a natural generalisation of single-item auctions when multiple items are available for sale from different sellers. As we discussed in Section 2, existing computational techniques cannot be applied to this setting due to continuous type spaces, while discretisation of the type space comes at the expense of computational feasibility. That is, the settings with more than a few discrete types and bid levels are beyond the computational reach of most techniques.

The purpose of analysing this setting is two-fold. First, we demonstrate the efficacy of our algorithm for a complex setting where no analytical solution exists, and give convergence results. Second, we demonstrate that our algorithmic technique can be used to contribute to the auction literature by providing an extensive empirical characterisation of the equilibrium bidding behaviour in simultaneous auctions. This empirical analysis is augmented in the next section, where we derive an analytical characterisation for a basic setting and show that the equilibria found for that setting match those that are found with our numerical approach.

In particular, in our experiments we focus on simultaneous Vickrey (i.e., second-price) auctions. However, any other pricing (e.g., first-price or all-pay) could be chosen as this does not affect the algorithm (but affects the equilibrium strategies). The auctions are simultaneous in that a bidder needs to make a decision on how much to bid in each auction without knowing any of the outcomes (unlike sequential auctions where the winner of an auction is known before a bid is placed in another auction). For this setting, it has been shown in prior decision-theoretic work [43] that, even though each individual auction is incentive compatible (bidding the true value for the item being auctioned is a dominant strategy), and even when the items are perfect substitutes (the bidder does not derive extra benefit from winning more than one item), a bidder is often better off bidding in multiple auctions and shading their bids, as opposed to choosing a single auction

¹¹ A similar thresholding procedure was applied to mixed strategies in Ganzfried et al. [53].

¹² We note that a variation of the algorithm has also been successfully applied to a more complex double auction setting in [54]. See Section 8 for more details.

and bidding truthfully. Furthermore, in the case of substitutable¹³ goods, the bidding strategies are typically non-monotonic in type, which makes finding the equilibrium bidding strategies a challenging task. In this section, we extend this work to a game-theoretic analysis in which all players can participate in all auctions, and the aim is to compute an equilibrium strategy. Here, we consider a wide range of combinatorial structures, including substitutes and complements.

6.1. Simultaneous Vickrey auctions

We consider a setting with k simultaneous sealed-bid single-item Vickrey auctions. The items sold in different auctions are heterogeneous. The set of auctions (equivalently, items) is denoted by $K = \{1, \dots, k\}$. The set of players N corresponds to the bidders. Each bidder has a *single-dimensional* privately-known type θ which is i.i.d. sampled from a c.d.f. F with continuous support on $[0, 1]$. F is assumed to be common knowledge. The finite action space is given by a set of joint bids defined as follows. Each auction has a finite set of admissible bids levels $B \subset \mathbb{R}_+$, and a bidder chooses a bid for each auction.¹⁴ For simplicity, we furthermore assume that all auctions have the same bid levels and these are equally spaced (both of these simplifications can be trivially relaxed but we choose this restriction to reduce the number of parameters to consider). Thus, the action space is $A = B^k$. Note that simultaneous auctions with discrete bids and continuous types are an instance of G-FACTs since the action space is finite, but the types are continuous. The only piece missing from a full specification of a Bayesian game is the utility function, which we define next.

While the type of a bidder is single dimensional, we assume that the bidders have combinatorial preferences: i.e., the items are heterogeneous and may range from perfect substitutes to perfect complements and combinations of these. This is achieved by a function $\phi : 2^K \rightarrow \mathbb{R}$, common to all bidders which specifies a *complementarity structure* of the auctions. The *value* that a bidder with type θ derives from winning a subset of items $\eta \subseteq K$ is given by the product $\phi(\eta)\theta$. Notice that the relative values of bundles are the same across bidders. In essence, the type is a scaling parameter: if bidder 1 has type x and bidder 2 has type $3x$, this means that bidder 2's value for each bundle is 3 times as high as the value of bidder 1. We acknowledge that single-dimensional types are more restrictive than multi-dimensional types that allow each bidder to have his own complementarity structure. Nevertheless, our restricted model is a good approximation for scenarios where items are likely to have a common complementarity structure (e.g., the bundle of left and right shoes is valuable, while each item in isolation is not).

As an example consider the case of two auctions ($K = \{1, 2\}$) that we study in detail in the rest of the paper. Let $\alpha = \phi(\{1\})$, $\beta = \phi(\{2\})$, and $\gamma = \phi(\{1, 2\})$ denote the value from winning only the first auction, only the second auction, and both auctions, respectively. Then, having $\alpha = \beta = \gamma = 1$ corresponds to a setting of perfect substitutes and free disposal. That is, a bidder does not gain from winning multiple items, but there is no cost either (not including any additional payments from winning multiple auctions). Our computational and analytical techniques do not rely on the assumption of free disposal, and, for completeness, we consider complementarity structures where free disposal does not hold: i.e., the values $0 \leq \gamma \leq \min(\alpha, \beta)$. At its extreme, $\gamma = 0$, we have valuations where winning both items results in zero value (this could be interpreted as the cost of disposal of the second item being equal to the independent value of the first item). Furthermore, $\alpha = \beta = 0$ and $\gamma = 1$ represents the case with perfect complements. That is, a bidder only receives utility from winning both items. Finally, setting $\gamma = \alpha + \beta$ means that the items are independent.

In addition, we can model auctions selling heterogeneous items. For example, $\beta = 2\alpha$ and $2\alpha \leq \gamma \leq 3\alpha$ model the case when the item sold in the second auction is twice as valuable as the item from the first auction, and these items exhibit some degree of substitutability. Such preferences could arise, for example, when the same type of item is sold in two different quantities (e.g., 1-liter and 2-liter cartons of milk), but having both is more than a bidder typically needs.

The assumption of a common complementarity structure ϕ enables us to model combinatorial valuations while keeping bidder types single-dimensional. Alternatively, true combinatorial valuations would endow each bidder with his own combinatorial structure, making each type $2^{|K|}$ -dimensional. This more general model is left open for future work (our extension to multi-parameter domains in Appendix A.2 may be helpful). Thus, in the present work we focus on the common complementarity structure.

Given a complementarity structure ϕ , the expected utility of a bidder from playing action $a \in A$ is:

$$\hat{u}(\theta, a, h) = \theta \sum_{\eta \subseteq K} \phi(\eta) q(a, \eta, h) - \text{cost}(a, h) \quad (5)$$

where:

- $q(a, \eta, h)$ is the probability that playing an action (bids) $a \in A$ results in winning exactly the set η of auctions, given the distribution of actions from the opponents, h . In Appendix B we show how this probability is calculated (note that this is not trivial since we have discrete bids and therefore need to take into account a tie breaking rule).

¹³ Two items are substitutable if the utility from winning both of them is less than the sum of the utilities for each individual item. Similarly, two items are *complementary* if the utility from winning both of them is more than the sum of the utilities for each individual item.

¹⁴ We argue that having a finite set of bids is not necessarily restrictive in practice, since bids are often rounded to an appropriate level (e.g. to the nearest dollar amount for small bids, the nearest ten-fold for larger bids, etc.). In addition, the set of admissible bids can be further restricted by the auctioneer to increase seller revenue [46,55].

Table 1
Experimental settings.

Parameter	Value(s)
Number of auctions (k)	2
Number of bidders (n)	2, 5, 10
Number of bid levels per auction ($ B $)	2, 5, 10
Type probability distribution (F)	<i>Uniform</i> ([0, 1])
Complementarity structure (α, β)	varies
Complementarity structure (γ)	0, 0.05, 0.1, ..., 2.95, 3
Initialisation of FP beliefs	<i>random</i>
Number of runs for each setting	30
Number of FP iterations per run	5000

- $cost(a, h)$ is the expected payment when placing action a and given distribution h . [Appendix B](#) shows how this is calculated (as can be seen in the appendix, note that the expected payment is simply the sum of the expected payment for each auction, and so is much easier to calculate since this can be done independently for each auction).

Importantly, note that expected utility \hat{u} is linear in θ , which allows us to apply the algorithms from [Section 5](#).

6.2. Numerical results

In this section we present equilibrium results obtained by running the fictitious play algorithm described in [Figs. 1, 3, 4, and 5](#). The algorithm in [Fig. 3](#) is instantiated using [Eq. \(5\)](#). The details of the tie-breaking rule appear in [Appendix B](#) (throughout the analysis we use the exact tie breaking rule unless specified otherwise). In the following, we start with the experimental setup in [Section 6.2.1](#). Then, in [Section 6.2.2](#) we measure the empirical convergence of the algorithm to an ϵ -Nash equilibrium. The actual equilibria obtained are first discussed with only 2 bid levels per auction in [Section 6.2.3](#) for homogeneous items, and in [Section 6.2.4](#) for heterogeneous items. In [Section 6.2.5](#) these results are extended to more than 2 bid levels.

6.2.1. Experimental setup

A game is specified by the number of auctions, the number of bidders, a set of possible bids, a complementarity structure, and a distribution of agents' types. In all of the experiments, we focus on 2 auctions, and a uniform distribution of types between 0 and 1. We begin the numerical investigation with the simplest possible setting: 2 bidders, 2 bid levels per auction, and complementarity structures where the individual value of each item is the same (i.e., the items sold at both auctions are identical). In this setting, we find an equilibrium for each degree of complementarity from substitutes to complements. The observed equilibria for this simple setting enable us to identify some properties that continue to hold in the more complicated setting we consider next: auctions with more than 2 bidders, and auctions selling different items. We then further expand the setting by considering more than 2 bid levels.

An overview of various experimental settings is given in [Table 1](#). Although we tested with many other values as well, these are representative of the results that we obtained. The bid levels in B are equally spaced between 0 and 1. This means that, if the number of bid levels is 2, then $B = \{0, 1\}$. On the other hand, if this is set to 5, then $B = \{0, 0.25, 0.5, 0.75, 1.0\}$, etc. Recall that the number of bid levels is per auction. This means that, for example, if this number is 10, then the total number of actions for a bidder when there are 2 simultaneous auctions is $10^2 = 100$. Furthermore, *random* initialisation of the FP beliefs means that the initial probability of each action is set randomly between 0 and 1, and then normalised so that the probabilities sum to one (note that this is different from having each action played with equal probability). These values are sampled anew for different runs of the same experiment. Therefore, the (only) difference between runs is the initial FP beliefs (since the algorithm itself is deterministic). The aim in having multiple runs is to see whether or not different initial beliefs result in different equilibrium strategies being computed. When multiple runs converge to the same ϵ -equilibria, we are more confident that a true equilibrium has been identified as we describe next. We run each experiment 30 times to obtain statistically significant results based on 95% confidence intervals.

6.2.2. Convergence and scalability

In this section we empirically analyse to what extent the results converge, and the computational runtime required as we scale the number of bidders and bid levels. These results provide a useful insight into the practical applications of the algorithm. In more detail, we measure convergence in terms of the size of ϵ in the ϵ -Nash equilibrium (see [Definition 2](#) in [Section 3](#)). Recall that the ϵ of a given strategy s is given by the difference between the utility obtained by playing a best response s^* to h_s and the utility from playing s when the action distribution is h_s : $\epsilon(s) = \tilde{u}(s^*, h_s) - \tilde{u}(s, h_s)$. In particular, we would like to measure the ϵ of the strategy that can be constructed from the current FP beliefs, h^t , using the `BeliefsToStrategy` algorithm. Thus, we set $s = \text{BeliefsToStrategy}(h^t)$ and $h_s = h^t$. In addition, to obtain a




Fig. 6. Percentage of runs converged within a given error in simulations with $\alpha = 0.7$, $\beta = 1$, and the parameters n , $|B|$, and γ taking each of the values described in Table 1. For each combination of parameters, 30 simulations were run.

unit-free measure of convergence so that we can compare different settings,¹⁵ we use a standard approach to normalise the difference, resulting in the so-called *relative error* [56]:

$$\text{error}(t) = \frac{\tilde{u}(s^*, h^t) - \tilde{u}(\text{BeliefsToStrategy}(h^t), h^t)}{\tilde{u}(s^*, h^t)}. \quad (6)$$

Note that the error is guaranteed to be between 0 (the equilibrium) and 1 (as far as a strategy can be from the equilibrium).

The results using this measure appear in Fig. 6, which shows the percentage of runs that converge to a given error within a number of iterations, for all settings described in Table 1, and where each setting is run 30 times, and $\alpha = 0.7$ and $\beta = 1$ (the results are very similar for other values of α and β). This figure shows that virtually all runs of the algorithm converge to ϵ -equilibrium with a small error. Moreover, as the number of iterations increases, the percentage of runs that are within ϵ of the equilibrium keeps increasing. This indicates that, on average, once an ϵ -equilibrium for a given ϵ is reached, running extra iterations does not lead to divergence.

A potential weakness of the ϵ -equilibrium concept is that, even though the gain from deviation may be very small, the ϵ -equilibrium strategy may be arbitrarily far away from an exact equilibrium strategy (see, e.g., [57]). We address this concern in two ways. First, we run the same settings starting from different initial beliefs. If the algorithm consistently converges to the same strategy,¹⁶ this increases our confidence that the true equilibrium is obtained (note that the converse is not true, since converging to different strategies could simply mean that there exist multiple equilibria). We found that, using our algorithm, all of the simulations for 2 bid levels (Sections 6.2.3 and 6.2.4), as well as the simulations with more bid levels for auctions selling weakly complementary items (i.e., $\gamma \geq \alpha + \beta$), converged in the latter stronger sense.

