Generalised Fictitious Play for a Continuum of Anonymous Players

Zinovi Rabinovich, Enrico Gerding, Maria Polukarov and **Nicholas R. Jennings**
School of Electronics and Computer Science,
University of Southampton,
Southampton SO17 1BJ
{zr, eg, mp3, nrj}@ecs.soton.ac.uk

Abstract

Recently, efficient approximation algorithms for finding Nash equilibria have been developed for the interesting class of *anonymous games*, where a player's utility does not depend on the identity of its opponents. In this paper, we tackle the problem of computing equilibria in such games with *continuous player types*, extending the framework to encompass settings with imperfect information. In particular, given the existence result for pure Bayes-Nash equilibiria in these games, we generalise the *fictitious play algorithm* by developing a novel procedure for finding a best response strategy, which is specifically designed to deal with continuous and, therefore, infinite type spaces. We then combine the best response computation with the general fictitious play structure to obtain an equilibrium. To illustrate the power of this approach, we apply our algorithm to the domain of simultaneous auctions with continuous private values and discrete bids, in which the algorithm shows quick convergence.

1 Introduction

As multiagent systems scale up, an individual's influence on the agents' interactions becomes ever smaller, and the resulting outcome depends on the aggregated actions taken by groups of agents (players). Now, the formal framework to model such situations is that of *games with a continuum of players*, which are also referred to as *large games*.[1] Typically, such games are *anonymous*, that is, the preferences of a player do not depend on the identities of its opponents. Rather, they only depend on action distributions over the population and the player's own action. This, in turn, is related to the assumption of perfect competition in large economies and multiagent systems with many participants, where any single individual has a negligible global effect. Relevant applications include the Internet, traffic routing and congestion settings, and auctions and markets.

[1]The informal intuition behind the terminology is that the number of players is so *large* that the set of players is viewed as a *continuous* mass, rather than discrete, separable individuals.

Against this background, in this paper we investigate *games with a continuum of anonymous players* (CAPs). Games with a continuum of players were first analysed in a pioneering paper by Schmeidler (1973) who proved the existence of pure strategy equilibria in these games. Later, Mas-Colell (1984), Rath *et al.* (1995) and Khan *et al.* (1997) found an alternative formulation for Schmeidler's model and simplified the existence proof. Interestingly, these results are also applicable to games with a finite set of players with continuous types, extending the framework to capture games with imperfect information, e.g. auctions with private evaluations.

However, besides these existence results, there are very few characterisation results for CAPs in the literature. Blonski (2001) provides necessary and sufficient conditions for an equilibrium distribution in CAPs with a finite action set. Daskalakis and Papadimitriou (2007) tackle the problem of computing Nash equilibria in anonymous games and develop efficient approximation algorithms for games with a finite set of players. However, the computation of equilibria in CAPs remains a relatively uncharted research direction, which is nevertheless important to the multiagent systems community because of the generality of CAPs and their relevance to the abovementioned applications.

To this end, our paper generalises the *fictitious play* (FP) algorithm (Brown, 1951), an iterative procedure whose convergence results in an equilibrium. In so doing, we develop the first FP-based algorithm which is applicable to CAPs. In particular, we present a novel procedure that efficiently computes a player's best response against a continuum of anonymous opponents, under some weak assumptions on the structure of the space of the players' utilities. We then combine the best response computation with the general FP structure to obtain an equilibrium.

Building on this, we apply our generalised FP algorithm to simultaneous auctions (Gerding *et al.*, 2007), where it quickly converges, producing a pure Bayes-Nash equilibrium. We choose this particular domain for its practical importance and theoretical interest, as this domain has been proven to be resistant to other computational techniques. In more detail, Reeves and Wellman (2004) provided a procedure to effectively compute best response strategies in two-player environments with utilities fully linear in player types and actions. While allowing both the private value and the action space to be continuous, the linearity assumption is extremely restrictive, and does not necessarily hold in many settings. In fact, as we show in this paper, linearity in actions is explicitly violated in our simultaneous auctions example. Furthermore, it is shown that the simple iterative best response procedure employed in (Reeves and Wellman, 2004) to compute the equilibrium, does not converge in our example domain.

