Game Theory with Restricted Strategies

(DRAFT: Please do not circulate.)

Michael Kearns
AT&T Labs
Florham Park, NJ. USA

Yishay Mansour
Tel-Aviv University
Tel-Aviv, Israel

Andrew Y. Ng
U.C. Berkeley
Berkeley, CA. USA

Satinder Singh
AT&T Labs
Florham Park, NJ. USA

Abstract

We present a theory of multiplayer games in which each player chooses their strategy from a limited class that may be significantly weaker than the class of all mixed strategies. As in much of AI and machine learning, such restrictions may be desirable for reasons of computational tractability. We define and study the notions of restricted Nash equilibria, local restricted Nash equilibria, and restricted security strategies. We illustrate a number of these concepts with two well-studied games from the literature, the *War of Attrition* and *Pricebots* games, in which both players are restricted to play mixed strategies that are mixtures of Gaussians.

1 Introduction

Modern research in machine learning and AI has largely adopted the view that restricted representations of the world, in addition to being necessary, may have a variety of conceptual and computational benefits. For example, Bayesian networks offer a language in which some but not all joint distributions may be very succinctly represented, and may additionally admit efficient algorithms for probabilistic inference. Function approximation and factored Markov decision processes (MDPs) in reinforcement learning are approaches to very large state spaces that sacrifice representing all possible policies or MDPs in exchange for faster learning and planning algorithms. Research in supervised learning long ago focused on the necessity and benefits of certain restricted function classes such as decision trees and neural networks.

A similar development of restricted languages, and algorithms that exploit them, needs to be occur in the area of multi-agent systems, where the basic mathematical underpinnings are provided by game theory (Owen 1995). Classical game theory provides a natural formalism for settings in which a collection of agents compete or cooperate in pursuit of individual or global rewards. However, it largely sidesteps the computational and representational issues that arise when the number of actions available to the players becomes very large, or the number of states in the game becomes large.

In this paper, we present first steps towards a theory of multiplayer games in which, due to large or complex action spaces, players choose their strategies from restricted classes that may be significantly handicapped in comparison to the class of all mixed strategies. While a number of earlier works provided interesting results for specific games and restrictions on strategies (Papadimitriou & Yannakakis 1994, Freund etal. 1995), here we attempt to develop a general theory in which both the game and the restrictions are treated generally. We define the general notion of a restricted Nash equilibrium, the natural generalization of the classical definition, and explore conditions for its existence. We also define and study the notion of restricted security strategies, which generalize the classical notion, and the notion of a *local* restricted Nash equilibrium, which is the analogue of local minima in areas like supervised learning. For each of these notions, we examine existence, and compare to the corresponding classical definition. We illustrate a number of these concepts with two well-studied games from the literature, the *War of Attrition* (Smith 1974) and *Pricebots* (Greenwald & Kephart 1999) games, in which both players are restricted to play mixed strategies that are mixtures of Gaussians.

2 Definitions

A two-player, general-sum game is specified by a pair of matrices R and C, which give the payoffs for the *row* player (Player 1) and the *column* player (Player 2). If player 1 chooses his ith action (row) and player 2 chooses her jth action (column), then the payoffs to players 1 and 2 are R_{ij} and C_{ij} respectively. A special case of interest is *zero-sum* games, in which $C = -R$. Note that, for simplicity, we have implicitly assumed that each player has only a finite number of actions. However, most of our results generalize easily to the case of infinite actions, and our examples and exposition will often use the terminology of infinite-action games. Generalizations to n-player games are also straightforward.

A randomized or (*mixed*) *strategy* for player 1 is a vector π_1 whose ith element specifies the probability that player 1 takes their ith action. (Thus, the elements of π_1 are non-negative and sum to 1.) A strategy for player 2 is defined similarly. When we wish to discuss actions more abstractly than as rows and columns, we will write $\pi(a)$ to denote the probability of taking action a under the strategy π. Note that when the two players play strategies (π_1, π_2), the expected payoffs to the two players are $\pi_1^T R \pi_2$ and $\pi_1^T C \pi_2$ respectively (where $\cdot^T$ is the transpose of $\cdot$). A *Nash equilibrium* is a pair of strategies

(π_1, π_2) so that

$$\forall \tilde{\pi}_1, \quad \pi_1^T R \pi_2 \geq \tilde{\pi}_1^T R \pi_2$$
$$\forall \tilde{\pi}_2, \quad \pi_1^T C \pi_2 \geq \pi_1^T R \tilde{\pi}_2$$

where the $\tilde{\pi}_b$ are restricted to be strategies (distributions over actions). To distinguish such equilibria from ones we introduce shortly, we will sometimes also call these *unrestricted Nash equilibria.*