Second, we compare the strategies with analytical results for settings where these can be derived. In particular, for a special case of 2 bidders, 2 auctions, and 2 bid levels, we are able to derive equilibria analytically (our derivation is discussed in Section 7.1). We see that the analytical results are identical to the equilibrium results obtained computationally in Section 6.2.3. Thus, in this special case, ϵ -equilibria obtained numerically are approximating exact pure-strategy equilibria. Although we do not have a theoretical proof of convergence in general, we note that the equilibria we obtain for variants of this special case (e.g., with more than 2 bidders or with heterogeneous items) follow the same structure, which we take as a reasonable evidence that these approximate equilibria are close (in terms of the strategy, not just the utility) to exact equilibria.

We now consider the amount of computation required, both in terms of the number of iterations, as well as time elapsed before convergence to an ϵ -Nash equilibrium. In these experiments, we choose $\text{error} = 0.01$ (i.e., where the error is no more than one percent of the total utility), but the results show similar trends for other values. In particular, we are interested to see how our algorithm scales for the simultaneous auctions setting with $k = 2$ auctions, when we increase the number of bidders and the number of discrete bid levels. Although each experiment is run 30 times with different initial beliefs, the results in this section show the average over the 15 runs with the lowest runtime. We do this because the results vary depending on the initial beliefs, and we see that, while in most cases the results converge within a couple of hundred iterations, there are a few outliers which skew the results and take much longer or do not converge to the required error within the maximum numbers of iterations (which was set to 2500 for these particular experiments). We avoid these outliers by taking the top half of the runs. Furthermore, we argue that, in practice, it is possible to run a number

¹⁵ Equilibrium utility for different complementary structures could be very different. Thus, the same absolute difference may constitute 1% of utility for one complementarity structure, and 200% for another.

¹⁶ As measured, for example, by a negligibly small Euclidean distance between the action distributions from different runs after a fixed number of iterations.

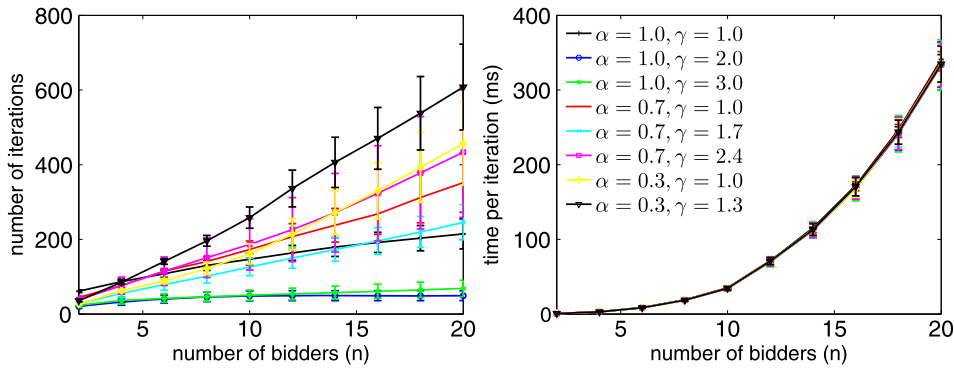


Fig. 7. Average number of iterations required to reach $error = 0.01$ (left) and time per iteration (right) for different values of n and different preference structures with $\beta = 1$. Here, we use the exact tie breaking rule (note that the number of bidders does not affect the time per iteration when using the approximate tie breaking rule). The number of bid levels is set to $|B| = 10$. Results are averaged over the 15 fastest of 30 runs. The error bars denote the 95% confidence intervals.




Fig. 8. Average number of iterations required to reach $error = 0.01$ (left) and time per iteration (right) for different values of $|B|$ and different preference structures with $\beta = 1$. The number of bidders is set to $n = 10$, and the approximate tie breaking rule is used. Results are averaged over the 15 fastest of 30 runs. The error bars denote the 95% confidence intervals.

of experiments in parallel to see which one converges first, which would have the same effect. All of the experiments were run on a Linux cluster with 2.27 GHz Nehalem processors and the simulation was implemented in Java.

There are two factors that determine the total computation time: the number of iterations required and the computation time for each iteration. The effect of the first factor can be seen in Figs. 7 (left) and 8 (left) which show the average number of iterations required to reach the equilibrium, as we increase the number of bidders, respectively the number of bid levels per auction, for a variety of preference structures. This result shows that the number of required iterations always increases as the number of bidders increases, but typically flattens out as we increase the number of bid levels. An increase in the number of iterations is indicative of the difficulty of the problem, and this suggests that problems with more bidders are more challenging to solve, which seems intuitive. In most cases, however, the increase seems to be linear or even sublinear, and so has relatively little impact on the final computation. Interestingly, the preference structure also has a large effect on the difficulty of the problem, and generally increasing the asymmetry between the two auctions increases the number of iterations required. On the other hand, increasing the number of bid levels merely increases the granularity and, for a given relative error, this has little effect on the number of iterations needed to converge.

Whereas the algorithm scales relatively well in terms of the number of iterations for the simultaneous auctions domain, it is less promising when we consider the computation required for each iteration. Here, the computation required is mainly due to computing the utility lines (the `UtilityLines` algorithm in Fig. 3, which requires finding the slope and intercept in Eq. (5)), and computing the best response (i.e., the `BestResponse` algorithm in Fig. 4).¹⁷ We first consider the effect of the number of bidders which, due to the tie breaking rule, affects the computation of utility lines. Note that, from Appendix B, we can see that, to compute the exact probability of winning, we need to consider all possible numbers of ties in each auction. As a result, for $m = 2$ auctions the computation required scales in the order $O(n^3)$ with the number of bidders. The empirical results in Fig. 7 (right) are for the same settings as before, and show the average real time (in seconds) required

¹⁷ Note that the computation of the `BeliefsToStrategy` algorithm is negligible compared to the other algorithms since the main part consists of sorting the actions by slope. Furthermore, the `BeliefsToStrategy` algorithm is only required to compute the relative error (Eq. (6)), and the strategy itself once the process has converged.




Fig. 9. Equilibrium strategy and corresponding utility for $\gamma = 1.4$ and $n = 2$.

to compute an iteration, as we increase the number of bidders. Note that, as we can expect, the number of bidders has a large impact on the computation, but the preference structure does not.

Clearly, we can do much better by simply using an approximation of the tie breaking rule, and a simple approximation which scales well with the number of bidders is given in Appendix B.1. Using this approximation, the increase in computation due to an increase in the number of bidders becomes negligible. Furthermore, we empirically consider the *additional* error (in terms of the ϵ -Nash) introduced by this approximation (we do so by computing the best response both with and without the approximation, and computing the error in both cases). From this we can determine that the error decreases and goes to zero as the number of bid levels goes to infinity, but empirically we find that the error is already very small for small numbers of bid levels. For example, the average additional error is less than 0.003 when the number of bid levels is 20.

In terms of the runtime when increasing the number of bid levels, in the case of $k = 2$ the number of actions is equal to $|B|^2$, where $|B|$ is the number of bid levels, and so the time complexity of a single iteration is at least $O(|B|^2)$. Furthermore, as discussed in Section 5.1, the time complexity for finding the best response using our algorithm is $O((|B|^2)^2) = O(|B|^4)$. This is consistent with the empirical results depicted in Fig. 8 (right) which show the time per iteration when using the approximate tie breaking rule, as the number of discrete bid levels increases.¹⁸ As a result, for $k = 2$, we can easily compute results for settings of 100 discrete bid levels per auction.

6.2.3. Equilibrium results for homogeneous items

We first consider a simple setting where the items sold at each auction are identical, the set of bids is $B = \{0, 1\}$, and there are only $n = 2$ bidders. Given that the auctions are identical, we set $\alpha = \beta = 1$, and vary the value of γ as specified in Table 1 and explained in Section 6.1. This value ranges from $\gamma = 0$, which models a setting of extreme substitutes without free disposal, to $\gamma = 3$ which corresponds to complements. In-between are perfect substitutes ($\gamma = 1$) and independent auctions ($\gamma = \alpha + \beta = 2$). In what follows we analyse the results after 5000 iterations of the fictitious play algorithm (this number was found to be sufficiently large for experiments to converge to a very small error).

To illustrate the results, Fig. 9 shows the strategy and corresponding utility lines generated by the *BeliefsToStrategy* procedure at the end of a particular run for a setting with $\gamma = 1.4$ (i.e., a representative value where agents have substitutable preferences). This figure shows that, for this setting, all 4 possible actions are played with non-zero probability. Moreover, as can be expected, agents with higher types bid higher. Specifically, agents with a low type play (0, 0); agents with a very high type play (1, 1), and in-between types play either (0, 1) or (1, 0). From this example we can see several interesting trends. First, the slopes of the utility lines are increasing as expected, except for actions (0, 1) and (1, 0) where the slopes seem to be identical. Second, the actions (0, 1) and (1, 0) seem to be played with equal probability (note that the type interval is of equal size and types are uniformly distributed). In fact, however, the fictitious play beliefs assign *almost* equal probabilities to (1, 0) and (0, 1), and the slopes are *almost* identical. Furthermore, the slopes oscillate: if at iteration t the action (1, 0) has a slightly higher slope, then at iteration $t + 1$ the action (0, 1) has a higher slope. This is because the best-response strategy also oscillates, and only one of the actions (1, 0) or (0, 1) is played with non-zero probability in best response, never both, and these two actions alternate. This illustrates why a special *BeliefsToStrategy* procedure is needed to find an equilibrium and why simply taking the best response does not result in an equilibrium strategy.

In more detail, this fluctuating behaviour mimics the FP dynamics in games of complete information such as matching pennies, where the best response for the mismatching player is heads whenever the probability of playing tails is above one

¹⁸ We note that there is considerable scope for optimising the code, e.g. by using a more efficient algorithm for computing the upper envelope or by detecting and removing dominated utility lines (see Section 5.1).




Fig. 10. Action distributions (i.e., FP beliefs after 5000 iterations) for auctions with $n = 2$ bidders selling homogeneous items.

Table 2
Strategy corresponding to beliefs h .

Value range	bid
$0 \leq v \leq h(0, 0)$	(0, 0)
$h(0, 0) \leq v \leq h(0, 0) + h(1, 0)$	(1, 0)
$h(0, 0) + h(1, 0) \leq v \leq h(0, 0) + h(1, 0) + h(0, 1)$	(0, 1)
$h(0, 0) + h(1, 0) + h(0, 1) \leq v \leq 1$	(1, 1)

half, and tails otherwise. There, the beliefs asymptotically approach equal probability of playing heads and tails yielding an equilibrium strategy. Similarly, here the FP beliefs for actions (1, 0) and (0, 1) become increasingly similar as the number of iterations increases. However, in contrast to complete information games where FP beliefs define a unique mixed strategy, in games of incomplete information, FP beliefs do not correspond to a single strategy. Thus, we need to convert FP beliefs into a strategy that induces these beliefs and is roughly a best response to them. The `BeliefsToStrategy` procedure accomplishes this goal as we proved in Section 5.2. Furthermore, in Section 7.1 we formally show that the strategy found by the `BeliefsToStrategy` procedure in fact corresponds to the analytically derived strategy for the case of $n = 2$ bidders, 2 bid levels, and 2 auctions. In particular, we can see that, in equilibrium the two actions (1, 0) and (0, 1) are always played with identical probability as expected.¹⁹

Next, we consider the equilibrium strategies for different values of γ . We can plot action distributions more concisely than equilibrium strategies and take advantage of `BeliefsToStrategy` to map each action distribution to a strategy. To this end, Fig. 10 plots action distributions for each value of the complementarity parameter γ between 0 and 3. This figure (and other figures that follow) shows action distributions (i.e., FP beliefs) after 5000 FP iterations, and averaged over 30 runs. We omit the error bars in the figures because the confidence intervals are very small and cannot be seen. This shows that, starting from different initial beliefs, the beliefs converge to the same action distribution. An equilibrium strategy can be recovered from the action distributions by applying the `BeliefsToStrategy` procedure. The resulting strategy appears in Table 2.

As can be seen, the action distributions appear to be continuous in the complementarity parameter γ (see Fig. 10). Furthermore, the values of γ can be partitioned into three intervals, each corresponding to a different set of actions being played with non-zero probability. For low values of γ (in the case of highly substitutable items), action (1, 1) is never played; in the mid-range, all 4 actions are played with non-zero probability; in the case of complementarities (1, 0) and (0, 1) are never played. We denote these intervals by $[0, \hat{\gamma}_1]$, $[\hat{\gamma}_1, 2]$, and $[2, \infty]$, where $\hat{\gamma}_1$ is the lowest value of γ for which the bid (1, 1) is played in equilibrium. As can be seen in Fig. 10, for this particular setting the value of $\hat{\gamma}_1 \approx 1.2$.²⁰ Furthermore, as soon as γ reaches the value of additive valuations ($\gamma = \alpha + \beta = 2$), the bids (1, 0) and (0, 1) are not played at all as the agents try avoiding winning a single item. Interestingly, this is consistent with existing analytical results in the literature for continuous bids whereby only equal-bid pairs are played for items that display complementarity [1] (when the auctions are identical).