For more complex domains, several approximation algorithms have been developed. For example, Jordan *et al.* (2008) formulated computation of a Nash equilibrium as a search problem and provided an estimation procedure for a pure equilibrium in games with large strategy spaces. Their solution, however, is not directly applicable to games with private information. This shortcoming has been addressed by the work of Vorobeychik and Wellman (2008) who extended search based computation to simulation based games. In their paper, simulated annealing was applied to compute approximated equilibria in games with private information. However, their algorithm, together with more amenable assumptions, required domain specific parameter selection and function design. In contrast, our FP based approach is generic and domain independent.

The rest of the paper is organised as follows. In Section 2 we formally define the class of games with a continuum of anonymous players. In Section 3 we present our generalised fictitious play algorithm for computing equilibria in CAPs. Section 4 is devoted to our experimental setting of simultaneous auctions, and the results are presented in Section 5. We conclude in Section 6.

2 Continuum of Anonymous Players

Here we formally define the class of CAPs, closely following the model of (Mas-Colell, 1984). Let A denote a compact metric space of actions available to each player, and let $\mathcal{M} = \mathcal{M}(A)$ denote the space of all probability measures on A endowed with the weak convergence topology. The utility of a player is defined by a continuous function $u : A \times \mathcal{M} \to \mathbb{R}$, mapping the player's action and the action distribution induced by choices of other players into a reward. The space of all continuous functions of the above form is denoted by $\mathcal{U}$ and equipped with the supremum norm.

Notice that a player type can be viewed as just an alternative (symbolic) name for the utility function, since the space of the former maps into the space of the latter. Following this intuition, a game with a finite number of players and a distribution over private types can be alternatively defined in terms of a distribution over the space of utility functions.

Definition 1 *A CAP is given by a probability distribution μ over the space $\mathcal{U}$ of continuous functions of the form $u : A \times \mathcal{M} \to \mathbb{R}$, where $\mathcal{M}$ is the space of distributions over a compact space of actions A.* **(Mas-Colell, 1984)**

Following this line of thought even further, it becomes convenient to formalise equilibrium in terms of distributions as well. Namely, an equilibrium τ of a game μ, termed a *Cournot-Nash Equilibrium* (CNE), is a distribution over the space of type-action pairs $\mathcal{U} \times A$, that satisfies the following:

1. $\tau_{\mathcal{U}} = \mu$
2. $\tau(\{(u, a) \| u(a, \tau_A) \geq u(A, \tau_A)\}) = 1$,

where $\tau_{\mathcal{U}}, \tau_A$ are the marginals of τ on $\mathcal{U}$ and A respectively. The intuition behind this definition is that, given a global distribution τ, selecting the best action in response to this distribution itself would not change it *en masse*.

The important result of (Mas-Colell, 1984) is the proof of existence of a *symmetric equilibrium*, where each player type is assigned exactly one action to follow.

Definition 2 *A CNE distribution τ for a CAP game μ is* ***symmetric*** *if there is a measurable function $h : \mathcal{U} \to A$ such that $\tau(\{(u, a) | a = h(u)\}) = 1$. i.e., players with the same characteristics play the same action.* **(Mas-Colell, 1984)**

The conditions for the existence of symmetric equilibria are that the game μ is non-atomic, giving zero probability to any specific player type to appear, and that the action space is discrete and finite. In the following sections we will adopt these conditions to simplify the material exposition, but we note that only the best response calculation requires them explicitly. Notice also that the existence of a function h makes this form of Cournot-Nash Equilibrium equivalent to a pure Bayes-Nash equilibrium of typed games. For the remainder of the paper we will use these two terms interchangably.

3 Fictitious Play for CAPs

In this section we outline the basics of FP algorithms (Brown, 1951; von Neumann and Brown, 1950) and introduce our generalised version for CAPs.