It is well-known that a Nash equilibrium exists for any game (Owen 1995). As discussed in the introduction, classical game theory implicitly assumes a setting in which players are able to represent any mixed strategy in tabular form, an assumption that clearly becomes less plausible when the number of actions is very large or infinite. Thus, as in much of machine learning and AI, we wish to reformulate game theory in a way that allows restricted or parametric representations of mixed strategies. We begin with the following definition, which posits restricted classes of strategies from which players must choose.

Definition. Let a two-player game and restricted *strategy classes* Π_1 and Π_2 be given. A *restricted Nash equilibrium* is a pair (π_1, π_2), $\pi_1 \in \Pi_1, \pi_2 \in \Pi_2$, such that

$$\forall \tilde{\pi}_1 \in \Pi_1, \quad \pi_1^T R \pi_2 \geq \tilde{\pi}_1^T R \pi_2$$
$$\forall \tilde{\pi}_2 \in \Pi_2, \quad \pi_1^T C \pi_2 \geq \pi_1^T C \tilde{\pi}_2$$

Thus, whereas in a Nash equilibrium no player has any incentive to switch to any other strategy, in a restricted Nash equilibrium, no player has any incentive to switch to any other strategy *within their strategy class* Π_i. Naturally, if $\Pi_1 \times \Pi_2$ contains any (unrestricted) Nash pair, then that Nash pair will also be a restricted Nash pair.

As an example of the use of restricted strategies in a case where the action space is infinite, suppose that the actions available to both players correspond to the real numbers (for example, we will shortly study a game in which each player must set a price for a product). Rather than allowing all distributions over the reals, we might naturally restrict players to choose a mixture of k Gaussians:

$$\pi(p) = w_1 \mathcal{N}(p|\mu_1, \sigma_1^2) + \cdots + w_k \mathcal{N}(p|\mu_k, \sigma_k^2) \qquad (1)$$

where $\mathcal{N}(p|\mu, \sigma_2) = (1/\sqrt{2\pi}\sigma)\exp(-(p-\mu)^2/2\sigma^2)$ is a Gaussian density with mean μ and variance σ^2, and the w_i's are the mixture weights ($w_i \geq 0, \sum_i w_i = 1$).

3 Existence of Restricted Equilibria

We begin with the most basic question we study, that of the existence of restricted Nash equilibrium. Let us first observe that, in contrast to the unrestricted setting, restricted equilibria need not exist in general. To see this, note that a special case of the restricted setting is the one in which Π_1 and Π_2 consist of all and only pure (that is, deterministic) strategies. Since it is well-known that pure strategy Nash equilibria do not always exist, this means that we certainly cannot expect restricted Nash equilibria to always exist either [1]. (Here we have used the fact that if a strategy pair is not Nash, then for at least one player, there must be a change to a pure-strategy that yields a higher payoff.) However, there are several natural settings in which we *can* guarantee the existence of restricted Nash equilibria; we detail one of these in the remainder of this section.

Theorem 1. *Let a two-player game and restricted strategy classes Π_1 and Π_2 be given. Suppose that Π_1 and Π_2 are both compact [2] and closed under mixture (that is, for all $\alpha \in [0,1]$ and $\pi, \pi' \in \Pi_i$, we have $\alpha\pi + (1-\alpha)\pi' \in \Pi_i$). Then the game has a restricted Nash equilibrium in the given strategy classes.*

The proof of this Theorem is based on the Kutakani fixed-point Theorem (Owen 1995), and closely follows the proof of Nash's original result. Due to space constraints, we instead state and prove here a special case of the theorem that is nearly as general as the original, and that will be useful to us later.

We say that a class of strategies Π is a *mixture class* if for some finite k, there exist k fixed distributions $p_1, \ldots, p_k$ (over the space of actions), so that Π consists of all and only mixture distributions over the p_i:

$$\Pi = \left\{ \pi : \pi(a) = \sum_i w_i p_i(a); w_i \geq 0, \sum_i w_i = 1 \right\}$$

For instance, the mixture of Gaussians class of strategies given in Equation (1) is a mixture class if the means μ_i and variances σ_i^2 are *fixed*, and the weights w_i are the free parameters of the class.