Table 3 further analyses the strategy and shows that equilibrium action distributions are monotone in γ within each of the intervals (the values for $\gamma = \infty$ are based on simulations for large (but finite) values of γ). Furthermore, the probability of playing the “highest” possible bid of (1, 1) increases as the items become more complementary. In fact, we observe that,


¹⁹ It is worth noting that, in the homogeneous case, there actually exists a continuum of equilibrium strategies. This is because the agents are indifferent between playing (1, 0) and (0, 1). For example, the strategy where the intervals for (0, 1) and (1, 0) in Fig. 9 are swapped is also an equilibrium. Also, there exist equilibria with more intervals. However, all of these equilibria result in the same action distribution, and this action distribution is unique (see also Section 7.1).

²⁰ We derive the exact value in Section 7.1.

Table 3

Equilibrium analysis for homogeneous items and 2 bid levels.

	$\gamma \in [0, \hat{\gamma}_1]$	$\gamma \in [\hat{\gamma}_1, 2]$	$\gamma \in [2, \infty]$	$\gamma = \infty$
$f(0, 0)$	decreases	increases	decreases	0
$f(1, 0) = f(0, 1)$	increases	decreases	0	0
$f(1, 1)$	0	increases	increases	1

**Fig. 11.** Action distributions for auctions selling homogeneous items.

in the limit, the probability of bidding (1, 1) approaches 1 from below. However, for any finite γ , the bid (0, 0) is played in equilibrium by the types that are small enough, resulting in a positive $h(0, 0)$.

In the remainder of this section, we show that our technique can be used to derive equilibria for more than 2 bidders. The results for 5 and 10 bidders are shown in Fig. 11. We observe the pattern identified for 2 bidders continues to hold for 5 and 10 bidders. In particular, we observe the same types of interval, where only the value of $\hat{\gamma}_1$ (the lowest γ for which the action (1, 1) is played) changes. Furthermore, the monotonicity results shown in Table 3 are identical for these intervals. Comparing across the graphs, we notice that $\hat{\gamma}_1$ increases with the number of bidders. That is, the items must display more complementarity for (1, 1) to be played in equilibrium when there is more competition. This is a result of an increasing cost associated with the bid (1, 1): the same strategy s results in a higher (second) price as the number of bidders playing the strategy increases. This leads to a higher expected cost of winning with the bid of 1, discouraging bidding 1 unless the type is sufficiently high. This effect is also reflected in a lower probability of playing (1, 0), (0, 1), and (1, 1) across γ when $n = 10$ compared to when $n = 5$.

6.2.4. Equilibrium results for heterogeneous items

Our next set of results considers auctions selling different items and where an agent has different valuations for these items. To illustrate the effect of the degree of asymmetry, we run two sets of experiments for different relative values of the items. In the first set, the value of the item sold in one auction is 0.7 of the value of the item sold in the other auction ($\alpha = 0.7$, $\beta = 1$). In the second set of experiments, one item is much less valuable: its value is only 0.3 of the value of the other item ($\alpha = 0.3$, $\beta = 1$). Notice that the case of additive valuations, beyond which the items become complementary, occurs at $\gamma = \alpha + \beta$, which is $\gamma = 1.7$ for the first case and $\gamma = 1.3$ in the second.

Action distributions for each setting with 2, 5 and 10 bidders are plotted in Fig. 12. As before, these results are averaged over 30 runs. Since the item sold in the second auction is more desirable, we can see that, in equilibrium, (0, 1) is played more often than (1, 0): the curve $h(0, 1)$ is above $h(1, 0)$ for all values of γ . Furthermore, we note that, even though the actions (1, 0) and (0, 1) are played on adjacent intervals, the switching of the optimal best response in each iteration, which we observed in the homogeneous case, does not occur here. After sufficiently many iterations, the bid (1, 0) is always selected by lower types, and the bid (0, 1) is always selected by the higher types. Thus, Table 2 still provides an equilibrium strategy.

We observe similarities in the equilibrium structure of homogeneous and heterogeneous cases. After identifying the regions where the set of bids played with non-zero probability does not change, we notice that as in the homogeneous case, the probability of each bid within a region is monotonic in γ . The bids (1, 0) and (0, 1) are not symmetric when items are heterogeneous, resulting in more regions. For $n = 5$ and $n = 10$, there are five regions summarised in Table 4. Comparing the graphs for $n = 5$ and $n = 10$, we notice that as in the homogeneous case, the probability of bidding (0, 0) is higher for $n = 10$ while the probabilities of the other bids are lower. Comparing two complementarity structures $\alpha = 0.7$ and $\alpha = 0.3$, we observe that the bid (1, 0) is played more often when the item is worth 0.7 than when it is worth 0.3 (symmetrically, the bid (0, 1) is played less often). This corresponds to a higher competition for the item when it is more desirable.

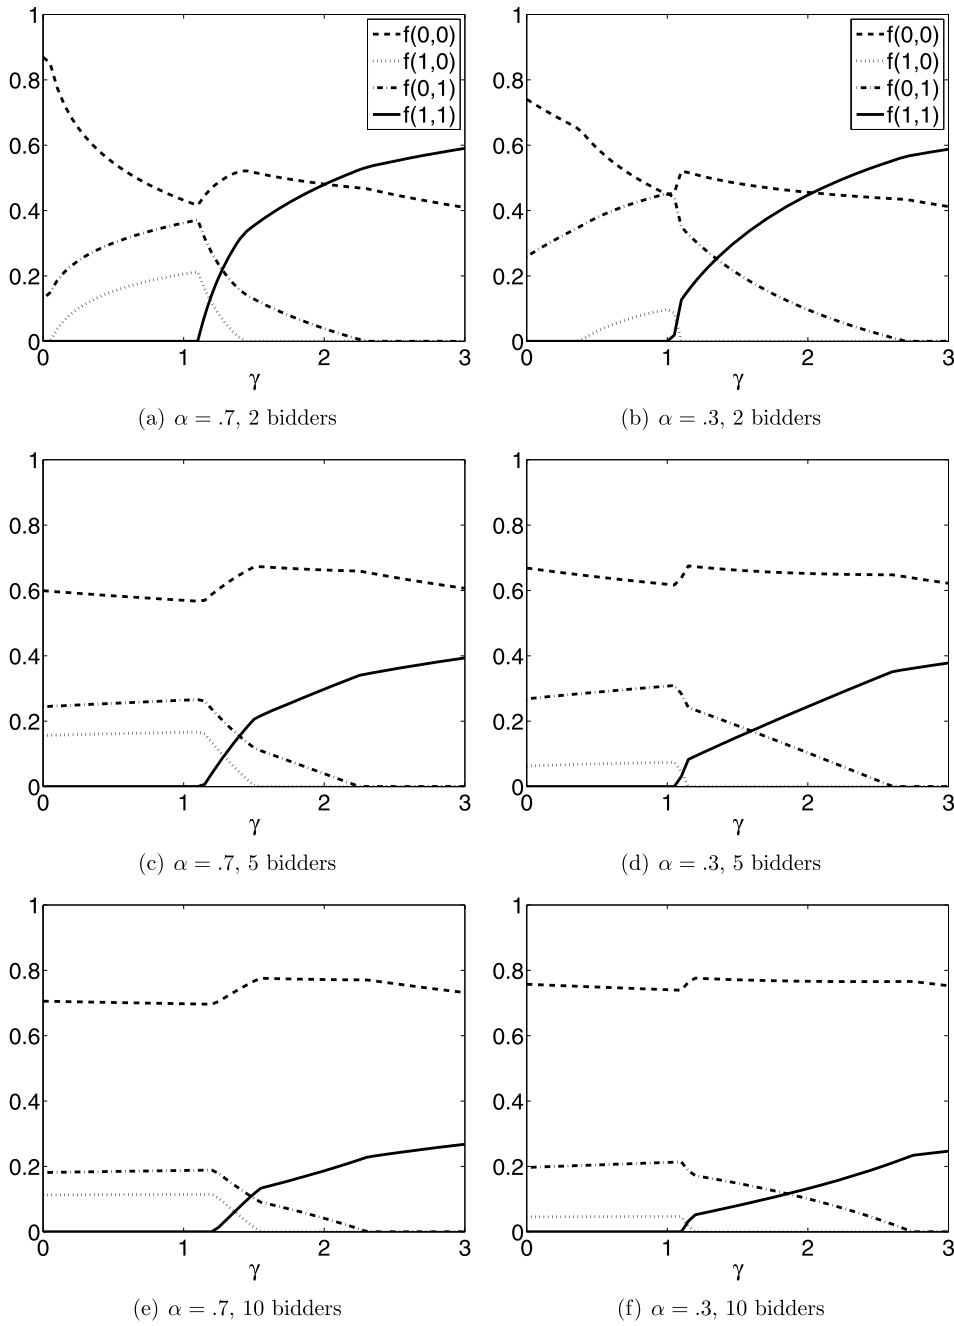


Fig. 12. Action distributions for auctions selling heterogeneous items.

Table 4

Equilibrium analysis for heterogeneous items and 2 bid levels. The bid (1, 1) is played with positive probability for $\gamma > \hat{\gamma}_2$. The bid (1, 0) is played with zero probability for $\gamma < \hat{\gamma}_1$ ($\hat{\gamma}_1$ is zero for $n = 5$ and $n = 10$) and $\gamma > \hat{\gamma}_3$. The bid (0, 1) is played with zero probability for $\gamma > \hat{\gamma}_4$.

	$\gamma \in [0, \hat{\gamma}_1]$	$\gamma \in [\hat{\gamma}_1, \hat{\gamma}_2]$	$\gamma \in [\hat{\gamma}_2, \hat{\gamma}_3]$	$\gamma \in [\hat{\gamma}_3, \hat{\gamma}_4]$	$\gamma \in [\hat{\gamma}_4, \infty]$	$\gamma = \infty$
$f(0, 0)$	decreases	decreases	increases	decreases	decreases	0
$f(1, 0)$	0	increases	decreases	0	0	0
$f(0, 1)$	increases	increases	decreases	decreases	0	0
$f(1, 1)$	0	0	increases	increases	increases	1




Fig. 13. Complementary items: equilibrium strategies for 5 bidders with 10 bid levels.




Fig. 14. Substitutable items: equilibrium strategies for 2 bidders with 5 bid levels.

6.2.5. Equilibrium results for auctions with more than two bid levels

The strategies analysed so far were limited to two bid levels. However, our technique can be applied to any number of bid levels. Here we discuss results for ten bid levels, but similar results hold for other bid levels. With more than two bid levels, there is no easy way to represent the results concisely for each value of γ as we did before (since the number of possible actions is large). Therefore, we select a few representative values of γ to illustrate the types of equilibria we find. To this end, Fig. 13(a) shows equilibria for homogeneous items with a small degree of complementarity. The bid submitted in each auction is plotted as a function of type. We see that, consistent with the two-bid and continuous case (see [10]), in the case of homogeneous items, the bids in both auctions are the same (i.e., the lines coincide) and are given by an increasing step function. In fact, we observe that for any $\gamma > \alpha + \beta$, the equilibrium follows this structure.

In the case of complementary heterogeneous items, the strategy follows the same form: the bid in each auction is an increasing step function, which is also consistent with the results for 2 bid levels (see, e.g., Fig. 9(a)). This can be seen in Fig. 13(b), which shows equilibrium strategies for heterogeneous items with a small degree of complementarity. We tried many other parameter settings (i.e., changing the number of bidders, bid levels, and complementarity structures with $\gamma > \alpha + \beta$), and equilibria for all of them followed this structure. Moreover, we noticed that, for high degrees of complementarity, the step functions coincide as in the homogeneous case.

Our results for substitutable items (i.e., $\gamma < \alpha + \beta$) are not as conclusive for two reasons. First, even though convergence to ϵ -equilibrium with a low error was always observed, multiple runs did not always produce the same equilibrium. Second, we could not discern any general patterns as we did in the complementary case. To illustrate this, we plot equilibrium strategies for two settings where multiple runs led to the same equilibrium. For weakly substitutable items, the equilibrium resembles the increasing step functions which characterised equilibria for complementary items. An example of this is in Fig. 14(a). However, when items are stronger substitutes, equilibrium is more difficult to describe. Fig. 14(b) shows equilibrium when items are close to being perfect substitutes.

Before proceeding to an analytical characterisation of equilibria, we note that all of the numerical results described in this section are for two simultaneous auctions. The algorithm is applicable to any number of auctions, however, we chose to study this case as it is already complex enough (and has not been solved before). Furthermore, computing the fair tie-

breaking rule for three or more auctions becomes too cumbersome (the case with two auctions is already complex enough as can be seen in [Appendix B](#)). We emphasise that settings with 3 and more simultaneous auctions can be studied with our FP algorithm given an approximation of the tie-breaking rule.

In the next section we provide analytical characterisation of equilibria. The equilibria that we derive there analytically confirm our numerical results from in [Section 6.2.3](#).

