Resource Selection Games with Unknown Number of Players

Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion – Israel Institute of Technology
Haifa 32000, Israel

Abstract

In the context of pre-Bayesian games we analyze resource selection games with unknown number of players. We prove the existence and uniqueness of a symmetric safety-level equilibrium in such games and show that in a game with strictly increasing linear cost functions every player benefits from the common ignorance about the number of players. In order to perform the analysis we define safety-level equilibrium for pre-Bayesian games, and prove that it exists in a compact-continuous-concave setup; in particular it exists in a finite setup.

1 Introduction

In a *resource selection system*, Γ, there is a set of m resources, $j \in \{1, \cdots, m\}$. Each resource j is associated with a cost function $w_j : \{1, 2, \cdots, \cdots\} \rightarrow \Re$, where $w_j(k)$ is the cost for every user of resource j if there are k users. Together with a set of n players a resource selection system defines a game in strategic form– *a resource selection game*, $\Gamma(n)$. The action set[1] of every player i in $\Gamma(n)$ is the set of resources M, and the cost of i depends, via the resource-cost functions, on the resource she chooses and on the number of other players who choose this resource. Thus, resource selection games are special types of congestion games [23, 18]. A resource selection game is also referred to as a simple congestion game.[2] In many situations the assumption that every player knows the number of players is not reasonable. The goal of the current paper is to analyze resource selection games with unknown number of players. One approach for analyzing such situations is the Bayesian approach, where it is assumed that the distribution of the random set of players is

[1] When dealing with games in strategic form the choice set of a player is referred to as an action set or as a strategy set. We use "actions" rather than "strategies" because we keep the notion of strategy to describe the choice set of a player in "bigger" games.

[2] Simple congestion games and their generalization to player-specific games were discussed in [22, 16, 8, 27, 11]. Simple congestion games were also discussed in the price of anarchy literature, e.g., [12, 2]

commonly known. In this approach one is interested in Bayesian equilibrium. [3] The goal of this paper is to analyze resource selection games in the non-Bayesian setup in which the players do not have probabilistic information about the number of the other players in the game. We will use what we call pre-Bayesian games. A pre-Bayesian game is defined to be a Bayesian game without the state probability.[4]

In a classical pre-Bayesian game there is a fixed set of players $I = \{1, 2, \cdots, n\}$, each of them is endowed with a set of actions, $x_i \in X_i$. There is a set of states $\omega \in \Omega$. The payoff of player i, $u_i(\omega, x_1, x_2, \cdots, x_n)$ depends on the realized state, on her choice of action, x_i, and on the choices of all other players. However, the realized state is not known to the players. Every player receives a state-correlated signal, $t_i = \tilde{t}_i(\omega)$ on which she conditions her action. A pre-Bayesian game becomes a Bayesian game when a commonly known probability measure on the set of states is added to the system. The above description is not general enough for our purpose because the set of players is fixed. We generalize the model by allowing the set of active players as well as the action sets to depend on the state.

In pre-Bayesian games one can deal with ex-post equilibrium, which generalize the notion of dominant strategy equilibrium. Such equilibria rarely exist, and consequently only very special type of pre-Bayesian games have been discussed in Game Theory or in Computer Science. Hyafil04 and Boutilier (2004)[10] suggested minimax regret equilibrium as a solution to general pre-Bayesian games. In such equilibrium, every agent, after receiving her informational signal is using the minimax regret criterion in the decision making problem obtained by fixing the behavior of all other agents. They proved that such an equilibrium exists in every finite pre-Bayesian game. In this paper we define another type of equilibrium concept, safety level equilibrium. In such an equilibrium, every player, after receiving her informational signal use the maximin payoff (or minimax cost) criteria in the decision making problem obtained by fixing the behavior of all other agents. We use a more general definition of pre-Bayesian games and prove that a safety level equilibrium always exists. Independently, [1] also define safety level equilibrium. However, they restricted attention to a very particular information structure in which all players receive the same signal. All of the above mentioned existence results rely on Kakutani fixed point theorem, which is a standard tool in game theory for proving equilibria existence results.[5]

---

[3]See [15, 14, 6] for such analysis in the context of auctions, and [19, 20, 21] for such analysis in the context of elections.

[4]Surprisingly, until recently such games did not have a name in the literature. Pre-Bayesian games have been also called *games in informational form* and *games without probabilistic information* [9, 7], *games with incomplete information with strict type uncertainty* [10], and *distribution-free games with incomplete information* [1].

[5]We note that an equilibrium concept for specific types of pre-Bayesian games were already defined in [26, 25, 24]

As we said, our main goal in this paper is to analyze resource selection games with unknown players, and we do it by analyzing safety-level equilibria in such games. However, for the sack of completeness we also define and prove existence of a third type of equilibrium, which is consistent with a large part of single agent setups of incomplete information in CS. We define competitive ratio equilibrium and prove that it exists.[6]

Once we have the right tool we proceed to analyze resource selection games with unknown number of players. We focus on a model with increasing resource cost functions. In order to derive results for the case in which the number of players is unknown we prove several results, some of them are interesting for themselves, about classical resource selection games with known number of players. In particular we prove that every resource selection game possesses a unique mixed-action symmetric equilibrium. When the number of players is k, the unique symmetric equilibrium mixed-action of every player is denoted by p^k. That is, $p^k = (p_1^k, p_2^k, \cdots, p_m^k)$, where p_j^k is the probability that a player chooses resource j.

In our pre-Bayesian model all active players know a common bound, n, on the number of active players, but the players do not know the true number of players, say k, $k \leq n$. Hence, the only signal a player receives is an "activity" signal. A state in this pre-Bayesian game is the set of active players. We prove that a resource selection game with unknown number of players has a unique symmetric safety level equilibrium. In this equilibrium, every active player is using the unique symmetric-equilibrium mixed-action, p^n in the game in which the number of players is commonly known and equals n.