In more detail, the standard algorithm consists of computing and applying a best response to a frequency estimate (termed *FP belief*) of the opponent's actions. The underlying assumption of the FP beliefs is that an opponent samples its actions from some fixed distribution, i.e. opponents are assumed to play a fixed mixed strategy. Under this assumption, keeping score of the relative appearance frequencies for different actions provides a good estimate of the opponents' strategy, and justifies the application of the best response to the FP belief. The algorithm, however, dictates performing the belief updates for all agents of the game, with the intuition that the agents will continually adapt to each other, eventually arriving at an equilibrium.

The FP algorithm has two types of convergence. First, it may converge in terms of the *strategy*, i.e. after a number of iterations, the best response strategy of each agent may stabilise. In this case, the collection of the players' best response strategies constitutes a pure Nash Equilibrium. Unfortunately, it is quite easy to construct a game where this best response stabilisation will not occur, which brings us to the second type of convergence. A game is said to have the *fictitious play property* (FPP), if FP beliefs converge (see e.g. (Monderer and Shapley, 1996)). The set of converged FP beliefs then constitutes a mixed Nash equilibrium.

However, the standard notion of FP belief does not include types. This renders the existing FP algorithm inapplicable to typed games. In what follows, we introduce the necessary changes (including the relationship between best response and type) and generalise FP to CAPs.

3.1 The Generalised FP Algorithm

Before we formally define FP as it relates to CAPs, let us make some preliminary observations to clarify its structure.

First, the concept of beliefs needs to be generalised. In the standard FP, a belief maps the identity of the opponent to an empirical frequency of action appearance. But in CAPs the opponent identities are represented by their utility functions, furthermore there is a continuum of such utilities. It then follows that a belief has to map from the space of utility functions to the space of action frequencies: $f : \mathcal{U} \to \mathcal{M}$. Alternatively, the belief can be represented as a distribution measure τ over the space $\mathcal{U} \times A$ with the limitation that $\tau_{\mathcal{U}}$, the marginal on $\mathcal{U}$, will coincide with μ, the game itself. Now, if the beliefs converge, they converge to a CNE, in a similar fashion to the convergence of the standard FP beliefs to a Nash equilibrium. Although the resulting CNE is not necessarily symmetric, if the action space is finite, the equilibrium can be symmetrised (or purified) (see e.g. (Radner and Rosenthal, 1982)) giving rise to a pure Bayes-Nash equilibrium.

Second, we need to reconsider the update of beliefs. Instead of taking the distributed view of the standard FP computation, in which every player maintains an independent set of beliefs and performs the update independently, in CAPs we have to compute the best response to the current beliefs (distribution) for all types of players and update all beliefs in the system to incorporate the best response.

We are now ready to introduce the generalised FP algorithm for CAPs (see Figure 1). The algorithm begins by initialising the beliefs, τ^0, to an arbitrary choice of actions performed by the population. For example, players of all types may select actions uniformly. The algorithm then enters a loop that continually computes a best response and updates the beliefs. At iteration t, the best response is computed (line 2) with respect to the population wide distribution of actions expressed by the marginal distribution, τ_A^t, of the belief τ^t. The detailed procedure for computation of the best response function, $h : \mathcal{U} \to A$, is presented in Section 3.2. Once the response h is obtained, its inherent joint type-action distribution is calculated (line 3), and the beliefs of the next iteration τ^{t+1} absorb it accordingly (line 4), with the standard update rate of $\alpha(t) = \frac{t}{t+1}$. Finally, if the beliefs indicate convergence (lines 5-6), the resulting distribution τ is returned.

Require:
Set iteration count $t = 0$
Set $\tau^0 = \mu \otimes m$ for some $m \in \mathcal{M}$
1: **loop**
2: Compute best response function: $h : \mathcal{U} \to A$
3: Compute inherent distribution:
$$\tau_h(\mathbf{u}, a) = \mu(h^{-1}(a) \cap \mathbf{u})$$
4: Update beliefs:
$$\tau^{t+1} = \alpha(t) * \tau^t + (1 - \alpha(t)) * \tau_h$$
5: **if** (Convergence precision reached) **then**
6: **return** $\tau = \tau^{t+1}$
7: **end if**
8: Set $t \leftarrow t + 1$
9: **end loop**

Figure 1: Generalised FP algorithm for CAPs.