7. Analytical characterisation of equilibria for linear utilities

This section provides an analytical characterisation of equilibria for the case when agents' utilities are linear in type (as defined in [Section 5](#)). In more detail, we reduce the problem of finding equilibria to solving systems of polynomial equations. While this characterisation holds for all games with linear utilities, deriving equilibrium relies on the ability to solve the systems exactly. We demonstrate that it can be done for simultaneous auction games studied in [Section 6.2.3](#). Specifically, we analytically derive the equilibria and prove their uniqueness for each complementarity structure in simultaneous auctions for two homogeneous items with two bid levels and two bidders.

We compare these equilibria to the empirical results in [Section 6.2.3](#). This comparison is important since, even though the empirical results generally converge to a very small error, the small error only means that deviating from the approximate equilibrium strategy results in at most a small benefit. However, there are no guarantees that the approximate equilibrium is similar to the theoretical Nash equilibrium (in terms of the action distributions). Given this, the results in this section confirm that the ϵ -Nash equilibria discovered for this setting are the same as the unique exact equilibria derived below, and provide a validation for the fictitious play approach (at least in the simultaneous auctions domain). This holds for the entire range of complementarity structures for identical auctions studied in [Section 6.2.3](#).

Our analytical characterisation begins with the analysis of best response. We continue using the best response representation from [Section 5.1](#). Thus, a best response is specified by a set of m' actions $A' \subseteq A$ ordered according to the slope and the corresponding intervals (represented by an increasing vector $c \in \mathbb{R}^{m'-1}$) on which each action is played: action a'_j is the best-response action on the interval given by $[c_{j-1}, c_j]$. This representation is without loss of generality for pure-strategy best response: the actions of any best response are increasing in the slope of the utility lines²¹ and a single action is played for each type. Whereas in [Section 5.1](#) we presented an algorithmic procedure for finding a best response (see [Fig. 4](#)), we now provide an analytical characterisation. First, we give a few observations that follow immediately from the best-response structure.

Recall from equilibrium results for single-item auctions with continuous actions that the equilibrium action (i.e., bid) increases in type. However, in our model the set of actions is discrete and may have a more complex underlying structure (as would be the case, for example, in the simultaneous auctions model, where bids are not single dimensional), and there is no self-evident total order among them. To remedy this, we consider slopes of the utility lines. These slopes provide a total order over actions. As we noted in [Observation 1](#), the slopes of actions played in a best response strategy increase in type. Since an equilibrium is a best response, the same applies to equilibrium actions, providing an equilibrium monotonicity condition.

Observation 2. *The slope of the utility envelope of an equilibrium strategy increases in type.*

A direct consequence of this is another observation.

Observation 3. *If all slopes are distinct in equilibrium, an action cannot be played on more than one interval.*

For a given action distribution h , we say that $a_i \preceq a_j$ if the slope of a_i 's utility line is smaller than the slope of a_j 's utility line. With this notation, a best response is characterised by the following lemma.

Lemma 1. *Given a set of available actions $A = \{a_1 \preceq \dots \preceq a_m\}$, the pair $(A' = \{a'_1 \preceq \dots \preceq a'_{m'}\}, c)$ is a pure-strategy best response to the action distribution h if and only if the following equations are satisfied:*

$$A' \subseteq A, \quad (7)$$

$$0 < c_1 < \dots < c_{m'-1} < 1, \quad (8)$$

$$\hat{u}(c_j, a'_j, h) = \hat{u}(c_j, a'_{j+1}, h) \quad \forall 1 \leq j \leq m' - 1, \quad (9)$$

$$\hat{u}(0, a'_1, h) \geq \hat{u}(0, a_k, h) \quad \forall a_k < a'_1, \quad (10)$$

$$\hat{u}(c_j, a'_j, h) \geq \hat{u}(c_j, a_k, h) \quad \forall 1 \leq j \leq m' \quad a'_j < a_k < a'_{j+1} \quad (11)$$

²¹ This follows from the fact that the best-response function (see [Eq. \(3\)](#)) is convex.

where $c_0 = 0$, $c_{m'} = 1$, and $a'_{m'+1}$ is a dummy action (i.e., it does not appear in A) and has a slope above a_m (this dummy action is used in Eq. (11)).

Proof. See Appendix C. \square

The analytical characterisation of a best-response provides a partial characterisation of equilibrium: each equilibrium strategy is a best response. To be an equilibrium, the best-response strategy must be a best-response to itself. We formalise this in the theorem below.

Theorem 4. A strategy s is a pure-strategy symmetric equilibrium of the game $\Gamma = (N, A, u(\cdot), \Theta, F(\cdot))$ with $\hat{u}(\theta, s(\theta), h_s)$ linear in θ if and only if:

$$s(\theta) = a'_j \mid \theta \in [c_{j-1}, c_j]$$

where $(A' = \{a'_1 \preceq \dots \preceq a'_{m'}\}, c)$ satisfies Eqs. (7)–(11) as well as:

$$h(a'_j) = F(c_j) - F(c_{j-1}) \quad \forall 1 \leq j \leq m', \quad (12)$$

$$h(a_j) = 0 \quad \forall a_j \notin A'. \quad (13)$$

Proof. A strategy is an equilibrium if and only if it is a best response (and thus can be represented by a pair (A', c)) to itself: i.e., to the action distribution it induces. Eqs. (7)–(11) ensure the strategy is a best response. The action distribution corresponding to a strategy (A', c) is easy to express analytically: the probability of playing an action $a'_j \in A'$ is the same as the probability that the type is from the interval $[c_{j-1}, c_j]$ (Eq. (12)) while the probabilities of all other actions are zero (Eq. (13)). \square

A direct way of searching for an equilibrium is for each possible subset of actions $A' \subseteq A$ to check whether there exist parameters c satisfying best-response (8)–(11) and action distribution (12)–(13) equations. Although in general the equations can be arbitrarily complex depending on the distribution of types and number of players, they are almost always numerically solvable (see, e.g., [58–61]). A complete analytical characterisation is tractable when the number of actions is small. In the next section, we provide such a characterisation for the simultaneous auctions setting with 2 bid levels studied in Section 6.2.3.

7.1. Two identical items, two bidders, two bids per auction, uniform distribution of types

In this section, we use the above characterisation to provide an analytical derivation of the equilibrium for the simultaneous auctions setting studied numerically in Section 6.2.3. Specifically, we restrict our attention to 2 auctions each selling an identical item and 2 bidders with types uniformly distribution between 0 and 1. Furthermore, there are 2 bid levels per auction: 0 and 1. The set of possible joint bids is therefore $A = \{(0, 0) (0, 1) (1, 0) (1, 1)\}$. As before, we set $\alpha = \beta = 1$. The only remaining complementarity parameter is γ , which determines how much more or less an agent values having both items. In the following, we analytically derive equilibria as a function of γ .

We start by making several observations. First, due to uniform distributions, the bid distribution in Eq. (12) induced by a strategy (A', c) becomes:

$$h(a'_j; c) = F(c_j) - F(c_{j-1}) = c_j - c_{j-1}.$$

Second, we note that the actions (0, 1) and (1, 0) must be played with equal probability in equilibrium. To see this, suppose that (1, 0) is played more often than (0, 1). Then, the probability that the second auction has the price of 0 is higher. However, since the agent is indifferent between winning either item the best response is to play (0, 1) more often. Therefore, in equilibrium, the probabilities of playing (1, 0) and (0, 1) are the same, and these actions have an identical utility line in the best response. As a result, the best response interval on which either of the bids is played is continuous. That is, if these bids are played in an equilibrium on the interval $[c_1, c_2]$, then there is a continuum of equivalent equilibria where:

$$s(\theta_i) = (0, 1) \quad \text{or} \quad (1, 0) \mid f(1, 0) = f(0, 1) = \frac{c_2 - c_1}{2} \quad \text{if } \theta_i \in [c_1, c_2]. \quad (14)$$

This explains the switching behaviour we observed in fictitious play (see Section 6.2.3): at the equilibrium point any order of bids (1, 0) and (0, 1) is acceptable. However, any small change away from $h((1, 0)) = h((0, 1))$, leads to a unique preferred order. For notational convenience, since the probabilities of playing (1, 0) and (0, 1) are the same in equilibrium, in the following we merge this action into a single action and refer to the merged action as (1, 0). Specifically, saying that the action (1, 0) is played on the interval $[c_{j-1}, c_j]$ means that the actions (1, 0) and (0, 1) are played with equal probabilities on this interval.

Given the above observation, we can identify a unique order of slopes for each of the three actions. Note that, regardless of h , the action $(1, 0)$ wins in all the cases when the action of $(0, 0)$ wins. Similarly, the action $(1, 1)$ wins in all the cases when the action $(1, 0)$ wins. Thus, $(0, 0)$ has the lowest slope, $(1, 0)$ is next, and $(1, 1)$ has the highest slope. In particular, notice that the order does not depend on γ . Now, following Observation 2, in equilibrium, the slope increases in type, which means that $(1, 0)$ (or $(0, 1)$) is played by higher types than $(0, 0)$, and $(1, 1)$ is played by higher types than $(1, 0)$. Note that this is consistent with the equilibrium strategy described in Table 2.

The next step is to see which actions are played in equilibrium, and with what probabilities. It is easy to see that for any action distribution, $(0, 0)$ is the best response for types that are low enough and, thus, is played with a positive probability in any equilibrium. However, it is never the case that $(0, 0)$ is the only action in the support. These observations imply that the possible sets of equilibrium actions are $\{(0, 0) (1, 0)\}$, $\{(0, 0) (1, 1)\}$, and the set of all actions $\{(0, 0) (1, 0) (1, 1)\}$. In fact, as we will show, each of these sets corresponds to an equilibrium for some range of complementarity structures.

As an example, consider the set $A' = \{(0, 0) (1, 0)\}$. In the notation of Lemma 1, $m' = |A'| = 2$ and, to establish the probability of each action being played we are looking for the intersection point $0 \leq c_1 \leq 1$ satisfying:

$$\begin{aligned} u(c_1, (0, 0), h(\cdot; c)) &= u(c_1, (1, 0), h(\cdot; c)), \\ u(1, (1, 0), h(\cdot; c)) &\geq u(1, (1, 1), h(\cdot; c)). \end{aligned}$$

A solution exists only for $0 < \gamma \leq 2(2 - \sqrt{2})$ and is unique:

$$c_1 = \frac{-4 - \gamma + \sqrt{16\gamma + \gamma^2}}{-4 + 2\gamma}.$$

Carrying out a similar analysis, we derive equilibria for the other 2 action sets. Details of these derivations can be found in Appendix D. These derivations show that there is a *unique* equilibrium (except for variations between actions $(0, 1)$ and $(1, 0)$) for each value of γ . More formally:

Theorem 5. *The simultaneous auctions game defined by 2 bidders, actions*

$$A = \{(0, 0) (0, 1) (1, 0) (1, 1)\}$$

uniform distribution of types in $[0, 1]$, and complementarity structure $\alpha = \beta = 1$ and $\gamma > 0$ has a unique²² equilibrium defined below for every value of γ .

For $0 < \gamma \leq 2(2 - \sqrt{2})$ the equilibrium is $A' = \{(0, 0) (1, 0)\}$ and $c = (c_1)$ where:

$$c_1 = \frac{-\gamma - 4 \pm \sqrt{\gamma^2 + 16\gamma}}{2(\gamma - 2)}.$$

For $2(2 - \sqrt{2}) < \gamma < 2$ the equilibrium is $A' = \{(0, 0) (1, 0) (1, 1)\}$ and $c = (c_1, c_2)$ where:

$$\begin{aligned} c_1 &= \frac{2(2 - 2\gamma + \sqrt{-\gamma^2 + \gamma^3})}{4 - 4\gamma + \gamma^2}, \\ c_2 &= \frac{-6\gamma^2 + 4\sqrt{(-1 + \gamma)\gamma^2} + 2\gamma(2 + \sqrt{(-1 + \gamma)\gamma^2})}{(-2 + \gamma)^2(-\gamma + \sqrt{(-1 + \gamma)\gamma^2})}. \end{aligned}$$

For $\gamma = 2$ the equilibrium is $A' = \{(0, 0) (1, 1)\}$ and $c = (.5)$.

For $2 < \gamma$ the equilibrium is $A' = \{(0, 0) (1, 1)\}$ and $c = (c_1)$ where:

$$c_1 = \frac{-6 - \gamma + \sqrt{-28 + 44\gamma + \gamma^2}}{4(-2 + \gamma)}.$$

Fig. 15 plots the action distributions defined in the theorem above. Notice that the graph is virtually identical (up to very fine precision) to Fig. 10 obtained via numerical simulations. As we conjectured in Section 6.2.3, the equilibrium probabilities of each action are continuous in γ . The structure identified in Table 3 is also confirmed. Using the analytical characterisation, we determined that the smallest value of γ for which the bid $(1, 1)$ is played in equilibrium is $2(2 - \sqrt{2})$. This provides the exact value of $\hat{\gamma}_1$, which we roughly estimated to be around 1.2 by looking at Fig. 10.

Furthermore, the value of $\gamma = 2$ corresponds to independent auctions and equilibrium can be found for each auction separately. The single auction equilibrium with possible bids 0 and 1 can be easily derived: the unique equilibrium is to bid 0 for the values below 0.5 and to bid 1 for the values above 0.5. Combining individual equilibrium strategies, we get the

²² We are treating all equilibria given by Eq. (14) as one.