Hence, when the number of players is k, every player uses p^k when the number of players is commonly known, and every player is using p^n when the number of players is unknown. Surprisingly, the lack of knowledge makes each of the players better off in the linear system, in which the resource cost functions are linear! That is, we show that in the linear model, when there are k players, and each of them is using p^n, the cost of each player is at most the cost he obtains in the unique symmetric equilibrium, p^k. Under very modest assumptions every player is strictly better off.

The above results are applicable to a mechanism design setup in which the organizer knows the number of active players, and the players do not know this number. If the goal of the organizer is to maximize revenue then he is better off revealing his private information.[7] If his goal is to maximize

in the context of work on artificial social systems, and in [13], where another equilibrium concept is defined for particular pre-Bayesian auctions. Both groups of authors used versions of safety-level equilibrium.

[6]The competitive ratio approach is relatively common in computer science in the context of single agent decision problems (see e.g., [3]. A recent paper of [5] applies the notion of competitive equilibrium to a particular problem of online algorithms.

[7]This is indirectly related to the Linkage Principle, in auction theory [17].

social surplus then he should not reveal the information. In order to estimate the gain of the players resulting from their ignorance we investigate in the last section the function $c^k(p^n)$ for $n \geq k$, where $c^k(p^n)$ is the cost for a player in a k-player setup when all players play the mixed-action associated with the symmetric equilibrium of the corresponding n-player setup.

2 Background

We remark that games can be analyzed with either payoff functions or cost functions. Most of the general theory has been developed with payoff functions. However, congestion games have been mainly analyzed using cost functions. Translating results proved in one of the setups to the other setup is obvious in most cases.[8] We follow the tradition of previous literature in the sense that the general theory is discussed with payoff functions, while resource selection games are discussed with cost functions.

A game in strategic form is a tuple, $\Gamma = (I, (X_i)_{i \in I}, (u_i)_{i \in I})$. I is a nonempty set of players, X_i is a nonempty set of actions for player i, and $u_i : X \to \Re$ is the payoff function of i, where $X = \times_{i \in I} X_i$ is the set of action profiles. Hence $u_i(x)$ is the payoff of player i when the profile of actions $x \in X$ is played. Γ is a *finite* game if I and X are finite sets.

Let $x \in X$ denote a profile of actions. For each $i \in I$ we let $x_{-i} = (x_j)_{j \in I \setminus \{i\}}$ denote the actions played by everyone but i. Thus $x = (x_i, x_{-i})$. An action profile $x \in X$ is in *equilibrium* if $u_i(x_i, x_{-i}) \geq u_i(y_i, x_{-i})$ for every player $i \in I$ and for every $y_i \in X_i$.

A *permutation* of the set of players is a one-to-one function from I onto I. For every permutation π and for every action profile $x \in X$ we denote by πx the permutation of x by π. That is, $(\pi x)_i = x_{\pi(i)}$ for every player i. Γ is a *symmetric game* if $X_i = X_s$ for all $i, s \in I$, and

$$u_i(\pi x) = u_{\pi(i)}(x)$$

for every player i, for every action profile x, and for every permutation π. A *symmetric action profile* is an action profile x such that $x_i = x_s$ for every $i, s \in I$. x is *a symmetric equilibrium* if it is both an equilibrium profile and a symmetric action profile .

For any finite set C, $\Delta(C)$ denotes the set of probability distributions over C. Let $\Gamma = (I, (X_i)_{i \in I}, (u_i)_{i \in I})$ be a finite game in strategic form. Every $p_i \in \Delta(X_i)$ is called a *mixed action* for i. $p_i(x_i)$ is the probability that player i plays action x_i. Every vector $p \in \Delta = \times_{i \in I} \Delta(X_i)$ defines a probability distribution over X; The probability, $p(x)$ of $x \in X$ is $\prod_{i \in I} p_i(x_i)$.

Let u_i^m be the expected payoff function defined on Δ by u_i. That is, $u_i^m(p) = E_p(u_i)$.

[8] An exception is the price of anarchy theory.

The game $\Gamma^m = (I, (\Delta(X_i))_{i\in I}, (u_i^m)_{i\in I})$ is called the *mixed extension* of Γ. A mixed action for player i in the game Γ is an action for player i in the game Γ^m.

A *mixed action equilibrium* in a finite game Γ is defined to be an equilibrium in Γ^m. It is well-known (see e.g, [4]) that every finite symmetric game in strategic form possesses a symmetric mixed action equilibrium.

3 Resource Selection Games with Known Number of Players

In a *resource selection system*, $\Gamma = (M, (w_j)_{j=1}^m)$ there is a set of resources, $M = \{1, \cdots, m\}$, $m \geq 1$. Every resource j is associated with a cost function $w_j : \{1, 2, \cdots\} \to \Re$. $w_j(k)$ is the cost for every user of resource j if there are k users. Together with the set of players $I_n = \{1, \cdots, n\}$ a resource selection system defines a game in strategic form– *a resource selection game* $\Gamma(I)$. In this paper only the number of players will be relevant (and not the "names" of the players) and therefore we will denote a resource selection game by $\Gamma(n)$. The action set of every player i in $\Gamma(n)$ is the set of resources M, and the cost of i depends on the resource she chooses and on the number of other players that choose this resource via the resource-cost functions. That is, $X_i^n = M$ for every $1 \leq i \leq n$, and $c_i^n(x) = w_{x_i}(\sigma_{x_i}(x))$, where for every resource j and for every action profile $x \in \times_{i=1}^n X_i^n = M^n$, $\sigma_j(x)$ is the number of all players s for which $x_s = j$. Obviously every resource selection game is a finite symmetric game.