The generic algorithm (Figure 1) can be simplified further if the domain of the players' utility functions is taken into account. Specifically, any $u \in \mathcal{U}$ has $\mathcal{M}$ as its domain, i.e. the best response is not based on the entire belief, τ^t, but rather on its action space marginal, τ_A^t. This allows us to reduce the supported beliefs from the space of distributions over $\mathcal{U} \times A$ to the space $\mathcal{M}$ of distributions over actions (see Figure 2, especially notice the change in line 3).

Require:
Set iteration count $t = 0$
Set $\tau_A^0 = m$ for some $m \in \mathcal{M}$
1: **loop**
2: Compute best response function: $h : \mathcal{U} \to A$
3: Compute the *marginal* distribution:
$$\tau_{h,A}(a) = \mu(h^{-1}(a))$$
4: Update beliefs:
$$\tau_A^{t+1} = \alpha(t) * \tau_A^t + (1 - \alpha(t)) * \tau_{h,A}$$
5: **if** (Convergence precision reached) **then**
6: **return** $\tau_A = \tau_A^{t+1}$
7: **end if**
8: Set $t \leftarrow t + 1$
9: **end loop**

Figure 2: Compact version of generalised FP.

We note that the algorithm presented in Figure 1 directly produces an equilibrium (line 6 returns a *complete* distribution), while the compact version (Figure 2) may require an additional step (line 6 returns a *marginal* distribution). In more detail, if the algorithm converges in strategies, then the simplified update procedure may be used directly to compute the equilibrium: we simply need to compute the best response, h_A, to the distribution τ_A, and then CNE is well defined by $\tau_{CNE}(\mathbf{u}, a) = \mu(h_A^{-1}(a) \cap \mathbf{u})$. If, however, the algorithm converges in beliefs, then an additional procedure is necessary to lift τ_A to τ_{CNE}. This procedure is similar to the CNE purification of Radner and Rosenthal (1982) and is based on the inclusion-exclusion principle. Notice, however, that in both cases the final outcome is a symmetric CNE.

3.2 Best Response Computation Procedure

Let us now introduce the specific procedure for computing the best response function h. Although, in general, computing the best response may be complex, if the dependence of the player's utility on its type is analytically simple, an efficient procedure can be composed. For ease of exposition, we begin by considering linear utilities; we then proceed and comment on more general cases.

In this context, the utility is linear in type if there is a function, $u_{gen} : \mathbb{R} \times A \times \mathcal{M} \to \mathbb{R}$, such that for all $\lambda \in \mathbb{R}$ the utility function of the player with type λ is given by $u_{gen}(\lambda, \cdot, \cdot) \in \mathcal{U}$, and the utility $u_{gen}(\lambda, a, m)$ is linear in λ. Now, if the action space is discrete and finite, then the optimal utility, as a function of private value, is piecewise linear. This is since, when computing the best response to the distribution of actions τ_A, the latter is considered to be fixed, and the value of an action, u_a, becomes a linear function of the private value. More specifically, $u_a(\lambda) = u_{gen}(\lambda, a, \tau_A)$. The best utility that a player with a given private value λ can achieve, is then $u_h(\lambda) = \max_{a \in A} u_a(\lambda)$, i.e. an upper envelope of a finite set of linear functions, and thus is piecewise linear.

As a result, the best response function, $h(\lambda) =$

$\arg\max_{a \in A} u_a(\lambda)$, can be written as a step function. In other words, it is possible to create a set of distinct intervals I, that cover the type space, i.e. $\bigcup_{\alpha \in I} \alpha = \mathbb{R}$, and also:

- For any $\alpha \in I$, if $\lambda_1, \lambda_2 \in \alpha$ then $h(\lambda_1) = h(\lambda_2)$
- For any distinct $\alpha_1, \alpha_2 \in I$, if $\lambda_1 \in \alpha_1, \lambda_2 \in \alpha_2$ then $h(\lambda_1) \neq h(\lambda_2)$

The set of intervals corresponds to the set of linear segments of the optimal utility function u_h, and the value of the best response is the action that creates the corresponding segment. Notice that, although the number of intervals may change over the course of a FP run, the maximum number of intervals $|I|$ only depends on $|A|$ and the shape of utility functions, and thus remains boundeded by a constant. In particular, for linear utility functions $|I|$ is bounded by $|A|$.