Fig. 15. Analytical results: action distributions for auctions with 2 bidders selling homogeneous items.

equilibrium strategy: bid (0, 0) for the values below 0.5 and (1, 1) for the values above. Each of the two equilibrium bids has the probability of 0.5 as can be seen in Fig. 15 for $\gamma = 2$.

Our analytical results show that, even though the approximate equilibrium strategy is not guaranteed to be similar to the actual Nash equilibrium strategy, in practice, we find that the empirical results are very close to the exact ones.

8. Conclusions

In this work we generalise FP to games of incomplete information with discrete actions and continuous types. We prove that, if FP beliefs converge, a *pure-strategy* Bayesian-Nash equilibrium can be constructed from the beliefs' limit point. Our algorithm recovers this equilibrium in case of (asymptotic) convergence. Furthermore, a pure ϵ -equilibria for any $\epsilon > 0$ can be obtained after a finite number of iterations.

The key distinguishing feature of our FP approach is that it works directly with continuous types and remains scalable in the number of agents and actions. This is in contrast to other currently available solvers (e.g., those listed in Section 2) that are typically only able to find equilibria in settings with discrete type spaces of small size or two players. Although, recent advances (such as graphical game representations and hybrid algorithms) allow discrete solvers to scale to larger type spaces, they nevertheless fail to accommodate cardinally larger continuous types. On the other hand, our algorithm is applicable to a large class of games with continuous type spaces and each iteration of FP can be computed efficiently. Furthermore, our FP algorithm can be applied to a wide range of auction settings, providing equilibrium calculations that otherwise would require specialised analyses and solution algorithms.

To illustrate the efficacy of our algorithm, we perform a set of numerical experiments, where FP was applied to a range of simultaneous auctions settings, where players have various combinatorial preferences for the items. The experiments show that FP converges to a very small ϵ in the settings we investigate, providing an empirical characterisation of equilibria in a complex domain for which no general theoretical results exist. We then analyse these equilibria in detail. The results show that, for weakly complementary items, as we vary the complementarity structure, the changes in equilibrium bids are continuous (there are no jumps). Furthermore, we observe that the bids are monotonic within each range of the complementarity parameter, i.e. where the support of equilibrium bids does not change. These characteristics continue to hold as we increase the number of bid levels and the number of bidders, although the position of the regions shift.

While the numerical results show convergence to ϵ -Nash equilibria with very small ϵ (in the order of less than 1% of the utility), there is no guarantee that this equilibrium is close to the true pure Nash equilibrium (which is known to exist for our setting). Therefore, to further support our results, in addition to the algorithm, we developed a full analytical characterisation for small settings. This shows that the equilibrium results are in fact unique, and correspond to those found by the FP algorithm.

We observe numerical convergence of FP to ϵ -equilibrium in all of our experiments, and to the same strategy in the experiments with complementary items. However, at the moment, we are not able to prove convergence analytically. The problem also proves elusive in games of complete information where results are available only for restricted settings (see, e.g., [7,15]). In fact, there are counter examples where FP is known not to converge. Identifying restricted settings of incomplete information where the generalised version of fictitious play provably converges remains open for future work.

Also open for future research are applications of the technique presented here to other domains (e.g., multi-unit or combinatorial auctions), both as a means of testing the convergence properties of FP, and as a means of obtaining numerical solutions to initiate a study of equilibrium properties in these domains. In fact, the technique outlined in this paper has been recently applied to compute equilibrium trading strategies in simultaneous double auctions in [54]. The authors show that in such settings, the FP algorithm consistently converges, allowing equilibrium trading strategies to be identified. There, the authors go even beyond a simple double auction to multiple simultaneous double auctions, where both buyers and sellers need to choose a double auction where they place bids and asks respectively. This setting is complex due to the presence

of both positive and negative networks effects; buyers are attracted to double auctions with many sellers, but would like to avoid competing buyers, and conversely for sellers.

Appendix A. Extensions

In this appendix we review some of the assumptions we used in the main body of the paper regarding the setting, utility functions and equilibria. Specifically, we consider, in turn, the limitations of a symmetric setting, single-dimensional type and linear dependency of the utility in type. In more detail, we show how the algorithm can be used to handle asymmetric settings and give a sketch of extensions to our algorithms to resolve the remaining limitations.

A.1. Asymmetric fictitious play

It is straightforward to extend the algorithms presented in this paper to asymmetric settings where each player has a potentially different action space, A , utility function, $u(\cdot)$, type space Θ , and distribution over types $F(\cdot)$, resulting in asymmetric equilibria. Formally, an *asymmetric* Bayesian game is defined by $\Gamma_{\text{asym}} = \langle N, \{A_i, u_i(\cdot), \Theta_i, F_i(\cdot)\}_{i \in N} \rangle$.²³ In equilibrium, each player $i \in N$ can have a different strategy $s_i(\cdot)$, resulting in action distribution $h_{s,i}(\cdot)$. Moreover, the expected utility function of a player i of type $\theta \in \Theta_i$ when playing action a_i in an asymmetric setting given the action distributions of *other* players $j \in N_{-i} = N \setminus i$ is defined as:

$$\hat{u}_i(\theta, a_i, \{h_{s,j}\}_{j \in N_{-i}}) = \mathbb{E}_{\{Y_j \sim h_{s,j}\}_{j \in N_{-i}}} [u_i(\theta, a_i, \{Y_j\}_{j \in N_{-i}})].$$

Similarly, the expected utility from playing a strategy $s'(\cdot)$ when players $j \neq i$ play strategies $s_j(\cdot)$ is $\tilde{u}_i(s', \{h_{s,j}\}_{j \in N_{-i}}) = \mathbb{E}_{\theta \sim X_i} [\hat{u}_i(\theta, s_i(\theta), \{h_{s,j}\}_{j \in N_{-i}})]$.

Given this we can define an *asymmetric* equilibrium as follows:

Definition 4. A strategy profile $s_i : \Theta_i \rightarrow A_i$, $i \in N$ is an asymmetric pure-strategy equilibrium of a game Γ_{asym} if:

$$\tilde{u}_i(s_i, \{h_{s,j}\}_{j \in N_{-i}}) \geq \tilde{u}_i(s', \{h_{s,j}\}_{j \in N_{-i}}) \quad \forall s' \in S_i, i \in N.$$

The remaining definitions can be modified analogously to the asymmetric setting.

We now turn to the FP algorithm in Fig. 1. To handle asymmetric settings, the algorithm now needs to compute the best response and maintain a separate set of beliefs for each player $i \in N$, which need to be updated separately (note that, if a subset of the players are symmetric, these can be grouped together into a single representative player). There are two approaches in which the beliefs can be updated: simultaneously or sequentially. In the former case, the best response for each player is calculated based on the beliefs from the previous iteration (at time t). In the latter case, the FP beliefs of each player are updated sequentially and these updated beliefs are used by the next player to calculate his best response. Although simultaneous updating is most commonly used in standard FP, Berger has shown that sequential updating actually has better convergence properties [62].

The amended FP algorithm containing both alternative updating rules is given in Fig. 16. Note that the `BestResponse` procedure requires a minor modification to the input to include the index of the player whose strategy we are computing, and the action distribution for each player $j \neq i$. However, no other modifications are needed to this algorithm. Furthermore, note that the `BeliefsToStrategy` procedure needs to be executed for each player. Finally, the convergence criterion needs to be modified since the best response can produce a different ϵ for each player.

A.2. Multi-dimensional types and non-linear utility

First, taking a closer look at the `FictitiousPlay` algorithm, depicted in Fig. 1, it is easy to see that it does not directly depend on the type space dimensionality. Rather, it was in the best response calculation and the `BeliefsToStrategy` procedure, where we made explicit use of our assumptions. Therefore, it is for these procedures that we need to relax our assumptions of type-space dimensionality and utility linearity.

Second, both the calculation of the best response and the `BeliefsToStrategy` procedure are based on a particular division of the type space: specifically, type space breakdown to action-equivalent subsets with respect to the best response upper envelope. Formally, for a belief h , we define for every action $a \in A$ a set $\tilde{N}_a(h) = \{\theta \in \Theta \mid \hat{u}(\theta, a, h) \geq \hat{u}(\theta, a', h) \quad \forall a' \in A\}$. Given a particular lexicographic ordering \preceq of actions, we can refine these sets into a collection of disjoint sets $\{I_a\}_{a \in \tilde{A} \subseteq A}$ (e.g., by setting $I_a = \tilde{N}_a \setminus \bigcup_{a' \preceq a} \tilde{N}_{a'}$, and purging empty I_a 's). Notice, that the collection $\{I_a\}_{a \in \tilde{A}}$ is a cover of the type space, so that $\bigcup_{a \in \tilde{A}} I_a = \Theta$. In fact, a specific collection would fully characterise the best response to a belief h . In particular, in

²³ Importantly, even though each agent has a different type distribution, $F_i(\cdot)$, we still require that these distributions are common knowledge. That is, we do not consider settings where some players have asymmetric beliefs about another player, and extending the FP algorithm to such settings is non-trivial.


```

Algorithm AsymmetricFictitiousPlay
Input: game  $\Gamma_{asym} = \langle N, \{A_i, u_i(\cdot), \Theta_i, F_i(\cdot)\}_{i \in N} \rangle$ ,
        initial beliefs  $h_i^0, i \in N$ , update rule  $\kappa$ 
Output: if converges, equilibrium strategy

1: set iteration count  $t = 0$ 
2: repeat
3:   for  $i \in N$ 
4a:    using simultaneous updating:
        strategy  $s = \text{BestResponse}(\Gamma_{asym}, i, \{h_j^t\}_{j \in N_{-i}})$ 
4b:    using sequential updating:
        strategy  $s = \text{BestResponse}(\Gamma_{asym}, i, \{h_j^{t+1}\}_{j < i}, \{h_j^t\}_{j > i})$ 
5:    compute the corresponding action distribution:
         $\forall a_i \in A_i : h_s(a_i) = \int_{s^{-1}(a_i)} f_i(x) dx$ 
6:    update beliefs of player  $i$ :
         $\forall a_i \in A_i : h_i^{t+1}(a_i) = \kappa(t)h_i^t(a_i) + (1 - \kappa(t))h_s(a_i)$ 
7:   end for
8:   set  $t = t + 1$ 
9: until converged
10: return  $\{\text{BeliefsToStrategy}(h_i^{t+1})\}_{i \in N}$ 

```

Fig. 16. Fictitious play algorithm for asymmetric games of incomplete information.

the case of single-dimensional utilities, this led to an interval structure. Notice, that `BeliefsToStrategy` simply utilises the collection $\{I_a\}_{a \in \tilde{A}}$ to define a policy $s(\theta) = \arg \max_{a \in \tilde{A}} \mathbb{1}_{I_a}(\theta)$.

Furthermore, formally the construction of the collection $\{I_a\}$ needs no assumption on the dimensionality, nor linearity of the utility function. Rather, these properties effect *only* the efficiency of that collection's representation. For instance, in the case of single-dimensional type space, the ordering of actions (by slope of their utility) created a linear fully ordered structure – the interval structure. On the other hand, for a two-dimensional type space, such an ordering is infeasible. However, alternative representations of such collections are possible. In fact, the field of computational geometry provides an extensive arsenal of such representations and algorithms ranging from envelopes of piecewise linear functions (see, e.g., [63]) to complex analytical curves (see, e.g., [64]).

Finally, notice that only the proof of Theorem 3 has made any use of the interval structure. Specifically, it relied on the fact that the *geometry* of the interval structure is similar for similar beliefs h^t and h^* . This statement, however, can be reproduced for any representation of the collection $\{I_a\}$, be that a set of intervals or a Delaunay triangulation, as long as it is consistent with some partial transitive ordering of actions for all beliefs h . Hence, by augmenting the representation of $\{I_a\}$, both the `BeliefsToStrategy` procedure and Theorem 3 can be adapted to hold for any dimensionality of the type space (or a non-linear utility).

Appendix B. Expected utility with fair tie-breaking

In this section we provide a computationally efficient procedure to compute the expected utility of an agent when there are $m = 2$ simultaneous auctions, when using a fair tie breaking rule. This tie breaking rule means that, if k players place the same bid in a particular auction, the probability of winning that auction is given by $1/k$. For convenience, we assume that both auctions have the same bid levels. Let B denote the set of discrete bids in a particular auction. Then the set of actions $A = B \times B$ available to each bidder is given by bid pairs $\mathbf{b} = (b_1, b_2) \in A$, where b_1 and b_2 are the bids in auctions 1 and 2 respectively.