Theorem 2. *Let a two-player game and restricted strategy classes Π_1 and Π_2 be given. Suppose that Π_1 and Π_2 are both mixture classes. Then the game has a restricted Nash equilibrium in the given strategy classes.*

Proof. Let Π_1 be the set of all mixtures over the distributions $p_1^1, \ldots, p_{k_1}^1$, and Π_2 be the set of all mixtures over $p_1^2, \ldots, p_{k_2}^2$. Let R and C denote the two payoff matrices. Now consider the following modified game. Player 1 has k_1 actions, and Player 2 k_2 actions, if players 1 and 2 choose actions i and j respectively, then their payoffs are $\hat{R}_{ij} = \mathrm{E}_{a_1 \sim p_i^1, a_2 \sim p_j^2}[R_{a_1,a_2}]$ and $\hat{C}_{ij} = \mathrm{E}_{a_1 \sim p_i^1, a_2 \sim p_j^2}[R_{a_1,a_2}]$. Thus, in the new game each player chooses a mixture component, and the payoffs are the expectations under the chosen components in the original game. This is simply a new k_1-by-k_2 matrix game, and thus by Nash's classical result, it must have an unrestricted Nash pair (w^1, w^2). The strategy w^b in the modified game corresponds to some strategy in Π_b in the original game, since Π_b contains all component mixtures. $\square$

It is worth noting that, assuming we have an algorithm for finding Nash equilibria for matrix games, then the proof above also gives us an algorithm for calculating restricted Nash equilibria for mixture classes of strategies. We can simply construct the k_1-by-k_2 modified game with payoffs $\hat{R}$ and

[1] At the end of this section we show that the classes simply containing all pure strategies is sufficient to make the notion of restricted equilibrium synonymous with unrestricted equilibrium.

[2] This is a technical condition that guarantees Π_i has no "holes", and means that if $\pi_1, \pi_2, \pi_3, \ldots$ is a sequence of strategies in Π_b that converges to a limit $\pi = \lim_{i\to\infty} \pi_i$, then $\pi \in \Pi_b$.

$\hat{C}$ as described in the proof, find a Nash equilibrium for it, and take its action probabilities to be our mixture weights in the original mixture strategy classes.

We conclude this section by showing that "sufficiently rich" strategy classes actually cause our restricted theory to coincide with the general, unrestricted theory.

Theorem 3. *Let a two-player game and strategy classes Π_1 and Π_2 be given. Suppose that Π_1 and Π_2 include all pure strategies for the players (in addition to perhaps other strategies). Then every restricted Nash pair (if any exist) will also be an unrestricted Nash pair.*

Proof. Assume for contradiction that there is some restricted Nash pair (π_1, π_2) that is not also an unrestricted Nash pair. Then by definition, player 1 (without loss of generality) must have an incentive to deviate from π_1. This also implies that there must be some action a_1 which player 1 would prefer over π_1, given that player 2 plays π_2. But by assumption, the pure strategy that always chooses a_1 is in Π_1. This contradicts (π_1, π_2) being a restricted Nash pair. □

4 Examples of Restricted Nash Equilibria

Before developing some additional notions of equilibria, we ground the definitions and theory of restricted Nash equilibria developed so far with a few examples.

The first examples are for the "War of attrition" game (Smith 1974). In this game, two animals fight over a prize, and each has to decide how many rounds to keep fighting before giving up. Each round in which both fight results in a -1 penalty to each animal, and on the first round in which only one animal decides to continue fighting, the one that continued wins a prize of $+20$, and the one that gave up gets nothing. If both happened to choose to give up on the same round, then neither gets the prize. The total payoff to each player is the total of the penalties for fighting, plus the prize (if it obtained it).

In the unique symmetric Nash equilibrium, each player chooses a random number of rounds to keep fighting according to a geometric distribution, shown by the dashed line in Figures 1a–c. We used mixture classes for our restricted strategy classes, picking the mixture components to be k fixed, evenly-spaced Gaussian distributions.[3] Following the procedure described in the proof of Theorem 2 to find the restricted Nash equilibria (and using the Gambit software package available from http://www.hss.caltech.edu/gambit/Gambit.html to solve the resulting, small, k-by-k games), all the symmetric restricted equilibria returned by the Nash solver for $k = 3, 5$, and 11 mixture components are shown in Figures 1a–c. The results are remarkably good, with the restricted Nash equilibria in the mixtures of Gaussians approximating the unrestricted Nash equilibria fairly well.