This piece-wise linear representation allows efficient computation of the global action frequency induced by the best response. Given type distribution, μ, the probability that the agent will actually have type in $\alpha \in I$ is $\mu(\alpha)$. Then, if all players use their best response, the action frequency induced by the group of players is $\tau_{h,A}(a) = \mu(I_a)$, where $I_a = \bigcup_{\alpha \in I, h(\alpha)=a} \alpha$.

In a general setting with non-linear utilities, a similar procedure can be used. In fact, as long as we can efficiently compute the maximisation envelope of the set of $u_a(\lambda)$ functions and inverse them to obtain interval division, the specific dependency of $u_{gen}(\lambda, \cdot, \cdot)$ on the player's type λ is superfluous. For example, the procedure works equally well for quadratic u_{gen}, or it would also be applicable for multidimensional type spaces, as would appear, for instance, in multi-item auctions (where a player's type is represented by a vector of values, one for each item).

4 Simultaneous Auctions

In this section we describe the simultaneous auctions problem, and show how the algorithm discussed earlier can be applied to this model. We choose this setting because, to date, there is no known pure Nash equilibrium solution. Furthermore, related research has shown that simple iterative approaches do not converge in this domain (Gerding *et al.*, 2007). At the same time, simultaneous auctions appear in many practical settings such as online auctions, where typically many similar goods are being sold at the same time, and, furthermore, such auctions are an effective way to allocate resources between agents and achieve coordination in multi-agent systems.

4.1 Auctions, CAPs and CNE

Auctions are usually anonymous in the sense that their outcome does not depend on the identity of the bidders, but only on the bids themselves. That is, practically any auction is an anonymous game. In addition, they possess another feature that is of special interest to us – namely, the fact that the auctioned item has a valuation that is personal and private to each of the agents. If in a generic anonymous game a player may know the exact set of utility functions that drive its opponents, in an auction only an estimate about the private value, and thus the utility, of other players can be obtained. It is common to assume that private values of all the players are independently sampled from some continuous distribution, which is known to all auction participants.

Under this assumption, a transformation occurs to the way a player considers its opponents. They stop being individuals and become a sample of a utility function's population, each appearing with respect to the probability density of the private value. As the opponent values are hidden, the agent has to consider the entire range of possible value assignments. As a result, even though the factual number of opponents may be finite, the player's decision is computed in response to a (virtual) continuum of players formed by the range of possible private values.

Given this, auctions may be captured using the notation and terminology of games with a continuum of anonymous players described earlier. Specifically, the action space, A, in auction settings is the space of all possible bids a player can place. Since the setting is anonymous, a utility function of an agent will have the form $u : A \times \mathcal{M} \rightarrow \mathbb{R}$, where $\mathcal{M}$ is the space of distributions over bids placed by the player's opponents. In turn, the specific shape of the utility function is determined by the private value, and the distribution of private values will shape the game μ, a distribution over the space $\mathcal{U}$. This observation allows us to apply CNE existence theorems and our generalised FP algorithm to auctions in a generic way, almost independently of the specifics of the auction process.

4.2 The Specific Setting

We consider a market consisting of m auctions $A_1, A_2, \ldots, A_m$, selling a single item each, and n bidders competing in these auctions. The items are complete substitutes, that is the bidders are indifferent between them and derive no additional utility from winning more than one item. A bidder derives a value v from obtaining one or more items, and these valuations are i.i.d. drawn from a *continuous* distribution with cumulative function F and density f.