Recall from Section 6.1 that the expected utility from playing an action $\mathbf{b} \in A$ is given by:

$$\hat{u}(\theta, \mathbf{b}, h) = \theta \sum_{\eta \subseteq K} \phi(\eta) q(\mathbf{b}, \eta, h) - \text{cost}(\mathbf{b}, h). \quad (\text{B.1})$$

In the following, we consider the left term first (which computes the expected value), followed by the right term (which computes the expected cost). Let $\Pr(W_i)$, $\Pr(W_i \cap W_j)$, and $\Pr(W_i \cap \bar{W}_j)$ denote the probability of winning auction i , the probability of winning both auctions i and j , and the probability of winning auction i but not j . Then for $K = \{1, 2\}$, we can rewrite the first term (ignoring the type) in Eq. (B.1) as follows:

$$\begin{aligned} \sum_{\eta \subseteq \{1, 2\}} \phi(\eta) q(\mathbf{b}, \eta, h) &= \alpha q(\mathbf{b}, \{1\}, h) + \beta q(\mathbf{b}, \{2\}, h) + \gamma q(\mathbf{b}, \{1, 2\}, h) \\ &= \alpha \Pr(W_1 \cap \bar{W}_2) + \beta \Pr(\bar{W}_1 \cap W_2) + \gamma \Pr(W_1 \cap W_2) \\ &= \alpha [\Pr(W_1) - \Pr(W_1 \cap W_2)] + \beta [\Pr(W_2) - \Pr(W_1 \cap W_2)] + \gamma \Pr(W_1 \cap W_2) \\ &= \alpha \Pr(W_1) + \beta \Pr(W_2) + [\gamma - \alpha - \beta] \Pr(W_1 \cap W_2), \end{aligned}$$

where $\alpha = \phi(\{1\})$, $\beta = \phi(\{2\})$, and $\gamma = \phi(\{1, 2\})$ as defined in Section 6.1. From the equation, we can see that, in order to calculate the expected value, it is sufficient to calculate $\Pr(W_1)$, $\Pr(W_2)$, and $\Pr(W_1 \cap W_2)$. In the following, we derive these probabilities based on the action distribution h .

Let $H(\mathbf{b}) = \sum_{\mathbf{b}' \in A: b'_1 < b_1 \& b'_2 < b_2} h(\mathbf{b}')$. It is convenient to think of the function H as a (multi-dimensional) cumulative distribution, where h is the corresponding probability mass function. In the following, we will also use the notation $H(b_1, b_2) = H(\mathbf{b})$ and $h(b_1, b_2) = h(\mathbf{b})$. Note that we define the inequalities in H to be strict. In addition, we use $\leq b_i$ to denote a non-strict relationship for a particular auction. For example, $H(b_1, \leq b_2) = \sum_{\mathbf{b}' \in A: b'_1 < b_1 \& b'_2 \leq b_2} h(\mathbf{b}')$.²⁴ Furthermore, let X_1 and X_2 denote random variables representing the bids placed by a bidder in auctions 1 and 2 respectively.²⁵ We can then use the functions H and h to define the following events:

1. $\Pr(X_1 < b_1 \cap X_2 < b_2) = H(b_1, b_2)$: win both auctions,
2. $\Pr(X_1 = b_1 \cap X_2 = b_2) = h(b_1, b_2)$: tie in both auctions,
3. $\Pr(X_1 = b_1 \cap X_2 < b_2) = H(\leq b_1, b_2) - H(b_1, b_2)$: tie in auction 1 and win auction 2,
4. $\Pr(X_1 < b_1 \cap X_2 = b_2) = H(b_1, \leq b_2) - H(b_1, b_2)$: win auction 1 and tie in auction 2,
5. $\Pr(X_1 > b_1 \cap X_2 > b_2) = [1 - H(\leq b_1, \leq b_2)]$: lose both auctions.

Note that the above events are mutually exclusive and always sum to one. If we define $H_i(b_i) = \sum_{\mathbf{b}' \in A: b'_i < b_i} h(\mathbf{b}')$ to be the cumulative bid distribution for a particular auction, we can similarly derive mutually exclusive events for a single auction:

1. $\Pr(X_i < b_i) = H_i(b_i)$: win auction i ,
2. $\Pr(X_i = b_i) = H_i(\leq b_i) - H_i(b_i)$: tie in auction i ,
3. $\Pr(X_i > b_i) = [1 - H_i(\leq b_i)]$: lose auction i .

The above provides the probabilities of certain events for a single other opponent. However, since there are $n - 1$ opponents, we need to calculate the distribution of the first-order statistic. When the bids are continuous, this is straightforward since (for example) $\Pr(X_i < b_i) = \Pr(X_i \leq b_i)$, and so $\Pr(W_i) = \Pr(X_i < b_i)^{n-1} = H_i(b_i)^{n-1}$. However, in the case of discrete bids, we also have to account for the tie breaking rule. For the single-auction case, $\Pr(W_i)$ becomes²⁶:

$$\Pr(W_i) = \sum_{x=0}^{n-1} \frac{1}{x+1} \binom{n-1}{x} \Pr(X_i = b_i)^x \Pr(X_i < b_i)^{n-1-x}. \quad (\text{B.2})$$

In the case of two auctions, this becomes much more complex since we need to enumerate over three possible events where ties occur: ties can occur in auction 1 only, in auction 2 only, or in both auctions. In the following, let x denote the number of bidders that correspond to first event, y to the second event, and z to the third event. Then, $\Pr(W_1 \cap W_2)$ becomes:

$$\begin{aligned} \Pr(W_1 \cap W_2) &= \sum_{x=0}^{n-1} \sum_{y=0}^{n-1-x} \sum_{z=0}^{n-1-x-y} \binom{n-1}{x} \binom{n-1-x}{y} \binom{n-1-x-y}{z} \frac{1}{x+z+1} \frac{1}{y+z+1} \\ &\quad \times \Pr(X_1 = b_1 \cap X_2 < b_2)^x \Pr(X_1 < b_1 \cap X_2 = b_2)^y \Pr(X_1 = b_1 \cap X_2 = b_2)^z \\ &\quad \times \Pr(X_1 < b_1 \cap X_2 < b_2)^{n-1-x-y-z}. \end{aligned} \quad (\text{B.3})$$

The above completes the computation of the expected value. We now show how to compute the final component of Eq. (B.1), the expected cost $\text{cost}(\mathbf{b}, h)$. We note that, unlike the probability of winning η out of K auctions, we can consider the expected costs for each auction separately thanks to linearity of expectation. Specifically, $\text{cost}(\mathbf{b}, h) = \text{cost}_1(b_1, h) + \text{cost}_2(b_2, h)$, where $\text{cost}_i(b_i, h)$ is the expected cost of auction i when bidding b_i in this auction. From [1] we know that for a single second-price auction, in the continuous case, the expected payment is equal to: $\text{cost}_i(b_i, h) = \int_0^{b_i} y g(y) dy$, where $g(y)$ is the density function of the first-order statistic of the bid distribution. However, with discrete bids, we need to consider two cases separately: when the second-highest bid in auction i is strictly less than b_i , then the bidder wins for sure. On the other hand, if the second-highest bid is equal to b_i , the probability of winning depends on the tie breaking rule. This results in the following equation:

²⁴ In practice, these functions can be implemented using a single look-up table, which can be computed in linear time and needs to be generated only once at the beginning of each FP iteration.

²⁵ Note that these random variables consider the bids of one of the bidders, and not *all* bidders. Since bidders are assumed to be symmetric it does not matter which particular one. Also, note that the variables are *interdependent* since the actions specify bid pairs. Therefore, we cannot assume that, e.g., $\Pr(X_1 < b_1 \cap X_2 < b_2) = \Pr(X_1 < b_1) \cdot \Pr(X_2 < b_2)$.

²⁶ Note that this is similar to Example 1 in Section 3, except that the distribution also takes into account the fact that the actions are joint bids over multiple auctions.

$$\text{cost}_i(b_i, h) = \sum_{x \in B | x < b_i} x[(H_i(\leq x))^{n-1} - (H_i(x))^{n-1}] + b_i[\Pr(W_i) - (H_i(x))^{n-1}], \quad (\text{B.4})$$

where $[(H_i(\leq x))^{n-1} - (H_i(x))^{n-1}]$ is the marginal bid distribution and corresponds to $g(y)$ in the continuous case. This completes the expected utility calculation.

As can be seen, due to the tie breaking rule, calculating the expected utility is computationally demanding. In particular, Eq. (B.3) considers all combination of ties which can simultaneously occur in 2 auctions (note that we cannot consider the auctions independently as the bid probabilities in the two auctions are correlated), and its computation is in the order of $O(n^3)$. This complexity increases rapidly for more than 2 auctions. Hence, in the next section we consider a way to approximate the expected utility concerning the tie breaking.

B.1. Approximate tie breaking

The above shows that, while computing the expected cost component of the expected utility is relatively easy since this can be done independently for each auction, the same is not true for calculating the expected value, which involves calculating the probability of winning every subset of auctions. In particular, the tie breaking rule increases the computation required and does not scale well with the number of bidders and the number of auctions. To address this problem, in this section we present an approximated tie breaking rule, which has been used in some of the experiments in this paper (specifically, in Section 6.2.2).

In more detail, the approximation is based on the observation that, in the case of 2 auctions, the exact probability of winning is always between $H(\leq b_1, \leq b_2)^{n-1}$ and $H(b_1, b_2)^{n-1}$. The first term over-estimates the probability of winning, whereas the second term underestimates it. Given this, the approximate probability of winning both auction is defined as:

$$\tilde{\Pr}(W_1 \cap W_2) = \lambda H(\leq b_1, \leq b_2)^{n-1} + (1 - \lambda) H(b_1, b_2)^{n-1}, \quad (\text{B.5})$$

where λ is a parameter which can be tuned. In the experiments, we used the value $\lambda = 1/3$ which performed well in general compared to the exact solution, although we did not excessively tune it. Note that this approach can also be applied to a single auction, and more than 2 auctions.

B.2. Two bidders

Here, we consider a special case of the equations with the exact tie breaking rule for $n = 2$. Instantiating Eq. (B.3) for $n = 2$, this results in the following 4 combinations of (x, y, z) : $(0, 0, 0)$; $(0, 0, 1)$; $(0, 1, 0)$; $(1, 0, 0)$. Expanding the equation we get

$$\begin{aligned} \Pr(W_1 \cap W_2) &= \underbrace{H(b_1, b_2)}_{0,0,0} + \underbrace{\frac{1}{4}h(b_1, b_2)}_{0,0,1} + \underbrace{\frac{1}{2}[H(b_1, b_2 + \epsilon) - H(b_1, b_2)]}_{0,1,0} + \underbrace{\frac{1}{2}[H(b_1 + \epsilon, b_2) - H(b_1, b_2)]}_{1,0,0} \\ &= \frac{1}{4}h(b_1, b_2) + \frac{1}{2}H(b_1, \leq b_2) + \frac{1}{2}H(\leq b_1, b_2). \end{aligned}$$

In the case of a single auction, we get:

$$\Pr(W_i) = H_i(b_i) + \frac{1}{2}[H_i(\leq b_i) - H_i(b_i)] = \frac{1}{2}[H_i(\leq b_i) + H_i(b_i)]. \quad (\text{B.6})$$

Instantiating Eq. (B.1), we get

$$\begin{aligned} u(\theta, \mathbf{b}, h) &= \theta(\alpha \Pr(W_1) + \beta \Pr(W_2) + (\gamma - \alpha - \beta) \Pr(W_1 \cap W_2)) - \text{cost}(\mathbf{b}, h) \\ &= \theta \left(\alpha \frac{1}{2}[H_1(\leq b_1) + H_1(b_1)] + \beta \frac{1}{2}[H_2(\leq b_2) + H_2(b_2)] \right. \\ &\quad \left. + (\gamma - \alpha - \beta) \left[\frac{1}{4}h(b_1, b_2) + \frac{1}{2}H(b_1, \leq b_2) + \frac{1}{2}H(\leq b_1, b_2) \right] \right) - \text{cost}(\mathbf{b}, h). \end{aligned}$$

The expected payment for auction i simplifies to

$$\text{cost}_i(b_i, h) = \left(\sum_{x \in B | x < b_i} x[H_i(\leq x) - H_i(x)] \right) + \frac{1}{2}b_i[H_i(\leq b_i) - H_i(b_i)].$$

B.3. Two bids

Further assume the only available bids levels are $B = \{0, 1\}$. This results in 4 possible bid pairs as actions. We denote the probability of encountering each bid pair, and the corresponding cumulative distribution as follows:

$$\begin{aligned} h(0, 0) &= x, & H(0, 0) &= 0, \\ h(1, 0) &= y, & H(1, 0) &= 0, & H(\leq 1, 1) &= x + y, \\ h(0, 1) &= z, & H(0, 1) &= 0, & H(1, \leq 1) &= x + z, \\ h(1, 1) &= 1 - x - y - z, & H(1, 1) &= x & H(\leq 1, \leq 1) &= 1. \end{aligned}$$

The corresponding distributions for single auctions are:

$$\begin{aligned} H_1(0) &= 0, & H_1(1) &= x + z, & H_1(\leq 1) &= 1, \\ H_2(0) &= 0, & H_2(1) &= x + y, & H_2(\leq 1) &= 1. \end{aligned}$$