[3] More accurately, since the actions are discrete, we chose "Gaussian-shaped" discrete distributions, $p_i(a) \propto \exp(-(a - \mu_i)^2/2\sigma_i^2)$ for $a \in \mathbb{Z}_0^+$. Other details: Technically, calculating the returns $\hat{R}$ (as in Theorem 2's proof) requires an infinite sum over the integers, that was approximated via a finite sum.

As our second example, we consider the Pricebots game defined by Greenwald and Kephart (1999), a simplified model of an electronic economy in which there are two kinds of buyers: *smart* buyers that check the prices of all sellers and select the one with the lowest price, and *naive* buyers which simply select a random seller to buy from. The percentage of smart buyers in the population is a parameter of this game. The actions of each seller are the continuous-valued prices that the seller can choose, and a Nash equilibrium is specified by a distribution over these prices. Greenwald and Kephart (1999) show interesting "price war dynamics" under reasonable assumptions on the behavior of the sellers.

Using mixtures over $k = 11$ evenly-spaced Gaussian distributions for our restricted strategy classes and again following the procedure outlined in the proof of Theorem 2,[4] the results for 30%, 50%, and 90% "smart" customers are shown in Figures 2a–c. The symmetric Nash equilibria for this problem (as derived by Greenwald & Kephart, 1999) are also plotted, and the graphs again show only the symmetric equilibria returned by our Nash solver. Once more, we find that the results are remarkably good. Only for the case of 90% smart customers does the restricted Nash equilibria fail to approximate the unrestricted Nash equilibria well, but this is not unreasonable since the Gaussian mixture strategy classes cannot really approximate the true Nash equilibrium closely anyway.

5 Approximating Nash Equilibria

In the above examples we found that the probability distribution over actions defined by the restricted Nash equilibria approximated the probability distribution over actions defined by the unrestricted Nash equilibria fairly well. Is this true in general? In this section, we present two results relating restricted Nash equilibria to unrestricted Nash equilibria. The first is a negative result, by means of an artificial example, in which the only restricted Nash equilibrium fails to approximate any unrestricted Nash equilibria, even though there are strategies within the restricted strategy classes that are capable of doing so. The second is a positive result that guarantees that so long as our strategy classes are "sufficiently rich," then every restricted Nash equilibrium will also be an ϵ-Nash equilibrium (defined next).

Definition. Let a two-player game be given. A pair of strategies (π_1, π_2) forms an ϵ-*Nash equilibrium* if

$$\begin{aligned} \forall \tilde{\pi}_1, &\quad \pi_1^T R \pi_2 \geq \tilde{\pi}_1^T R \pi_2 - \epsilon \\ \forall \tilde{\pi}_2, &\quad \pi_1^T C \pi_2 \geq \pi_1^T C \tilde{\pi}_2 - \epsilon \end{aligned}$$

In other words, an ϵ-Nash equilibrium occurs when each player can improve on her current payoff by at most ϵ by deviating from her current strategy.

We now construct an artificial example in which the only restricted Nash equilibrium fails to approximate any unrestricted Nash equilibria — even though there are strategies within the restricted strategy classes that are capable of doing

[4] Experimental details: Here, 0 and 1 were the low and high prices, and the Gaussians were truncated to this range (and renormalized) as necessary. Numerical integration was used to calculate all the payoffs, and we used 3 sellers.

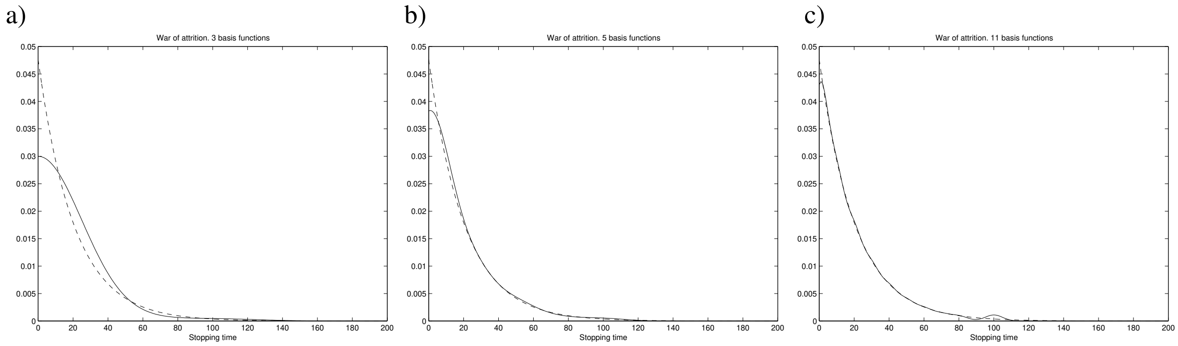

Figure 1: War-of-attrition: dashed line shows unrestricted Nash and solid line shows restricted Nash. a) with mixtures of 3 Gaussians, b) with mixtures of 5 Gaussians, and c) with mixtures of 11 Gaussians.