We assume that the bids are *discrete* and that the size of the bid space (i.e., the allowable bids) is finite. For simplicity, we assume that the bids are equally spaced, and, without loss of generality, that bids consist of integer values in the range $[0 : k]$. In the following, let $B = [0 : k]$ denote the bid space of a single auction, $B^m = [0 : k]^m$ the joint bid space over all auctions, and $\mathbf{b} = (b_1, b_2, \ldots, b_m) \in B^m$ a bid vector which specifies a bid for each auction. Furthermore, we assume that bidder valuations range between $(0, \mathbf{v}]$. Finally, throughout we assume that bidders are risk neutral.

We focus on second-price sealed bid *simultaneous auctions*, in which the bidders need to submit all their bids before the outcome of any auction is observed. Without loss of generality, we assume that all bidders place bids (possibly zero) in all auctions. In this setting, a bidding strategy is a function $S : (0, \mathbf{v}] \rightarrow B^m$ that maps a value into a vector of bids. Now, in order to calculate a bidder's expected utility given its bids and valuation, a bidder requires information about the actions of other players. However, although the actions are not known to the bidder, a bidder maintains beliefs about the bids of others. That is, a bidder maintains the probability that a certain bid occurs in an auction.

Notice that in our setting the bids placed in different auctions are correlated, and hence we are required to consider *joint* beliefs, i.e. joint probability distributions. We use the following notation. Let $X_1, X_2, \ldots, X_m$ denote discrete affiliated random variables representing the bids placed by a player, and let $P(\mathbf{b})$ denote their joint distribution. Similarly, we use $P_i(b_i)$ to denote the bid distribution for a single auction i, and $P_I(\mathbf{b})$ the joint distribution of bids in a subset $I \subseteq [A_1 : A_m]$ of auctions.

The player's utility function consists of two parts: the expected benefits, which is the valuation multiplied by the probability of winning *at least* one auction, minus the expected costs, the latter being the sum of the expected payments for each individual auction conditional on winning that auction:

$$U(v, \mathbf{b}) = v\ P_W(\cup_{i=1}^{m} A_i|\mathbf{b}) - \sum_{i=1}^{m} P_W(A_i|\mathbf{b}) C(A_i|b_i),$$

where $P_W(\cdot)$ stands for a probability of winning, and $C(\cdot)$ is the expected cost paid in case of winning. Note that this cost only depends on the bid in auction A_i, and not on the bids in other auctions:

$$C(A_i|b_i) = \frac{1}{P_i(b_i)^{n-1}} \sum_{x=1}^{b_i} (x-1)[P_i(x)^{n-1} - P_i(x-1)^{n-1}],$$

where $P_i(\cdot)^{n-1}$ stands for the highest order statistics, which, in fact, defines the probability of winning in auction A_i as a function of the bid placed in the auction.[2] Thus, $P_W(A_i|\mathbf{b}) = P_i(b_i)^{n-1}$, and $P_W(\cap_{i \in I} A_i|\mathbf{b}) = \times_{A_i \in I} P_i(b_i)^{n-1} = P_I(\mathbf{b})^{n-1}$ is the probability of winning all of the auctions in subset I. Finally,

$$P_W(\cup_{i=1}^{m} A_i|\mathbf{b}) = \sum_{j=1}^{m} (-1)^{j-1} \sum_{I \subset [1:m]\ s.t. |I|=j} P_I(\mathbf{b})^{n-1},$$

where $|I|$ is the cardinality of I. Clearly, the utility function is linear in the continuous valuation, v, and our generalised FP algorithm can be applied directly to this problem.

5 Empirical Evaluation

To demonstrate the effectiveness of the generalised FP algorithm we performed a set of experiments and applied our algorithm to the simultaneous auctions domain described above. The success of FP in these experiments is threefold. First, the algorithm converged in this non-trivial setting, which makes it a viable solution to a set of complex auction domains. Second, the algorithm showed quick convergence, which makes it an empirically efficient solution, in spite of its weak theoretical convergence properties. Third, this is the first time a pure Bayes-Nash equilibrium could be obtained for simultaneous auctions with continuous private values.