This then results in the following expected utility for each action:

$$\begin{aligned} \hat{u}(\theta, (0, 0), h) &= \theta \left[\alpha \left(\frac{1}{4}x + \frac{1}{2}z \right) + \beta \left(\frac{1}{4}x + \frac{1}{2}y \right) + \gamma \frac{1}{4}x \right], \\ \hat{u}(\theta, (1, 0), h) &= \theta \left[\alpha \left(-\frac{1}{4}y + \frac{1}{2}z + \frac{1}{2} \right) + \beta \frac{1}{4}y + \gamma \left(\frac{1}{4}y + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - z), \\ \hat{u}(\theta, (0, 1), h) &= \theta \left[\alpha \frac{1}{4}z + \beta \left(\frac{1}{2}y - \frac{1}{4}z + \frac{1}{2} \right) + \gamma \left(\frac{1}{4}z + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - y), \\ \hat{u}(\theta, (1, 1), h) &= \theta \left[\alpha \left(-\frac{1}{4}x - \frac{1}{4}y + \frac{1}{4}z + \frac{1}{4} \right) + \beta \left(-\frac{1}{4}x + \frac{1}{4}y - \frac{1}{4}z + \frac{1}{4} \right) + \gamma \left(\frac{3}{4}x + \frac{1}{4}y + \frac{1}{4}z + \frac{1}{4} \right) \right] \\ &\quad - 1 + x + \frac{1}{2}(y + z). \end{aligned}$$

B.4. Identical auctions

The auctions are identical when $\alpha = \beta$. Since we do not restrict γ , we can without loss of generality set $\alpha = \beta = 1$.

$$\begin{aligned} \hat{u}(\theta, (0, 0), h) &= \theta \left[\frac{1}{2}x + \frac{1}{2}y + \frac{1}{2}z + \gamma \frac{1}{4}x \right], \\ \hat{u}(\theta, (1, 0), h) &= \theta \left[\frac{1}{2}z + \frac{1}{2} + \gamma \left(\frac{1}{4}y + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - z), \\ \hat{u}(\theta, (0, 1), h) &= \theta \left[\frac{1}{2}y + \frac{1}{2} + \gamma \left(\frac{1}{4}z + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - y), \\ \hat{u}(\theta, (1, 1), h) &= \theta \left[\left(-\frac{1}{2}x + \frac{1}{2} \right) + \gamma \left(\frac{3}{4}x + \frac{1}{4}y + \frac{1}{4}z + \frac{1}{4} \right) \right] - 1 + x + \frac{1}{2}(y + z). \end{aligned}$$

We are interested in utilities under symmetric equilibria. In all such equilibria the probabilities of playing bids (1, 0) and (0, 1) are the same: i.e., $y = z$.

$$\begin{aligned} \hat{u}(\theta, (0, 0), h) &= \theta \left[\frac{1}{2}x + y + \gamma \frac{1}{4}x \right], \\ \hat{u}(\theta, (1, 0), h) &= u(\theta, (0, 1)) = \theta \left[\frac{1}{2}y + \frac{1}{2} + \gamma \left(\frac{1}{4}y + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - y), \\ \hat{u}(\theta, (1, 1), h) &= \theta \left[\left(-\frac{1}{2}x + \frac{1}{2} \right) + \gamma \left(\frac{3}{4}x + \frac{1}{2}y + \frac{1}{4} \right) \right] - 1 + x + y. \end{aligned}$$

Appendix C. Proof of Lemma 1

Proof. We need to show that (A', c) specifies an upper envelope of the utility lines $\{\hat{u}(\theta, a_j, h)\}_{j=1}^m$. Eqs. (7) and (8) simply limit attention to a convenient best-response representation, which, as we noted is without loss of generality. A pair (A', c) satisfying these two equations defines a function:

$$g(\theta) = \hat{u}(\theta, a'_j, h) \quad \text{where } j \mid \theta \in [c_{j-1}, c_j].$$

First we show the “if” direction. Eq. (9) guarantees that the function is continuous: at each *intersection* point c_j (where the function $g(\theta)$ switches from one utility line to another), the values of adjacent utility lines are the same. This function is an upper envelope if no utility line lies above it. Since each equation is a line, it is sufficient to check that at each intersection point c_j the value of $g(\theta)$ is at least as high as the value of all other utility lines. In fact, it is enough to check each of the lines $A \setminus A'$ at exactly one c_j as we argue next. For a line $a_k \in A \setminus A'$ we can uniquely identify a'_j and a'_{j+1} such that $a'_j < a_k < a'_{j+1}$. Eq. (11)²⁷ checks that the utility from playing a_k at the type c_j is below the utility from playing a'_j (which, by Eq. (9), is the same as the utility from playing a'_{j+1}). Formally, we have:

$$\hat{u}(c_j, a_k, h) \leq \hat{u}(c_j, a'_j, h) = \hat{u}(c_j, a'_{j+1}, h).$$

We first consider:

$$\hat{u}(c_j, a_k, h) \leq \hat{u}(c_j, a'_{j+1}, h). \quad (\text{C.1})$$

Since two lines intersect at most once, and after the intersection, the line with the higher slope is on top, Eqs. (9) together with $a'_1 < \dots < a'_{m'}$ imply:

$$\hat{u}(c_{j+1}, a'_{j+1}, h) \leq \hat{u}(c_{j+1}, a'_{j+2}, h) \quad \forall \theta \geq c_{j+1},$$

$$\hat{u}(c_{j+2}, a'_{j+2}, h) \leq \hat{u}(c_{j+2}, a'_{j+3}, h) \quad \forall \theta \geq c_{j+2},$$

...

$$\hat{u}(c_{m'-1}, a'_{m'-1}, h) \leq \hat{u}(c_{m'-1}, a'_{m'}, h) \quad \forall \theta \geq c_{m'-1}.$$

Recalling that $a_k < a'_{j+1}$ and applying the argument above to Eq. (C.1), we get:

$$\hat{u}(\theta, a_k, h) \leq \hat{u}(\theta, a'_{j+1}, h) \quad \forall \theta \geq c_j. \quad (\text{C.2})$$

Noting that $\hat{u}(c_j, a'_j, h) \leq \hat{u}(c_{j+1}, a'_{j+1}, h)$ and combining Eq. (C.2) with the inequalities above, we get:

$$\hat{u}(c_i, a_k, h) \leq \hat{u}(c_i, a'_j, h) \quad \forall i \geq j+1. \quad (\text{C.3})$$

In other words, if a utility line with a smaller slope than a'_{j+1} is below a'_{j+1} at c_j , then it is also below $g(\theta)$ for all types above c_j . It remains to consider the case of $a'_j < a_k$:

$$\hat{u}(c_j, a_k, h) \leq \hat{u}(c_j, a'_j, h). \quad (\text{C.4})$$

Analogously to the argument above, before two lines intersect, the one with the lower slope is above the one with the higher slope. Hence, Eq. (9) and $a'_1 < \dots < a'_{m'}$ imply:

$$\hat{u}(c_{j-1}, a'_j, h) \leq \hat{u}(c_{j-1}, a'_{j-1}, h) \quad \forall \theta \leq c_{j-1},$$

$$\hat{u}(c_{j-2}, a'_{j-1}, h) \leq \hat{u}(c_{j-2}, a'_{j-2}, h) \quad \forall \theta \leq c_{j-2},$$

...

$$\hat{u}(c_1, a'_2, h) \leq \hat{u}(c_1, a'_1, h) \quad \forall \theta \leq c_1.$$

Recalling that $a_k < a'_{j+1}$, and applying the argument above to Eq. (C.4), we get:

$$\hat{u}(\theta, a_k, h) \leq \hat{u}(\theta, a'_j, h) \quad \forall \theta \leq c_j. \quad (\text{C.5})$$

Noting that $\hat{u}(c_{j-1}, a_k, h) \leq \hat{u}(c_{j-1}, a'_j, h)$ and combining Eq. (C.5) with the inequalities above, we get:

$$\hat{u}(c_i, a_k, h) \leq \hat{u}(c_i, a'_j, h) \quad \forall i \leq j. \quad (\text{C.6})$$

By Eqs. (C.3) and (C.6), the line for a_k is below $g(\theta)$ at all intersection points c_i and therefore at all types. This argument applies to all lines a_k , proving that (A', c) is a best response.

The “only if” part of the proof is trivial: violating Eq. (9) for any j results in a discontinuous function, while a best response must be continuous. Furthermore, if any of the inequalities in Eqs. (10) or (11) do not hold, then a better response is available, again contradicting the best response. \square

²⁷ Eq. (10) covers the special case of a'_1 , the first action in A' : the value of $g(\theta)$ at 0 is checked to be at least as high as the values of all utility lines with slopes below a'_1 .

Appendix D. Derivation of equilibrium strategies

Here we provide a derivation of the equilibria discussed in Section 6 and plotted in Fig. 15. The equilibrium bid distribution for a symmetric equilibrium (A', c) is

$$\begin{aligned} h(b_{\pi(j)}; c) &= F(c_j) - F(c_{j-1}) \quad \forall 1 \leq j \leq |A'|, \\ h(b_j; c) &= 0 \quad \forall b_j \notin A'. \end{aligned}$$

Plugging in the uniform distribution of types we get

$$h(b_{\pi(j)}; c) = c_j - c_{j-1} \quad \forall 1 \leq j \leq |A'|.$$

As before, we use $x, y, \gamma, 1 - x - 2y$ to denote $h(0, 0), h(1, 0), h(0, 1)$, and $h(1, 1)$ respectively.

We present a complete derivation for $A' = \{(0, 0), (1, 0)\}$. The piecewise linear strategy (A', c) consists of 2 intervals: bid $(0, 0)$ is played on the interval $\theta \in [0, c_1]$, bid $(1, 0)$ on the interval $\theta \in [c_1, 1]$. The strategy is given by a parameter $0 < c_1 < 1$. Using Lemma 1, (A', c) must satisfy

$$u(c_1, (0, 0), h(\cdot; c)) = u(c_1, (1, 0), h(\cdot; c)), \quad (D.1)$$

$$u(1, (1, 0), h(\cdot; c)) \geq u(1, (1, 1), h(\cdot; c)), \quad (D.2)$$

$$h(0, 0) = x = c_1, \quad h(1, 0) = h(0, 1) = y = \frac{1 - c_1}{2}, \quad h(1, 1) = 0. \quad (D.3)$$

Since we fixed A' , the only degrees of freedom are c and γ . We solve the first equation for c_1 as a function of γ and plug it into the second inequality to find the range of γ supporting the equilibrium (A', c) .

Expanding $u(c_1, (0, 0), h(\cdot; c)) = u(c_1, (1, 0), h(\cdot; c))$, we obtain

$$c_1 \left[\frac{1}{2}x + y + \gamma \frac{1}{4}x \right] = c_1 \left[\frac{1}{2}y + \frac{1}{2} + \gamma \left(\frac{1}{4}y + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - y).$$

Replacing x and y according to Eqs. (D.3) and simplifying we obtain

$$c_1^2(\gamma - 2) + c_1(\gamma + 4) - 2 = 0.$$

The roots of the quadratic equation are

$$c_1 = \frac{-\gamma - 4 \pm \sqrt{\gamma^2 + 16\gamma}}{2(\gamma - 2)}.$$

The roots are real only for $\gamma \leq -16$ and $\gamma \geq 0 \wedge \gamma \neq 2$. We are only interested in the roots that satisfy $0 < c_1 < 1$. The root

$$\frac{-\gamma - 4 - \sqrt{16\gamma + \gamma^2}}{2(\gamma - 2)}$$

is never between 0 and 1: negative for $\gamma = -16$ and decreasing as γ decreases; equal to 1 for $\gamma = 0$ and increasing for $\gamma \in [0, 2)$; approaching -1 from below for $\gamma > 2$. The other root

$$c_1 = \frac{-\gamma - 4 + \sqrt{16\gamma + \gamma^2}}{2(\gamma - 2)} \quad (D.4)$$

is approaching zero from below for $\gamma \leq -16$ as γ decreases. This root falls into the feasible range $0 < c_1 < 1$ for $\gamma > 0$ (undefined for $\gamma = 2$): equals 1 at $\gamma = 0$ and approaches zero from above as γ increases. Therefore, only Eq. (D.4) for $\gamma > 0$ may be supporting an equilibrium. To be an equilibrium, it must satisfy Eq. (D.2). We find the values of $\gamma > 0$ where the condition is satisfied

$$\left[\frac{1}{2}y + \frac{1}{2} + \gamma \left(\frac{1}{4}y + \frac{1}{2}x \right) \right] - \frac{1}{2}(1 - x - y) \geq \left[\left(-\frac{1}{2}x + \frac{1}{2} \right) + \gamma \left(\frac{3}{4}x + \frac{1}{2}y + \frac{1}{4} \right) \right] - 1 + x + y.$$

Replacing x and y according to Eqs. (D.3), then replacing c_1 with Eq. (D.4), we get after simplification

$$\frac{5\gamma^2 + \gamma\sqrt{16\gamma + \gamma^2} - 24\gamma + 16}{16(2 - \gamma)} \geq 0.$$

Solving the resulting inequality for γ , we obtain $0 < \gamma \leq 2(2 - \sqrt{2})$ giving us a full characterisation of equilibrium for $A' = \{(0, 0), (1, 0)\}$: the strategy (A', c) with c defined in Eq. (D.4) is an equilibrium for $0 < \gamma \leq 2(2 - \sqrt{2})$. There are no other equilibria with the support $A' = \{(0, 0), (1, 0)\}$.