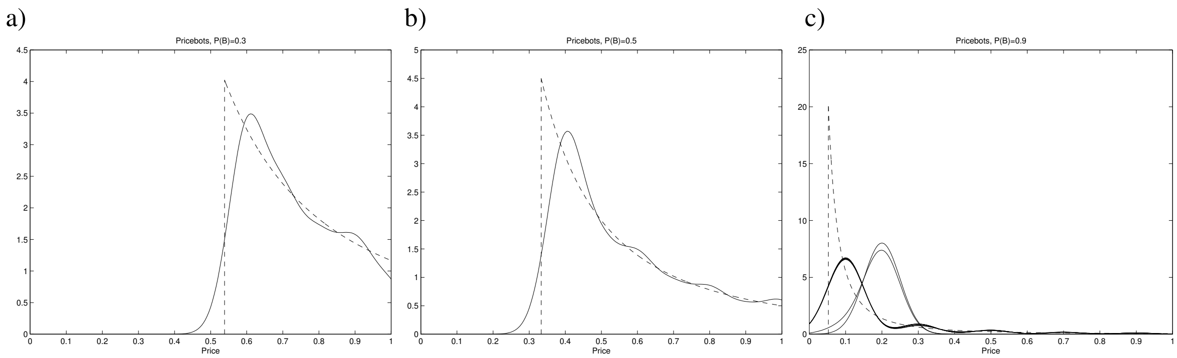

Figure 2: Pricebots: dashed line shows unrestricted Nash and solid line shows restricted Nash. a) with 30% smart customers, b) with 50%smart customers, and c) with 90% smart customers.

so. Consider the 4-action game with payoffs

$$R = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1-\epsilon_1/2 & 1-\epsilon_1/2 & 1-\epsilon_1/2 & 1-\epsilon_1/2 \\ 1 & 1 & 1+2\epsilon_2 & 1+2\epsilon_2 \end{bmatrix}$$

and $C = R^T$. Let $\Pi_1 = \Pi_2 = \{\pi^a, \pi^b\}$ each be a set of two strategies, where $\pi^a = [1-\epsilon_1, \epsilon_1, 0, 0]^T$ and $\pi^b = [0, 0, 1, 0]^T$. (That is, π^a picks action 1 with probability $1-\epsilon_1$, and action 2 with probability ϵ_1.) The pair (π^a, π^a) is an ϵ_1-Nash pair, whose strategies are moreover just ϵ_1 in variation distance away from the Nash pair in which both players chose the pure strategy of action 1. However, whether player 2 is playing π^a or π^b, player 1 can obtain strictly higher return by picking π^b. A similar statement holds for player 2. Thus, the unique restricted Nash pair is (π^b, π^b). However, this is not an ϵ_2-Nash pair, since given the other player's strategy, each player is able to obtain at least an additional $2\epsilon_2$ payoff by switching to action 4. Moreover, action 4 strongly dominates action 3, and hence no strategy in any (unrestricted) Nash pair could assign positive probability to action 3 (since there would always be an incentive to switch to action 4), which implies that the variation distance between π^b and any strategy from an (unrestricted) Nash pair is 1.

Theorem 4. *Let a two-player game with restricted strategy classes be given. Suppose that for some ϵ, Π_1 is "sufficiently rich", in that it can do nearly as well as the best (unrestricted) response to every strategy in Π_2:*

$$\forall \pi_2 \in \Pi_2, \exists \pi_1 \in \Pi_1, \forall \pi_1' : {\pi_1'}^T R \pi_2 \leq \pi_1^T R \pi_2' + \epsilon.$$

Here π_1' ranges over all unrestricted strategies. Suppose that Π_2 is similarly "sufficiently rich" for the same ϵ. Then every restricted Nash equilibrium (if it exists) will also be an ϵ-Nash equilibrium.