We now proceed and present our experimental setting and results in more detail. Since the auctions in our setting are identical, many equivalent equilibria may occur, only differing by the order of the auctions (clearly, in the case of two auctions, bidding high in one auction and low in the other results in the same payoff to a player as for doing the reverse). To eliminate such repetitions, we assume that different players may have different orderings for the auctions. Thus, for any individual bidder, it appears as though other players are placing their bids randomly between the auctions. Anonymising the auctions in this way, besides eliminating equivalent equilibria, has the advantage of reducing the action space, which, in turn, makes the calculation of the best response more efficient. Namely, without loss of generality, the space of actions can be replaced by the space of nondecreasing bid vectors, i.e. vectors in which a bid in auction A_i is greater than a bid in auction A_j only if $i > j$.

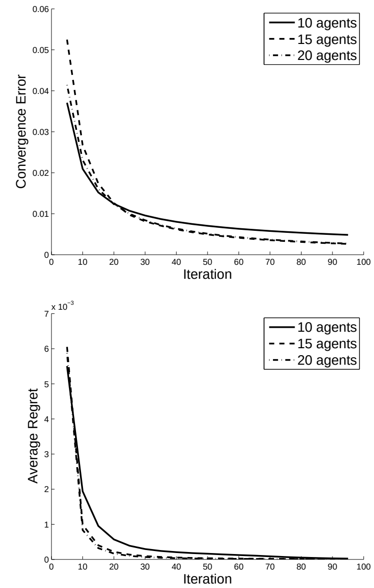

Figure 3: Convergence error (top) and average regret (bottom) for 4 auctions.

Furthermore, we assume that the players' private values for an auctioned item are uniformly distributed in the $[0, 1]$ interval, and that the bid space is discretised to form 10 distinctive bid levels. Even though the number of distinct bids seems to be small, the space of all possible joint bids is very large, namely $\Omega\left(\frac{|B|^m}{m!}\right)$, where m is the number of auctions and $|B|$ is number of bid levels.

To evaluate the performance of our algorithm, we simulate and run the simultaneous auctions domain with a varying number of auctions and bidders, and measure convergence of the algorithm using two indicators. First, we measure the convergence in beliefs, by calculating the *convergence error*, CE, which is determined by the infinity norm of the difference between two consecutive action distribution estimates τ_A^t and τ_A^{t+1}: $CE = \max_{a \in A} |\tau_A^{t+1}(a) - \tau_A^t(a)|$. If $CE < \frac{1}{t}$, the algorithm converges in beliefs.

Second, we compute the *average regret*, where regret is the difference between the utility obtained by a bidder if every-

[2] The tie breaking rule we employ in cases where two or more players place the same highest bid in an auction, are omitted from this version of the paper, due to space limitations.

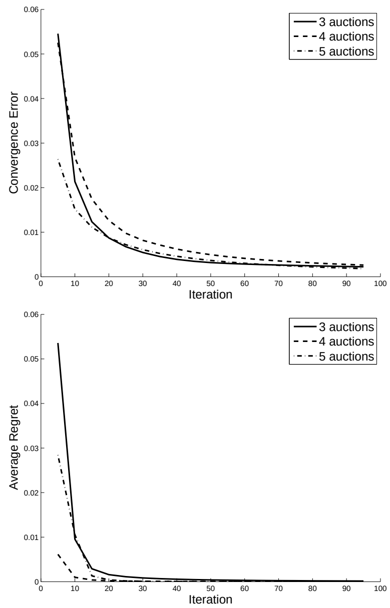

Figure 4: Convergence error (top) and average regret (bottom) for 15 bidders.

one is playing the same strategy at time t, and the utility of a player who 'deviates' and plays a best response given the current beliefs. This difference is then averaged over the entire range of player types to produce the average regret. The average regret serves as an indicator of the convergence in strategies (as opposed to the convergence in beliefs).

In more detail, Figures 3 and 4 depict the convergence error and average regret for a varying number of bidders in 4 simultaneous auctions, and for a varying number of auctions with 15 bidders, respectively.[3] These figures show that the convergence error drops exponentially fast as the algorithm proceeds. Furthermore, it converges even faster with respect to the regret factor.[4] We conjecture that the relative speedup of the regret factor convergence follows from the fact that similar, though distinct, policies may produce the same regret. As a result, the policy continues to change gradually, still keeping the beliefs error CE away from zero, while the regret has already reached low values.