Let $p \in \Delta(M)$ be a mixed action of an arbitrary player. That is, $p = (p_1, \cdots, p_m)$, where p_j is the probability that a player who uses the mixed action p will select resource j. We denote the support of p by *supp*(p). That is *supp*$(p) = \{j \in M | p_j > 0\}$. Denote by $c^n(p, j)$ the expected cost of a player that chooses resource j when each of the other $n-1$ players in $\Gamma(n)$ is using p. Let $c^n(p)$ be the expected cost of every player when each of the n players in $\Gamma(n)$ is choosing p.

For every $n \geq 1$, and for every $0 \leq \alpha \leq 1$. Let $Y_\alpha^n \sim Bin(n, \alpha)$ be a binomial random variable. That is, $f_\alpha^n(k) = P(Y_\alpha^n = k) = \binom{n}{k}\alpha^k(1-\alpha)^{n-k}$ for every $0 \leq k \leq n$. Let $F_\alpha^n(k) = P(Y_\alpha^n \leq k)$ be the distribution function of Y_α^n. Obviously

$$c^n(p, j) = E(w_j(1 + Y_{p_j}^{n-1})), \tag{1}$$

where E stands for the expectation operator. That is,

$$c^n(p, j) = \sum_{s=0}^{n-1} w_j(s+1) f_{p_j}^{n-1}(s). \tag{2}$$

Let $(q, \cdots, q) \in \Delta(M)^n$ be a symmetric mixed-action equilibrium profile in $\Gamma(n)$. We will refer to q as a *symmetric-equilibrium action*.

Theorem 1 *Every resource selection game with at least two players ($n \geq 2$) with increasing*[9] *resource cost functions possesses a unique symmetric mixed-action equilibrium.*

In order to prove Theorem 1 we need some preparations.

Lemma 1 *Let $n \geq 1$. $F^n_\alpha(k)$ is a strictly decreasing function of α for every $0 \leq k \leq n-1$.*

Proof: We show that the derivative of $F^n_\alpha(k)$ by α is negative for every $0 \leq k \leq n-1$.

$$\frac{\partial F^n_\alpha(k)}{\partial \alpha} = \sum_{i=0}^{k} \binom{n}{i} [i\alpha^{i-1}(1-\alpha)^{n-i} + (n-i)\alpha^i(1-\alpha)^{n-i-1}] =$$

$$\sum_{i=0}^{k} \binom{n}{i} [i\alpha^{i-1}(1-\alpha)^{n-i} - (n-i)\alpha^i(1-\alpha)^{n-i-1}].$$

Let $f(k) \triangleq \frac{\partial F^n_\alpha(k)}{\partial \alpha}\alpha(1-\alpha)$. Therefore

$$f(k) = \sum_{i=0}^{k} \binom{n}{i} [(1-\alpha)i\alpha^i(1-\alpha)^{n-i} - \alpha(n-i)\alpha^i(1-\alpha)^{n-i}] =$$

$$\sum_{i=0}^{k} \binom{n}{i} [\alpha^i(1-\alpha)^{n-i}(i-n\alpha)].$$

Observe that $\sum_{i=0}^{k} \binom{n}{i}\alpha^i(1-\alpha)^{n-i}i = \sum_{i=1}^{k} p(Y^n_\alpha \geq i) - k(1 - p(Y^n_\alpha \leq k))$. Therefore

$$f(k) = \sum_{i=0}^{k} \binom{n}{i} \alpha^i(1-\alpha)^{n-i}(i-n\alpha) = \sum_{i=1}^{k} p(Y^n_\alpha \geq i) - k + (k-n\alpha)p(Y^n_\alpha \leq k).$$

Obviously $\sum_{i=1}^{k} p(Y^n_\alpha \geq i) - k < 0$. We distinguish between two cases:

1. $k - n\alpha \leq 0$: This case will yield immediately that $f(k)$ is negative.

2. $k - n\alpha > 0$: Observe that $f(n) = 0$. We look at the difference between $f(k+1)$ and $f(k)$. $f(k+1) - f(k) = (k+1-n\alpha)p(Y^n_\alpha = k+1) > 0$. Since the differences are positive and $f(n) = 0$ then $f(k) < 0$.

□

Lemma 2 *Let $\Gamma(n)$, $n \geq 2$ be a resource selection game. Let $q, p \in \Delta(M)$ be mixed actions, and let $j \in M$ be a resource such that w_j is increasing in $\{1, 2, \cdots, n\}$. If $p_j > q_j$ then $c^n(p, j) > c^n(q, j)$.*

[9] That is, $w_j(k) < w_j(k+1)$ for all j and k.

Proof: We have to show that $c^n(p,j)$ is increasing in p_j.

By manipulating (2)

$$c^n(p,j) = \sum_{k=1}^{n-1}\left((w_j(k) - w_j(k+1))\sum_{l=0}^{k-1} f_{p_j}^{n-1}(l)\right) + w_j(n)\sum_{l=0}^{n-1} f_{p_j}^{n-1}(l) =$$

$$w_j(n) - [\sum_{k=1}^{n-1}(w_j(k+1) - w_j(k))F_{p_j}^{n-1}(k-1)],$$

where the last equality follows from the fact that $\sum_{k=0}^{n-1} f_{p_j}^{n-1}(k) = 1$.

By lemma 1, $F_{p_j}^{n-1}(k)$ is strictly decreasing in p_j for every $k = 0, ..., n-2$.In addition, w_j is strictly increasing, and therefore $c^n(p,j)$ is strictly increasing in p_j. $\square$

Proof of Theorem 1

Suppose in negation that there is more than one mixed-action symmetric equilibrium in $\Gamma(n)$. Let q and p be two symmetric equilibrium actions with $p \neq q$. Since $p \neq q$ there exists $j \in M$ with $q_j \neq p_j$. W.l.o.g $q_j > p_j$. Therefore there exist a resource $r \in M$ such that $r \neq j$ and $q_r < p_r$. We get a contradiction from the following sequence of inequalities: $c^n(q,j) > c^n(p,j) \geq c^n(p,r) > c^n(q,r) \geq c^n(q,j)$, where the strict inequalities hold hold by Lemma 2 and the other inequalities hold because q and p are equilibrium actions. $\square$

For every $n \geq 1$ we will denote the unique symmetric equilibrium in $\Gamma(n)$ by p^n, and we denote by $c^n = c^n(p^n)$ the equilibrium cost of a player in $\Gamma(n)$.