Proof. The result follows immediately from the assumptions and definitions. Given a restricted Nash pair (π_1, π_2), by assumption of the strategy classes being "sufficiently rich," player 1's π_1 is at most ϵ-worse than the best response to π_2, and similarly for player 2. But this is exactly the definition of an ϵ-Nash equilibrium. □

6 Local Restricted Nash Equilibria

While, as in the sections above, the restriction of player strategies to parametric classes is a step towards a computationally and representationally efficient theory, we expect that in many natural cases it may still not suffice to render the exact computation of "optimal" solutions tractable. This view is consistent with findings in machine learning and other areas of AI. (For instance, a five-layer neural network with a fixed architecture cannot represent all possible real-valued functions, yet computing the weights that minimize training error remains intractable in the worst case.) In such settings, one often settles for algorithms that can ensure at least local optimality. We thus wish to introduce and examine a notion of local optimality, or local equilibrium, for restricted strategy classes in game theory. We begin with the following definition.

Definition. Let a two-player game and smoothly parameterized strategy classes $\Pi_b = \{\pi_\theta : \theta \in \mathbb{R}^n\}$, $b = 1, 2$ be given. A pair $(\pi_{\theta_1}, \pi_{\theta_2})$, $\pi_{\theta_1} \in \Pi_1, \pi_{\theta_2} \in \Pi_2$, is called a *local* restricted Nash equilibrium if there exists an $\epsilon > 0$ so that

$$\forall \tilde{\theta}_1, ||\tilde{\theta}_1 - \theta_1||_1 \leq \epsilon : \quad \pi_{\theta_1}^T R \pi_{\theta_2} \geq \pi_{\tilde{\theta}_1}^T R \pi_{\theta_2}$$

$$\forall \tilde{\theta}_2, ||\tilde{\theta}_2 - \theta_2||_1 \leq \epsilon : \quad \pi_{\theta_1}^T C \pi_{\theta_2} \geq \pi_{\theta_1}^T C \pi_{\tilde{\theta}_2}$$

In other words, whereas in a restricted Nash pair, neither player had an incentive to switch to any other strategy in his class, in a *local* restricted Nash pair, neither player has any incentive to switch to any other strategy *within a small neighborhood* of her current strategy. Thus, each player is at a "locally" optimal choice of strategies, given her opponent's choice. Clearly, any restricted Nash pair is also a local restricted Nash pair. However, there are also natural settings in which local restricted Nash equilibria will exist that may not also be restricted Nash equilibria. Rather than giving the most general possible result, we present just one simple example here, involving a case where player 1's action space is *factored*, meaning that it is a composite action composed from independent choices of atomic actions. For example, imagine a variation on the Pricebots game in which player's action involves choosing a both a price and quality for the product.

Proposition 5. *Let a two-player game be given, and assume that Player 1's action a_1 is* factored *— thus, $a_1 = (a_{11}, a_{12})$, where $a_{11} \in A_{11}$, $a_{12} \in A_{12}$. Let Player 1's strategy class Π_1 contain all and only product distributions over $A_{11} \times A_{12}$. Let Player 2's strategy class be unrestricted. Then this game has a local restricted Nash equilibria.*

Proof. Consider the following three-player game with players called 1.1, 1.2 and 2. In this game, player 1.1 chooses an action from $a_{11} \in A_{11}$, player 1.2 chooses an action from $a_{12} \in A_{12}$, and player 2 chooses an action from A_2. The payoffs to both players 1.1 and 1.2 are the same as for player 1 when the play is $((a_{11}, a_{12}), a_2)$ in the original game, and the payoff for player 2 is the same as for player 2 in the original game. By Nash's classical result, this three-player game must have a Nash triple $(\pi_{1.1}, \pi_{1.2}, \pi_2)$. Returning to the original two-player game, the pair of strategies $(\pi_{1.1} \times \pi_{1.2}, \pi_2)$ must be a local Nash pair. More precisely, since in the three-player game, neither player 1.1 nor 1.2 had any incentive to deviate from $\pi_{1.1}$ or $\pi_{1.2}$ respectively, in the original two-player game, player 1 also has no incentive to make any *local* change to either the $\pi_{1.1}$ or the $\pi_{1.2}$ component of his strategy. However, such a local restricted Nash pair may not be a restricted Nash pair, since there may still be a way for player 1 to simultaneously change *both* $\pi_{1.1}$ and $\pi_{1.2}$ to improve his payoff. □

Note that the proof again provides an algorithm for finding local restricted Nash equilibrium (constructing the three-player game, and solving for a Nash equilibrium) under the assumed setting. Lastly, invoking ideas from Theorem 2, it is also straightforward to extend this result to (for example) the case in which player 1 is choosing (independent) mixture distributions over A_{11} and A_{12}, or where both players have factored actions.