Additional data analysis confirms that the generalised FP algorithm converges in our setting in the strong sense – that is, it converges to a best response strategy, and the corresponding response function h results in a pure Bayes-Nash equilibrium. Moreover, it confirms and expands upon previous conjectures (such as Gerding *et al.* (2008)) on the quality and properties of equilibria in simultaneous auctions. For instance, rather than forming completely distinct graphs, bidding strategies exhibit bifurcation behaviour, holding the same bid value for large intervals in the private value space.[5]

[3] We obtain similar results in other settings.

[4] The variance of convergence rate across the experimental runs was below 10^{-5}.

[5] We omit the details due to space limitations.

6 Conclusions

In this paper we presented a generic procedure for determining the best response computation in anonymous games with continuous player types. Specifically, we constructed a generalised version of the FP algorithm for this setting. We then used the fact that CAPs encompass a significant number of games with private information and applied our generalised FP algorithm to the setting of simultaneous auctions. The algorithm experimentally showed quick convergence and provided, for the first time, a pure Bayes-Nash equilibrium solution for simultaneous auctions with continuous private values.

For the future, we seek to extend this work in the following directions. First, although we have shown convergence empirically for a specific domain, it remains to be seen whether it is possible to derive theoretical guarantees for the FP algorithm to converge in the auction domain, or rather that FP converges generally in CAPs. Our preliminary studies show that, if types can be grouped based on the best response equivalence, FP may not converge, which suggests that additional conditions are needed to obtain convergence. Second, we intend to extend our algorithm to capture continuous (and therefore, infinite) action spaces.

References

M. Blonski. Equilibrium characterization in large anonymous games. Working paper, University of Manheim, 2001.

G. W. Brown. Iterative solutions of games by fictitious play. In *Activity Analysis of Production and Allocation*. Wiley, 1951.

C. Daskalakis and C. H. Papadimitriou. Computing equilibria in anonymous games. In *FOCS*, 2007.

E. H. Gerding, R. K. Dash, D. C. K. Yuen, and N. R. Jennings. Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *AAMAS*, pages 267-274, 2007.

E. H. Gerding, Z. Rabinovich, A. Byde, E. Elkind, and N.R. Jennings. Approximating mixed nash equilibria using smooth fictitious play in simultaneous auctions. In *AAMAS*, 2008.

P. R. Jordan, Y. Vorobeychik, and M. P. Wellman. Searching for approximate equilibria in empirical games. In *AAMAS*, 2008.

M. A. Khan, K. P. Rath, and Y. Sun. On the existence of pure strategy equilibria in games with a continuum of players. *J. of Economic Theory*, 76(1):13–46, 1997.

A. Mas-Colell. On a theorem of Schmeidler. *J. of Mathematical Economics*, 13:201–206, 1984.

D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *J. of Economic Theory*, 68(1):258–265, 1996.

R. Radner and R. W. Rosenthal. Private information and pure-strategy equilibria. *Mathematics of Operations Research*, 7(3):401–409, 1982.

K. P. Rath, Y. Sun, and S. Yamashige. The nonexistence of symmetric equilibria in anonymous games with compact action spaces. *J. of Mathematical Economics*, 24:331–346, 1995.

D. M. Reeves and M. P. Wellman. Computing best-response strategies in infinite games of incomplete information. In *UAI*, 2004.

D. Schmeidler. Equilibrium points of nonatomic games. *J. of Statistical Physics*, 7(4):295–300, 1973.

J. von Neumann and G. W. Brown. Solutions of games by differential equations. In *Contributions to the Theory of Games*, pages 73–79. Princeton University Press, 1950.

E. Vorobeychik and M. P. Wellman. Stochastic search methods for Nash equilibrium approximation in simulation-based games. In *AAMAS*, pages 1055–1062, 2008.