We say that a resource cost function w_j is *convex* if it can be extended to a convex function on $[1, \infty)$.

The following lemma will be useful later.

Lemma 3 *Let $\Gamma = (m, (w_j)_{j=1}^m)$ be a resource selection system, with increasing and convex cost functions. There exists an integer $N \geq 2$, $N = N(\Gamma)$ such that for every $n \geq N$, the unique symmetric-equilibrium action in the game $\Gamma(n)$, $p^n \in \Delta(M)$ has a full support. That is, $p_r^n > 0$ for every $1 \leq r \leq m$.*

Proof: Recall that p^n is the unique symmetric-equilibrium action in $\Gamma(n)$, and that $c^n = c^n(p^n)$ is the symmetric-equilibrium cost of every player. As p^n is in equilibrium, $c^n(p^n,j) = c^n$ for every $j \in supp(p^n)$. For every resource j we denote by w_j the convex extension of w_j to $[0, \infty)$. As w_j is convex,

$$c^n(p^n,j) = E(w_j(1 + Y_{p_j^n}^{n-1})) \geq w_j(1 + E(Y_{p_j^n}^{n-1})) =$$

$$w_j(1 + p_j^n(n-1)),$$

where the first equality follows from (1), the inequality follows from the convexity of w_j, and the last equality follows from the well-known fact that

$$E(Y_\alpha^n) = \alpha n. \tag{3}$$

Obviously, there exists $j \in supp(p)$ for which $p_j^n \geq \frac{1}{m}$. For this resource j

$$c^n = c^n(p^n, j) \geq w_j(1 + \frac{1}{m}(n-1)) \geq$$

$$\min_{r=1}^{m} w_r(1 + \frac{1}{m}(n-1)).$$

Since w_j is increasing and convex, $\lim_{n\to\infty} w_j(n) = \infty$ for every resource j. Therefore $\lim_{n\to\infty} c^n = \infty$. Hence, there exists N such that for every $n \geq N$ $c^n > \max_{j=1}^m w_j(1)$. We claim that for every $n \geq N$, $p_r^n > 0$ for every $1 \leq r \leq m$. Indeed, if $p_r^n = 0$ for some r then because $c^n > w_r(1)$, a player will decreases her cost by deviating from p^n to r (assuming every other player is using p^n). This contradicts p^n being a symmetric-equilibrium action.

4 Equilibrium in Pre-Bayesian Games

In this section we define a general model of pre-Bayesian games, and present and prove existence results for safety-level, and competitive-ratio equilibria in such games. We also give the definition of minimax-regret equilibrium (existence proved by Hyafil and Boutilier [10]).

A *pre-Bayesian game* is a tuple $H = (\Omega, I, (Z_i)_{i\in I}, J, (X_i)_{i\in I}, (u_i(\omega, *)_{\omega\in\Omega, i\in I}, (T_i)_{i\in I})$, where,

- Ω is a set of *states*.

- I is a finite set whose elements are called *potential players*.

- $J : \Omega \to 2^I \setminus \{\emptyset\}$ is the function determining the set of active players. That is, i is *active* at ω if $i \in J(w)$.

  Let Ω_i be the subset of states at which i is active. That is, $\Omega_i = \{\omega \in \Omega | i \in J(w)\}$.

- Z_i is the set of potential actions of i.

  Only a subset of Z_i will be available to i at a given state.

- For every potential player i, $X_i : \Omega_i \to 2^{Z_i} \setminus \{\emptyset\}$ is the variable determining the set of actions that are available for i at ω.

  Let $X(\omega) = \times_{i\in J(\omega)} X_i(w)$ be the set of action profiles that can be generated at ω.

- For every potential player i, and for every $\omega \in \Omega_i$, $u_i(\omega, *) : X(\omega) \to \Re$ is the *payoff function* of i at the state ω.

- T_i is the set of possible signals that i may receive.

  We will sometimes follow the traditional approach in economics and refer to $t_i \in T_i$ as a *type* of i.

- $\tilde{t}_i : \Omega_i \to T_i$ is the *signaling function of i.*

  That is, whenever i is active at ω she receives the signal $\tilde{t}_i(\omega)$.

  Without loss of generality we assume that every type is possible. That is, $\tilde{t}_i(\Omega_i) = T_i$.

  Let $\Omega_i(t_i)$ be the set of all states ω that generate the type t_i. That is, $\Omega_i(t_i) = \{\omega \in \Omega_i | \tilde{t}_i(\omega) = t_i\}$.

- $\tilde{t}_i(\omega) = \tilde{t}_i(\omega')$ implies $X_i(\omega) = X_i(\omega')$.

  That is, players know their set of actions; In other words, X_i is a constant function over $\Omega_i(t_i)$ for every $t_i \in T_i$.

  This assumption allows us to define the set of available actions for i when his type is t_i. Indeed we define $X_i(t_i) = X_i(\omega)$ for any arbitrary $\omega \in \Omega_i(t_i)$.

The pre-Bayesian game proceeds as follows. Nature chooses $\omega \in \Omega$, and the game in strategic form $\Gamma(\omega) = (J(w), (X_i(\omega))_{i \in J(\omega)}, (u_i(\omega, *))_{i \in J(\omega)})$ is played. Every active player in the game, $i \in J(\omega)$ receives the signal $\tilde{t}_i(\omega)$, which reveals part of the information about the game. The minimal information is the set of available actions at ω. Thus, every $i \in J(\omega)$ chooses $x_i \in X_i(\omega)$ and his payoff is $u_i(\omega, x)$, where $x = (x_i)_{i \in J(\omega)}$.