7 Security Strategies and Zero-Sum Games

In classical game theory, the zero-sum case and its generalization to security levels in the general-sum case enjoy a special status, as there is a unique "value" to the game and elegant algorithms for its computation. In this section we explore the extension of these notions to the restricted strategy setting. We start by reminding the reader of the definition of security levels and security strategies (Fudenberg & Tirole 1991).

Definition. Let a two-player game with payoff matrices R and C be given. A *security strategy* for player 1 is a strategy π_1 such that

$$\forall \tilde{\pi}_1, \min_{\pi_2} \pi_1^T R \pi_2 \geq \min_{\pi_2} \tilde{\pi}_1^T R \pi_2$$

where π_2 must of course be a distribution over the actions of player 2. The quantity $\max_{\pi_1} \min_{\pi_2} \pi_1^T R \pi_2$ is called the *security level* for player 1. Security strategies and security levels for player 2 are defined analogously. In zero sum games, a pair (π_1, π_2) of security strategies for the two players is also a Nash pair (also called a max-min strategy-pair).

For restricted strategy classes, the analogous definitions of restricted security strategies and levels are also straightforward:

Definition. Let a two-player game with restricted strategy classes be given. A *restricted security strategy* for player 1 is a strategy $\pi_1 \in \Pi_1$ so that

$$\forall \tilde{\pi}_1 \in \Pi_1, \min_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2 \geq \min_{\pi_2 \in \Pi_2} \tilde{\pi}_1^T R \pi_2$$

The quantity $\max_{\pi_1 \in \Pi_1} \min_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2$ is called the *restricted security level* for player 1.

While we saw in the previous section that restricted Nash equilibria do not always exist, it is straightforward to guarantee the existence of restricted security strategies and restricted security levels for the case of finite action spaces. Here we give a generalization to the continuous action case.

Proposition 6. *Let a two-player general-sum game with restricted strategy classes be given. Suppose that Π_1 and Π_2 are compact. Then restricted security strategies and restricted security levels are well-defined and exist.*

Proof. We show they exist for player 1. By the compactness of the Π_b and the continuity of $\max_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2$ as a function of π_1, both

$$\max_{\pi_1 \in \Pi_1} \min_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2$$

$$\arg\max_{\pi_1 \in \Pi_1} \min_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2$$

are well-defined and exist. □

Interestingly, however, even for zero-sum games, a pair of restricted security strategies does not, unlike the unrestricted case, necessarily form a Nash pair or restricted Nash pair.

Proposition 7. *There exists a zero-sum game with restricted strategy classes, and two strategies $\pi_1 \in \Pi_1, \pi_2 \in \Pi_2$ that are restricted security strategies for players 1 and 2, but* not *a Nash pair or restricted Nash pair.*

Proof. Consider the penny-matching game, in which each player chooses either action "heads" or actions "tails." If the players' choices match, player 1 receives payoff +1, and player 2 receives -1; otherwise, the payoffs are reversed. Let Π_1 and Π_2 be the sets of all pure strategies. Then any pair $\pi_1 \in \Pi_1, \pi_2 \in \Pi_2$ is a restricted security strategy (with restricted security level -1), but clearly no such pair is a Nash or a restricted Nash pair. □

We now describe two simple algorithms for finding restricted security strategies. In unrestricted games, linear programming is a standard way of finding (unrestricted) security strategies, or Nash pairs in zero-sum games (Fudenberg & Tirole 1991). We now demonstrate that these ideas generalize to the setting of restricted security strategies for the case of mixture classes.

Proposition 8. *Let a two-player game with restricted strategy classes be given. Suppose that both Π_1 and Π_2 are mixture classes. Then there exists an efficient algorithm (using linear programming) that will find the restricted security strategies and levels for both players. If, moreover, the game is a zero-sum game, then such a pair of restricted security strategies will also be a restricted Nash equilibrium.*

Proof (Sketch). As in the proof of Theorem 2, such a game can be reduced to a matrix game in which each player chooses a mixture component. By solving for the (unrestricted) security strategies and levels in this new game and taking those to be the mixture weights chosen by the players for their strategy classes in the original game, we see that, in a similar manner to Theorem 2, we would have the desired restricted security strategies and levels. If the game is zero-sum, then the reduced game is also zero-sum, and in particular the pair of security strategies found would be a Nash pair for the reduced game, and hence (following Theorem 2) a restricted Nash pair for the original game. □

More generally, even in the case of arbitrary, non-mixture strategy classes, we have the following iterative algorithm for finding restricted security strategies and levels in finite-action games:[5]

1. Let π_0 be some strategy from Π_1, and let $S = \emptyset$.
2. Repeat for $i = 1, 2, \ldots$:
   (a) Let $\pi' = \arg\min_{\pi \in \Pi_2} \pi_{i-1}^T R \pi$.
   (b) Let $S := S \cup \pi'$
   (c) Let $\pi_i = \arg\max_{\pi \in \Pi_1} \min_{\pi' \in S} \pi^T R \pi'$.