Equilibria for $A' = A = \{(0, 0) (1, 0) (1, 1)\}$ and $A' = \{(0, 0) (1, 1)\}$ are derived similarly.²⁸ The systems of equations that need to be solved are

$$\begin{aligned} u(c_1, (0, 0), h(\cdot; c)) &= u(c_1, (1, 0), h(\cdot; c)), \\ u(c_2, (1, 0), h(\cdot; c)) &= u(c_2, (1, 1), h(\cdot; c)), \\ 0 < c_1 < c_2 < 1, \\ h(0, 0) &= c_1, \quad h(1, 0) = h(0, 1) = \frac{c_2 - c_1}{2}, \quad h(1, 1) = 1 - c_1 - \frac{c_2 - c_1}{2} \end{aligned}$$

and

$$\begin{aligned} u(c_1, (0, 0), h(\cdot; c)) &= u(c_1, (1, 1), h(\cdot; c)), \\ u(c_1, (0, 0), h(\cdot; c)) &\geq u(c_1, (1, 0), h(\cdot; c)), \\ 0 < c_1 < 1, \\ h(0, 0) &= c_1, \quad h(1, 0) = h(0, 1) = 0, \quad h(1, 1) = 1 - c_1 \end{aligned}$$

respectively.

The results above show that the equilibrium for each value of γ is unique (see Fig. 15).

We omit algebraic derivations and only present the solutions describing equilibria. The bids $A' = A = \{(0, 0) (1, 0) (1, 1)\}$ support the following unique equilibrium on $2(2 - \sqrt{2}) < \gamma < 2$

$$\begin{aligned} c_1 &= \frac{2(2 - 2\gamma + \sqrt{-\gamma^2 + \gamma^3})}{4 - 4\gamma + \gamma^2}, \\ c_2 &= \frac{-6\gamma^2 + 4\sqrt{(-1 + \gamma)\gamma^2} + 2\gamma(2 + \sqrt{(-1 + \gamma)\gamma^2})}{(-2 + \gamma)^2(-\gamma + \sqrt{(-1 + \gamma)\gamma^2})}. \end{aligned}$$

The bids $A' = \{(0, 0) (1, 1)\}$ support the following unique equilibrium on $2 < \gamma$

$$c_1 = \frac{-6 - \gamma + \sqrt{-28 + 44\gamma + \gamma^2}}{4(-2 + \gamma)}.$$

References

- [1] V. Krishna, Auction Theory, 2nd edition, Academic Press, 2009.
- [2] R.D. McKelvey, A.M. McLennan, T.L. Turocy, GAMBIT: Software tools for game theory, <http://www.gambit-project.org>, 2010.
- [3] R.D. McKelvey, A. McLennan, Computation of equilibria in finite games, in: Handbook of Computational Economics, vol. 1, North-Holland, Amsterdam, 1996, pp. 87–142.
- [4] B. von Stengel, Equilibrium computation for two-player games in strategic and extensive form, in: Algorithmic Game Theory, Cambridge University Press, 2007, pp. 53–78 (Ch. 3).
- [5] G.W. Brown, Iterative solutions of games by fictitious play, in: Activity Analysis of Production and Allocation, Wiley, 1951, pp. 374–376.
- [6] J. Robinson, An iterative method of solving a game, Annals of Mathematics 54 (1951) 296–301.
- [7] D.K. Levine, D. Fudenberg, Theory of Learning in Games, MIT Press, 1998.
- [8] A. Mas-Colell, M.D. Whinston, J.R. Green, Microeconomic Theory, Oxford University Press, New York, 1995.
- [9] C.F. Camerer, Behavioral Game Theory: Experiments in Strategic Interaction, Princeton University Press, 2003.
- [10] V. Krishna, R.W. Rosenthal, Simultaneous auctions with synergies, Games and Economic Behavior 17 (1996) 1–31.
- [11] E.H. Gerding, Z. Rabinovich, A. Bye, E. Elkind, N.R. Jennings, Approximating mixed Nash equilibria using smooth fictitious play in simultaneous auctions, in: The 7th International Conference on Autonomous Agents and Multiagent Systems, 2008, pp. 1577–1580.
- [12] M. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, 1994.
- [13] D. Monderer, L.S. Shapley, Potential games, Games and Economic Behavior 14 (1996) 124–143.
- [14] U. Berger, Fictitious play in $2 \times N$ games, Journal of Economic Theory 120 (2005) 139–154.
- [15] D. Monderer, L.S. Shapley, Fictitious play property for games with identical interests, Journal of Economic Theory 68 (1) (1996) 258–265.
- [16] S. Hart, A. Mas-Colell, A simple adaptive procedure leading to correlated equilibrium, Econometrica 68 (5) (2000) 1127–1150.
- [17] S. Hart, Adaptive heuristics, Econometrica 73 (5) (2005) 1401–1430.
- [18] A. Jafari, A. Greenwald, D. Gondek, G. Ercal, On no-regret learning, fictitious play, and Nash equilibrium, in: Proceedings of the 18th International Conference on Machine Learning, 2001, pp. 226–233.
- [19] S. Govindan, R. Wilson, Computing Nash equilibria by iterated polymatrix approximation, Journal of Economic Dynamics and Control 28 (2004) 1229–1241.
- [20] C. Papadimitriou, T. Roughgarden, Computing correlated equilibria in multi-player games, Journal of the ACM 55 (3) (2008) 1–29.
- [21] A.C. Chapman, A. Farinelli, E.M. de Cote, A. Rogers, N.R. Jennings, A distributed algorithm for optimising over pure strategy Nash equilibria, in: Proceedings of the 24th AAAI Conference on Artificial Intelligence, 2010, pp. 749–755.
- [22] S. Hoda, A. Gilpin, J. Pe  a, T. Sandholm, Smoothing techniques for computing Nash equilibria of sequential games, Mathematics of Operations Research 35 (2) (2010) 494–512.

²⁸ The algebraic derivations in this section were performed in Mathematica 7.0 [60].

- [23] A. Gilpin, S. Hoda, J. Pe  a, T. Sandholm, Gradient-based algorithms for finding Nash equilibria in extensive form games, in: Proceedings of the 3rd International Conference on Internet and Network Economics, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 57–69.
- [24] D. Koller, N. Megiddo, B. von Stengel, Efficient computation of equilibria for extensive two-person games, *Games and Economic Behavior* 14 (1996) 247–259.
- [25] S. Singh, V. Soni, M.P. Wellman, Computing approximate Bayes-Nash equilibria in tree-games of incomplete information, in: Proceedings of the 5th ACM Conference on Electronic Commerce, ACM, New York, NY, USA, 2004, pp. 81–90.
- [26] A.X. Jiang, K. Leyton-Brown, Bayesian action-graph games, in: Advances in Neural Information Processing Systems, vol. 23, Curran Associates, Inc., 2010, pp. 991–999.
- [27] D. Koller, B. Milch, Multi-agent influence diagrams for representing and solving games, *Games and Economic Behavior* 45 (1) (2003) 181–221.
- [28] S. Govindan, R. Wilson, A global Newton method to compute Nash equilibria, *Journal of Economic Theory* 110 (2003) 65–86.
- [29] G. van der Laan, A.J.J. Talman, L. van der Heyden, Simplicial variable dimension algorithms for solving the nonlinear complementarity problem on a product of unit simplices using a general labelling, *Mathematics of Operations Research* 12 (3) (1987) 377–397.
- [30] N.D. Stein, A. Ozdaglar, P.A. Parrilo, Separable and low-rank continuous games, *International Journal of Game Theory* 4 (2008) 457–474.
- [31] D.M. Reeves, M.P. Wellman, Computing best-response strategies in infinite games of incomplete information, in: Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence, AUA Press, Arlington, Virginia, United States, 2004, pp. 470–478.
- [32] M. Zinkevich, M. Littman, The AAAI computer poker competition, *Journal of the International Computer Games Association* 29 (3) (2006).
- [33] M. Zinkevich, M. Bowling, M. Johanson, C. Piccione, Regret minimization in games with incomplete information, in: Advances in Neural Information Processing Systems, vol. 20, Curran Associates, Inc., 2008, pp. 905–912.
- [34] S. Ganzfried, T. Sandholm, Computing equilibria in multiplayer stochastic games of imperfect information, in: Proceedings of the 21st International Joint Conference on Artificial Intelligence, 2009, pp. 140–146.
- [35] J. Hawkin, R. Holte, D. Szafron, Automated action abstraction of imperfect information extensive-form games, in: Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, 2011, pp. 681–687.
- [36] S. Ganzfried, T. Sandholm, Computing equilibria by incorporating qualitative models, in: Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, 2010, pp. 183–190.
- [37] Z. Rabinovich, E.H. Gerding, M. Polukarov, N. Jennings, Generalised fictitious play for a continuum of anonymous players, in: Proceedings of the 21st International Joint Conference on Artificial Intelligence, 2009, pp. 245–250.
- [38] R. Rosenthal, R. Wang, Simultaneous auctions with synergies and common values, *Games and Economic Behavior* 17 (1) (1996) 32–55.
- [39] B. Szentes, R. Rosenthal, Three-object two-bidder simultaneous auctions: Chopsticks and tetrahedra, *Games and Economic Behavior* 44 (2003) 114–133.
- [40] R.P. McAfee, Mechanism design by competing sellers, *Econometrica* 61 (6) (1993) 1281–1312.
- [41] M. Peters, S. Severinov, Competition among sellers who offer auctions instead of prices, *Journal of Economic Theory* 75 (1997) 141–179.
- [42] E.H. Gerding, A. Rogers, R.K. Dash, N.R. Jennings, Sellers competing for buyers in online markets: Reserve prices, shill bids, and auction fees, in: Proceedings of the 20th International Joint Conference on Artificial Intelligence, 2007, pp. 1287–1293.
- [43] E.H. Gerding, R. Dash, A. Bye, N.R. Jennings, Optimal strategies for bidding agents participating in simultaneous Vickrey auctions with perfect substitutes, *Journal of Artificial Intelligence Research* 32 (2008) 939–982.
- [44] E.H. Gerding, R.K. Dash, D.C.K. Yuen, N.R. Jennings, Bidding optimally in concurrent second-price auctions of perfectly substitutable goods, in: Sixth International Joint Conference on Autonomous Agents and Multiagent Systems, 2007, pp. 267–274.
- [45] M.S.-Y. Chwe, The discrete bid first auction, *Economics Letters* 31 (4) (1989) 303–306.
- [46] T. Mathews, A. Sengupta, Sealed bid second price auctions with discrete bidding, *Applied Economics Research Bulletin Special Issue (Auctions) I* (2008) 31–52.
- [47] W. Rudin, *Real and Complex Analysis*, McGraw-Hill, 1986.
- [48] W. Hildenbrand, *Core and Equilibria of a Large Economy*, Princeton University Press, 1974.
- [49] A. Mas-Colell, On a theorem of Schmeidler, *Journal of Mathematical Economics* 13 (1984) 201–206.
- [50] R. Radner, R.W. Rosenthal, Private information and pure-strategy equilibria, *Mathematics of Operations Research* 7 (3) (1982) 401–409.
- [51] D. Fudenberg, D.K. Levine, Consistency and cautious fictitious play, *Journal of Economic Dynamics and Control* 19 (1995) 1065–1089.
- [52] J. Hershberger, Finding the upper envelope of n line segments in $O(n \log n)$ time, *Information Processing Letters* 33 (4) (1989) 169–174.
- [53] S. Ganzfried, T. Sandholm, K. Waugh, Strategy purification and thresholding: Effective non-equilibrium approaches for playing large games, in: Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, 2012.
- [54] B. Shi, E. Gerding, P. Vytelingum, N.R. Jennings, An equilibrium analysis of competing double auction marketplaces using fictitious play, in: 19th European Conference on Artificial Intelligence (ECAI), 2010, pp. 575–580.
- [55] E. David, A. Rogers, J. Schiff, S. Kraus, M. Rothkopf, N.R. Jennings, Optimal design of English auctions with discrete bid levels, *ACM Transactions on Internet Technology* 7 (2) (2007), Article 12.
- [56] M. Abramowitz, I. Stegun (Eds.), *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover Publications, 1964.
- [57] Y. Shoham, K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2008.
- [58] M. Courtois, A. Klimov, J. Patarin, A. Shamir, Efficient algorithms for solving overdefined systems of multivariate polynomial equations, in: *Advances and Cryptology – EUROCRYPT*, in: Lecture Notes in Computer Science, vol. 1807, Springer, Berlin, 2000, pp. 392–407.
- [59] B. Sturmfels, *Solving Systems of Polynomial Equations*, CBMS Regional Conference Series in Mathematics, American Mathematical Society, 2002.
- [60] Wolfram Research Inc. *Manual*, 2010.
- [61] D. Eberly, *Solving Systems of Polynomial Equations*, GeometricTools, 2000.
- [62] U. Berger, Brown's original fictitious play, *Journal of Economic Theory* 135 (1) (2007) 572–578.
- [63] H. Edelsbrunner, L.J. Guibas, M. Sharir, The upper envelope of piecewise linear functions: Algorithms and applications, *Discrete and Computational Geometry* 4 (1) (1989) 311–336.
- [64] J.A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision*, Cambridge University Press, 1999.