A *strategy* for a potential player i in the pre-Bayesian game H is a function $s_i : T_i \to Z_i$ that satisfies: $s_i(t_i) \in X_i(t_i)$ for every $t_i \in T_i$. Hence, i follows the strategy s_i if whenever he is active and receives the signal t_i he chooses the action $s_i(t_i)$. The set of strategies of i is denoted by $\Sigma_i = \Sigma_i(H)$, and the set of strategy profiles is denoted by $\Sigma = \Sigma(H)$, that is $\Sigma = \times_{i \in I} \Sigma_i$.

Let H be a pre-Bayesian game. We say that H is *finite* if the set of states, and the sets of potential actions are finite.

We say that H is a *compact-continuous pre-Bayesian game* if the set of states are finite, Z_i is a compact subset of some Euclidean space[10], $X_i(t_i)$ is a compact subset of Z_i for every type t_i, and $u_i(\omega, *)$ is continuous on $X(\omega)$ for every $\omega \in \Omega_i$.

[10] Or of some linear topological space.

A compact continuous pre-Bayesian game is *(quasi) concave* if Z_i is a convex set, $X_i(t_i)$ is a convex subset of Z_i for every type t_i, and $u_i(\omega, *, x_{-i})$ is (quasi) concave on $X_i(\omega)$ for every $\omega \in \Omega_i$ and every $x_{-i} \in X_{-i}(\omega)$.

Let $H = (\Omega, I, (Z_i)_{i\in I}, J, (X_i)_{i\in I}, (u_i(\omega, *)_{\omega\in\Omega, i\in I},$
$(T_i)_{i\in I})$ be a finite pre-Bayesian game. We are about to define the mixed extension, H^m of H by allowing every active player at $\omega \in \Omega$ to choose a mixed action $\mu_i \in \Delta(X_i(\omega))$. However we need the following definition.

For every $\omega \in \Omega_i$ let $\Delta_i(\omega)$ be the set of all probability distributions $\mu_i \in \Delta(Z_i)$ that vanishes outside $X_i(\omega)$. Obviously, $\Delta_i(\omega)$ can be identified with $\Delta(X_i(\omega))$.

The pre-Bayesian game $H^m = (\Omega, I, (\Delta(Z_i))_{i\in I}, J, (\Delta_i)_{i\in I}, (u_i^m(\omega, *)_{\omega\in\Omega, i\in I}, (T_i)_{i\in I})$ is called the *mixed extension* of H, where $u_i^m(\omega, *)$ is the expected-payoff function defined by u_i. Obviously, H^m is a concave pre-Bayesian game. Every strategy of i in H^m is called a *mixed-strategy for i in H*.

4.1 Safety-level equilibrium

Let H be a compact-continuous pre-Bayesian game. For every profile, $s \in \Sigma$ of strategies in H, and for every state ω we denote by $s[\omega]$ the profile of actions chosen in $X(w)$ when every player i uses s_i. That is, $s[\omega] = (s_j(\tilde{t}_j(\omega))_{j\in J(\omega)}$. As usual, $s_{-i}[\omega]$ is obtained from $s[\omega]$ by removing the action chosen by i.

Let $i \in I$, $t_i \in T_i$, $x_i \in X_i(t_i)$, and let $s_{-i} = (s_j)_{j\in I\setminus i}$ be a profile of strategies of all players but i. The worst case cost of i is

$$W_i(t_i, s_{-i}, x_i) = \min_{\omega\in\Omega_i(t_i)} u_i(\omega, x_i, s_{-i}[\omega]).$$

Obviously $W_i(t_i, s_{-i}, *)$ is continuous on $X_i(t_i)$. We say that $x_i^* \in X_i(t_i)$ is *optimal* for type t_i given s_{-i} if the maximal value of $W_i(t_i, s_{-i}, x_i)$ over $x_i \in X_i(t_i)$ is attained at x_i^*. A strategy of player i, s_i, is a *safety-level best-response to s_{-i}* if for every type t_i, $s_i(t_i)$ is optimal for t_i given s_{-i}. A strategy profile $s = (s_i)_{i\in I}$ is called a *safety-level equilibrium* if for every i, s_i is a safety-level best-response to s_{-i}.

Hence, s is a safety-level equilibrium if and only if for every ω, and for every player i, which is active at ω, $s_i(\tilde{t}_i(\omega))$ is optimal for $\tilde{t}_i(\omega)$ given s_{-i}.

A safety-level equilibrium in a pre-Bayesian game with exactly one state is simply a Nash equilibrium in this game. We next show that safety-level equilibria exist in every quasi concave game.

Theorem 2 *Every quasi concave pre-Bayesian game possesses a safety-level equilibrium.*

Proof: Note that for every i Σ_i is a closed convex and compact subset of $Z_i^{T_i}$. Therefore Σ is a closed convex and compact set because it is a cartesian products of such sets. For every $s \in \Sigma$ and for every $i \in I$ let $B_i(s) \subseteq \Sigma_i$ be the set of all $d_i \in \Sigma_i$, which are best response to s_{-i}. Let $B(s) = \times_{i \in I} B_i(s) \subseteq \Sigma$. It is standard to check that the correspondence $s \to B(s)$ satisfies the conditions of Kakutani's fixed point theorem. That is, it is upper hemicontinuous, and $B(s)$ is a nonempty compact convex subset of Σ for every $s \in \Sigma$. Therefore there exists a fixed point s^*, that is, $s_i^* \in B_i(s^*)$ for every $i \in I$. Obviously such a fixed point is a safety-level equilibrium.

Hence, if H is a finite game, H^m possesses a safety-level equilibrium. Every such an equilibrium is called a *mixed- strategy safety-level equilibrium* in H.

4.2 Minimax-regret equilibrium

Minimax-regret equilibria were defined by Hyafil and Boutilier in [10]. We use their definition for our general model of pre-Bayesian games.