In order for this algorithm to be well-defined, we must assume that the optimization problems in steps 2(a) and 2(c) can be computationally solved. This holds, for example, if the two strategy classes are convex (closed under mixture).

[5]Generalizing this result to (say) games with infinite/continuous (but bounded) action is also possible if we assume that the payoffs are Lipshitz continuous in the choices of actions; see the extended version of this paper.

It is also easy to prove that this algorithm will converge to the correct answer.

Theorem 9. *Let a two-player finite-action game with compact restricted strategy classes be given, and let $s = \max_{\pi_1 \in \Pi_1} \min_{\pi_2 \in \Pi_2} \pi_1^T R \pi_2$ be the restricted security level of player 1. Then running the above algorithm, we have that $\min_{\pi' \in \Pi_2} \pi_i^T R \pi'$ converges to s as $i \to \infty$.*

Proof. Assume the contrary. Assume without loss of generality that the payoffs are bounded in absolute magnitude by 1. After each execution of step 2(a) of the algorithm, we have a new lower bound on s (given by $\pi_{i-1}^T R \pi'$), and after each execution of step 2(c), we have a new upper bound on s (given by $\min_{\pi' \in S} \pi_i^T R \pi'$). The sequence of upper bounds we get as we run the algorithm is non-increasing, and being bounded (by the minimum payoff) below, must therefore converge to some limit s_U. Now if the theorem does not hold, then our sequence of lower bounds either does not converge, or converges to a different value than the sequence of upper bounds. In either case, there therefore exists $\epsilon > 0$ so that on an infinite number of iterations the lower bounds are at least ϵ below s_U. i.e. There exists an infinite (increasing) sequence $i_1, i_2, \ldots$ so that $\pi_{i-1}^T R \pi' \leq s_U - \epsilon$. on iterations $i = i_1, i = i_2, \ldots$ of the algorithm. Lastly, consider some such pair, $\pi_{i_j - 1}$ and $\pi_{i_k - 1}$, $i_j < i_k$. By our assumptions, there must exist a π' (say, the one generated on the i_j-th iteration of the algorithm) such that $\pi_{i_j-1}^T R \pi' \leq s_U - \epsilon \leq \pi_{i_k-1}^T R \pi' - \epsilon$. Hence, $d(\pi_{i_j-1}, \pi_{i_k-1}) \geq \epsilon/2$, where $d(\cdot,\cdot)$ is the variation distance between distributions. But this implies that we have an infinite sequence of strategies $\pi_{i_1-1}, \pi_{i_2-1}, \ldots$, each pair of which is at least some lower-bounded variation distance from each other. Since each strategy is a point in the finite-dimensional simplex, which is bounded, this is not possible.[6] □

References

Freund Y., Kearns M., Mansour Y., Ron D., Rubinfeld R., and Schapire R.E. (1995). Efficient Algorithms for Learning to Play Repeated Games Against Computationally Bounded Adversaries. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*.

Fudenberg D. and Tirole J. (1991). *Game Theory*. MIT Press.

Greenwald A. R. and Kephart J. O. (1999). Shopbots and Pricebots. In *Proceedings of Sixteenth Joint International Conference on Artificial Intelligence*, pages 506-511

Haussler D. (1992). Decision-theoretic generalizations of the PAC model for neural networks and other applications. *Information and Computation*, 100:78–150.

Owen G. (1995). *Game Theory*. Academic Press, UK.

Papadimitriou C.H. and Yannakakis M. (1994). On complexity as bounded rationality. In *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing*, pages 226–733.

Smith J.M. (1974). The theory of games and evolution of animal conflicts. *Journal of Theoretical Biology*. 47:209–221.

[6]This is a classical fact from information theory. Another way of stating this using terminology from computational learning theory is that the simplex of distributions over the action space has finite packing number with respect to the ℓ_1. (See, e.g. Haussler 1992, or the full paper).