Let H be a compact-continuous pre-Bayesian game. Let $i \in I$, $t_i \in T_i$ and let $s_{-i} = (s_j)_{j \in I \setminus i}$ be a profile of strategies of all players but i. Let $\omega \in \Omega_i(t_i)$. The *regret* of $x_i \in X_i(t_i)$ at w is defined as:

$$R(x_i, w, t_i, s_{-i}) =$$

$$\max_{z_i \in X_i(t_i)} [u_i(w, z_i, s_{-i}[\omega]) - u_i(w, x_i, s_{-i}[\omega]].$$

The *maximal regret* of $x_i \in X_i(t_i)$ over all $\omega \in \Omega_i(t_i)$ is denoted by $MR(x_i, t_i, s_{-i})$. That is:

$$MR(x_i, t_i, s_{-i}) = \max_{\omega \in \Omega_i(t_i)} R(x_i, w, t_i, s_{-i}).$$

We say that $y_i \in X_i(t_i)$ is a minimax regret strategy at t_i given s_{-i} if the minimal value of $MR(x_i, t_i, s_{-i})$ over $x_i \in X_i(t_i)$ is attained at y_i. A strategy s_i is a minimax regret best response to s_{-i} if for every type t_i, $s_i(t_i)$ is a minimax regret strategy at t_i given s_{-i}. s is a *minimax regret equilibrium* if for every player i s_i is a minimax regret best response to s_{-i}. In [10] the authors show:

Theorem 3 (Hyafil and Boutilier) *Every quasi concave pre-Bayesian game possesses a minimax-regret equilibrium.*

4.3 Competitive ratio equilibrium

Competitive ratio equilibrium resembles the minimax-regret equilibrium. They differ only in the definition of regret.

Let H be a compact-continuous pre-Bayesian game, where all payoff functions are positive. Let $s_{-i} = (s_j)_{j \in I \setminus i}$ be a profile of strategies of all players but i, let t_i be an active type of i, and let $\omega \in \Omega(t_i)$. Replace the definition of $R(x_i, w, t_i, s_{-i})$ in the previous section with the following:

$$\hat{R}(x_i, w, t_i, b_{-i}) = \max_{z_i \in X_i(t_i)} \frac{u_i(w, z_i, s_{-i}[\omega])}{u_i(w, x_i, s_{-i}[\omega])}.$$

Note that $\hat{R}$ is well defined since all payoff functions are positive. A strategy profile s is a competitive-ratio equilibrium if for every player i s_i is a minimax-regret best response to s_{-i} with respect to the regret function $\hat{R}$.

Theorem 4 *Every quasi concave game in pre-Bayesian game with positive payoffs functions possesses a competitive ratio equilibrium.*

The proof follows from Theorem 3 by applying the logarithmic function to the payoff functions.

5 Resource Selection Games with Unknown Number of Players

Consider a fixed resource selection system, Γ with the set of resources $M = \{1, \cdots, m\}$, $m \geq 1$, and resource cost functions $(w_j)_{j=1}^m$.

We proceed to describe our model of resource selection games with unknown number of players. Let $I = \{1, 2, \cdots, n\}$, $n \geq 1$. be the set of potential players. The set of states, Ω is the set $2^I \setminus \{\emptyset\}$ of all nonempty subsets of I. The set of active players at state K is K itself. That is $J(K) = K$ for every $K \in \Omega$. The set of actions of player i is the set of resources. That is, $Z_i = M$ for every player i. The set of types of $i \in I$ is a singleton, $T_i = \{active\}$. That is, an active player knows that he is active. The non-active players do not receive any signal. For every player $i \in I$ and every $\omega \in \Omega_i$ $X_i(\omega) = M$. For every $i \in I$ and for every $K \in \Omega$ such that $i \in K$ the cost function of i $c_i(K, *) : Z \to \Re$ is defined by $c_i(K, z) = w_{z_i}(\sigma_{z_i}^K(z))$, where $\sigma_{z_i}^K(z)$ is the number of all players $l \in K$ for which $z_l = z_i$. The above pre-Bayesian game is finite. We denote its mixed extension by $H_\Gamma(n)$, and we referred to $H_\Gamma(n)$ as a *resource selection game with unknown number of players.* A strategy of player i in $H_\Gamma(n)$ can be described by a mixed action $q[i] \in \Delta(M)$. That is, when receiving the signal $\{active\}$, i uses $q[i]$. Since we deal with costs and not with payoffs we use minimax rather than maximin in the definition of safety-level equilibrium. let $\mu = (q[1], \cdots, q[n])$

be a strategy profile in $H_\Gamma(n)$, and let i be a player. Let $\mu[-i] = (q[l])_{l \in I \setminus \{i\}}$ be the profile of strategies of the other players. Observe that Ω_i is the set of all nonempty subsets of I that contain i. If all cost functions are non-decreasing, and if i believes that all other players are using the profile $q[-i]$, then it is obvious that the worst case scenario for i is obtained in the state I. Thus we have:

Lemma 4 *Let Γ be a resource selection system in which the resource cost functions are non-decreasing. Let $\mu \in \Delta(M)^n$. μ is a safety-level equilibrium in $H_\Gamma(n)$ if and only if μ is a mixed-action equilibrium in $\Gamma(n)$.*

Proof: Assume μ is a mixed-action equilibrium in $\Gamma(n)$. Let i be an active player. By the comment we made before the statement of the lemma,

$$\min_{p[i] \in \Delta(M)} \max_{K \in \Omega_i} c_i(K, p[i], \mu[-i]) =$$

$$\min_{p[i] \in \Delta(M)} c_i(I, p[i], \mu[-i]).$$

Because μ is a mixed-action equilibrium in $\Gamma(n)$, the min in the right hand-side of the above formula is attained at $q[i]$. Therefore μ is a safety-level equilibrium in $H_\Gamma(n)$. An analogous argument proves the if part of the lemma.

Theorem 5 *Let Γ be a resource selection system in which the resource cost functions are non-decreasing. $H_\Gamma(n)$ has a unique symmetric safety-level equilibrium. In this symmetric safety-level equilibrium every player is using the strategy p^n, where p^n is the unique symmetric-equilibrium action in $\Gamma(n)$.*

Proof: The proof follows directly from Theorem 1 and Lemma 4.

By Theorem 5, each of the players in $H_\Gamma(n)$ is using the strategy p^n, where p^n is the unique symmetric-equilibrium mixed action in $\Gamma(n)$. However, the cost of each active player in $H_\Gamma(n)$ is not $c^n = c^n(p^n)$, it depends on the true state. If the true state is K, that is K is the set of active players, and $|K| = k$, the cost of each active player i is $c^k(p^n)$. It is worthy to compare this cost with the cost $c^k = c^k(p^k)$ that every player in K would have paid had the players in K known the state. We make these comparison in linear models. We say that a resource selection system is *linear* if for every resource j there exists a constant d_j such that $w_j(k) = w_j(1) + (k-1)d_j$ for every $k \geq 1$. For every number of players, n, the associated resource selection game, $\Gamma(n)$, as well as the associated resource selection game with unknown player set, $H_\Gamma(n)$ will be called *linear* too. Note that in a linear system, w_j is increasing if and only if $d_j > 0$, and w_j is non-decreasing if and only if $d_j \geq 0$.

Theorem 6 *Let Γ be a linear resource selection system with increasing resource cost functions. For every $k \geq 2$ let p^k be the unique symmetric equilibrium in $\Gamma(k)$. There exist an integer $N = N(\Gamma)$, $N \geq 2$ such that for all $n > k \geq N$:*

1. $c^k(p^k) \geq c^k(p^n)$.
2. *All inequalities above are strict if and only if there exists $j_1, j_2 \in M$ such that $w_{j_2}(1) \neq w_{j_1}(1)$*

Proof of Theorem 6

1. Let $j \in M$. By lemma 3 there exist an integer N such that for every $n \geq N$ the unique symmetric equilibrium of $\Gamma(n)$ has full support. By (1) and (3),

   $c^k = c^k(p^k) = w_j(1) + (k-1)d_j p_j^k$ for every $k \geq N$. Let n, k be integers such that $n > k \geq N$. We begin by showing that $c^k(p^k) \geq c^k(p^n)$. Let $c^k = c^k(p)$.

$$c^k(p^n) = \sum_{j=1}^{m} p_j^n c^k(p^n, j) = \sum_{j=1}^{m} p_j^n [c^n - (n-k)d_j p_j^n] = c^n - (n-k)\sum_{j=1}^{m} d_j (p_j^n)^2. \quad (1)$$

   It remains to show that

$$c^n - c^k \leq (n-k)\sum_{j=1}^{m} d_j (p_j^n)^2 \quad (2)$$

   We have

$$p_j^n = \frac{c^n - w_j(1)}{d_j(n-1)}, \quad p_j^k = \frac{c^k - w_j(1)}{d_j(k-1)}.$$

   Let $A = \sum_{j=1}^m \frac{1}{d_j}$ and $B = \sum_{j=1}^m \frac{w_j(1)}{d_j}$. Since $\sum_{j=1}^m p_j^k = \sum_{j=1}^m p_j^n = 1$ we have that $c^n = \frac{(n-1)+B}{A}$ and $c^k = \frac{(k-1)+B}{A}$. Hence

$$c^n - c^k = \frac{n-k}{A}.$$

   Together with (2) it remains to show that $A\sum_{j=1}^m d_j(p_j^n)^2 \geq 1$. Let $L = \sum_{j=1}^m \frac{(w_j(1))^2}{d_j}$. We have

$$A\sum_{j=1}^{m} d_j (p_j^n)^2 = A\sum_{j=1}^{m} \frac{(c^n - w_j(1))^2}{d_j(n-1)^2} = \frac{A}{(n-1)^2}[(c^n)^2 A - 2c^n B + L] =$$

$$\frac{1}{(n-1)^2}[((n-1)+B)^2 - 2((n-1)+B)B + LA] =$$

$$\frac{1}{(n-1)^2}[(n-1)^2 - B^2 + LA] = 1 - \frac{B^2 - LA}{(n-1)^2}. \quad (3)$$

   It remains to show that $LA - B^2 \geq 0$. This is immediate since for every couple of resources $j, r \in M$ such that $j \neq r$, $(w_j(1))^2 + (w_r(1))^2 \geq 2w_j(1)w_r(1)$.

2. Observe that equality holds if and only if $w_j(1) = w_r(1)$ for every $j, r \in M$.

□

Theorem 6 is applicable to a mechanism design setup in which the organizer knows the number of active players, and the players do not know this number. If the goal of the organizer is to maximize revenue then he is better off revealing his private information. If his goal is to maximize social surplus, then he should not reveal that information. In order to estimate the gain of the players resulting from their ignorance we analyze the function $c^k(p^n)$.

Theorem 7 *Let Γ be a linear resource selection system with increasing resource cost functions. There exist N such that for every $n \geq N$ the following assertions hold:*

1. $p_j^n = \frac{n-1+B-w_j(1)A}{Ad_j(n-1)}$, *where* $A = \sum_{j=1}^m \frac{1}{d_j}$, *and* $B = \sum_{j=1}^m \frac{w_j(1)}{d_j}$.

2. *The minimal social cost in $\Gamma(k)$ attained with symmetric mixed-action profiles is attained at p^{2k-1}. Consequently, $c^k(p^n)$ is minimized at $n = 2k-1$.*

Proof:

1. Let $q_j^n = \frac{1}{Ad_j} + \frac{B-w_j(1)A}{Ad_j(n-1)}$ for every $j \in M$. Let N be the smallest n such that $q_j^n > 0$ for every $j \in M$. Observe that for every $n \geq N$ $q_j^n > 0$ for every $j \in M$. Notice that for every n, $\sum_{j=1}^m q_j^n = 1$. Let $n \geq N$. Let $p^n = q^n$. It is enough to show that if all players but i play the mixed-action p^n then player i is indifferent between all resources. That is $c^n(p^n, j) = c^n(p^n, r)$ for every $j, r \in M$. $c^n(p^n, j) = w_j(1) + (n-1)d_j p_j^n = \frac{n-1}{A} + \frac{B}{A}$ which doesn't depend on j.

2. Let $k \geq N$. We need to show that $\sum_{j=1}^m p_j c^k(p,j)$ is minimized at $p = p^{2k-1}$ (subject to p being a probability distribution).

$$\sum_{j=1}^m p_j c^k(p,j) = \sum_{j=1}^m p_j^2 d_j(k-1) + p_j w_j(1).$$

This is a convex program. Therefore it is enough to find a p that satisfies the Kuhn Tucker conditions: there exist a (Lagrange multiplier) λ such that for every j:

$$2p_j d_j(k-1) + w_j(1) = \lambda.$$

With the condition that $\sum_{j=1}^m p_j = 1$ we have $\lambda = \frac{2(k-1)+B}{A}$. Therefore $p_j = \frac{2(k-1)+B-w_j(1)A}{2(k-1)Ad_j}$. Observe that p_j has the same form of p^n where $n = 2(k-1)+1 = 2k-1$ which completes the proof.

□

6 Conclusions

This paper tackles two fundamental issues in the study of multi-agent systems:

1. Providing solution concepts for multi-agent interactions with incomplete information, where exact probabilistic information is unavailable.

2. Incorporating uncertainty about the number of participants into the context of resource selection systems, one of the most fundamental settings in multi-agent systems.

Our results turned out to be both general and illuminating. On one hand, we provided general existence results for safety-level/competitive ratio equilibria. On the other hand, we were able to apply them to the context of resource selection games. Most importantly, our study led to highly surprising results, showing the positive effect of ignorance about the number of participants on the agents' costs and the system surplus.

Together, our work provides a foundational rigorous study of congestion games with incomplete information, while presenting general tools for the analysis of multi-agent systems where exact probabilistic information is not available.

References

[1] M. Aghassi and D. Bertsimas. Robust Game Theory. 2004.

[2] I. Ashlagi, D. Monderer, and M. Tenneholtz. On the value of correlation. In *Proceedings of the 21th Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 34–41, Corvallis,Oregon, 2005. AUAI Press.

[3] A. Borodin and R. El-Yaniv. *On-Line Computation and Competitive Analysis.* Cambridge University Press, 1998.

[4] S.F. Cheng, D. M. Reeves, Y. Vorobeychik, and M.P. Wellman. Notes on Equilibria in Symmetric Games. Technical Report, University of Michigan, 2004.

[5] R. Engelberg and J. Naor. Equilibria in Online Games. Unpublished manuscript, 2005.

[6] R. M. Harstad, J.H Kagel, and D. Levin. Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters*, 33:35–40, 1990.

[7] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.

[8] R. Holzman and N. Law-yone (Lev-tov). Strong equilibrium in congestion games. *Games and Economic Behavior*, 21:85–101, 1997.

[9] R. Holzman and D. Monderer. Characterization of Ex Post Equilibrium in the VCG Combinatorial Auctions. *Games and Economic Behavior*, 47:87–103, 2004.

[10] N. Hyafil and C. Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI-04)*, pages 268–277, Arlington, Virginia, 2004. AUAI Press.

[11] H. Konishi, M. Le Breton, and S. Weber. Equilibrium in a Model with Partial Rivalry. *Journal of Economic Theory*, 72:225–237, 1997.

[12] E. Koutsoupias and C. Papadimitriou. Worst-Case Equilibria. In *In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[13] D. Levin and E. Ozdenoren. Auctions with Uncertain Number of Bidders. To appear in 'Journal of Economic Literature', 2004.

[14] S. Matthews. Comparing auctions for risk averse buyers: a buyers point of view. *Econometrica*, 55:633–646, 1987.

[15] R. P. McAfee and J.McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43:1–19, 1987.

[16] I. Milchtaich. Congestion Games with Player-Specific Payoff Functions. *Games and Economic Behavior*, 13:111–124, 1996.

[17] P. Milgrom and R. Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50:1089–1122, 1982.

[18] D. Monderer and L.S. Shapley. Potential Games. *Games and Economic Behavior*, 14:124–143, 1996.

[19] R. B. Myerson. Population uncertainty and Poisson games. *International Journal of Game Theory*, 27:375–39, 1998a.

[20] R. B. Myerson. Extended Poisson games and the Condorcet jury theorem. *Games and Economic Behavior*, 25:111–131, 1998b.

[21] R. B. Myerson. Large Poisson games. *Journal of Economic Theory*, 94:7–45, 2000.

[22] T. Quint and M. Shubic. A Model of Migration. working paper. Cowles Foundation, Yale University, 1994.

[23] R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[24] Y. Shoham and M. Tennenholtz. Artificial Social Systems. *Computers and Artificial Intelligence*, 14:533–562, 1995.

[25] Y. Shoham and M. Tennenholtz. On Social Laws for Artificial Agent Societies: Off-Line Design. *Artificial Intelligence*, 73:231–252, 1995.

[26] Moshe Tennenholtz. *Efficient Representation and Reasoning in Multi-Agent Systems*. PhD Thesis, Weizmann Institute of Science, Israel, 1991.

[27] M. Voorneveld, P. Borm, F. van Megen, S. Tijs, and G. Facchini. Congestion Games and potentials reconsidered. *International Game Theory Review*, 1:283–299, 1999.