The Value of Ignorance about the Number of Players

Noga Alon, Reshef Meir and Moshe Tennenholtz
Microsoft Research
nogaa@tau.ac.il, reshef.meir@mail.huji.ac.il, moshet@microsoft.com

Introduction

In the game known as *the El-Farol bar*, each player decides whether to come to the bar or stay at home (Arthur 1994). Normally, every player prefers to enjoy the party, however if the bar becomes congested beyond its capacity, the police closes the place, inflicting severe trauma on all party-goers.

A variant of the problem (which is perhaps more realistic) is where the discomfort of the party-goers increases as the bar becomes more congested, even before capacity is exhausted. If the number of players is known, then in equilibrium the bar will always be packed up to the point of full capacity (just enough to avoid police intervention), or until discomfort is sufficiently high so that players are indifferent between partying and staying home.

A possible solution is to hide the exact number of players in the game. Facing uncertainty, some players may prefer to stay home so as to avoid the chance of exceeding capacity. Consequently, the bar will only become reasonably congested, so that at least those who arrive can enjoy the party.

Minority games applied the El Farol Bar as a metaphor for various economic situations (Challet, Marsili, and Zhang 2001). Our variant above and the *Kolkata Paise Restaurant problem* (Chakrabarti 2007), where players choose from multiple restaurants, are useful analogies to real problems like congestion of roads and of service providers—typically modeled as *congestion games* (Rosenthal 1973).

Our starting point is the model of Meir et al. (2012), where agents have uncertainty over the actual number of participants in a congestion game, formalized as a prior distribution over subsets of players. In this work we consider the idea of *partial information revelation* (signaling), for such games where the number of participants is unknown.

We expose a counter-intuitive phenomenon, where hiding the number of players may result in a *significant improvement* in welfare. We construct several examples of congestion games that demonstrate such improvement when hiding all or some information. We complement our results by proving that by hiding information the welfare can increase by a factor of at most n (the number of agents). As our examples demonstrate improvement by a similar factor, n is tight bound (and almost tight under i.i.d. participation).

Preliminaries

Let $[n] = \{0, 1, 2, \ldots, n\}$. A *Resource Selection Game (RSG)* is defined by a set of n agents N, and a set of resources F, each coupled with a cost function $c_j : [n] \to \mathbb{R}_+$. We denote the costs of resource $x \in \mathcal{F}$ by a cost vector $c_x = (c_x(1), c_x(2), \ldots, c_x(n))$. Each agent has a set of allowed strategies $\mathcal{F}_i \subseteq \mathcal{F}$. A *strategy profile* is a vector of strategies $\mathbf{A} = (A_1, \ldots, A_n)$, where $A_i \in \mathcal{F}_i$. For every profile $\mathbf{A}$ we denote by $N_x = \{i \in N : A_i = x\}$ the set of agents that selected resource x, and $n_x = |N_x|$.

The cost (negative utility) of agent i in profile $\mathbf{A}$ is $cost_i(G, \mathbf{A}) = c_x(n_x)$, where $x = A_i$. The *social cost* (or total cost) of a profile $\mathbf{A}$ is: $sc(G, \mathbf{A}) = \sum_{i=1}^{n} cost_i(G, \mathbf{A}) = \sum_{x \in \mathcal{F}} n_x c_x(n_x)$.

All the definitions in the paper can be naturally extended to general congestion games, where each agent can select a subset of resources rather than a single resource.

A profile $\mathbf{A}$ in G is a (pure) *Nash equilibrium* (PNE) if no agent can gain by departing from $\mathbf{A}$. All congestion games are potential games, and thus admit a pure Nash equilibrium (Rosenthal 1973). In this work we restrict our attention to pure Nash equilibria.

Uncertainty over participants. Following Meir et al. (2012), we extend an RSG G with participation probabilities for every agent, which may be correlated in general. We have a vector $\mathbf{p} \in \Delta(2^N)$, s.t. $p(S)$ is the probability that exactly the set S of agents participate.

Thus in the modified game $G^{\mathbf{p}}$, the cost of using resource x to agent j is $c^{\mathbf{p}}_{j,x}(N_x) = \sum_{R \subseteq N : j \in R} p(R) c_x(|R \cap N_x|)$.

Meir at al. (2012) proved that for any congestion game G, and any vector $\mathbf{p}$, $G^{\mathbf{p}}$ has a weighted potential function, and thus has a pure Nash Equilibrium. Meir (2013) later showed that $G^{\mathbf{p}}$ is itself a congestion game.

Signaling. The conceptual contribution of the current work is the introduction of *signaling schemes*. We assume that there is a central authority aware of the actual number of participants, which has an opportunity to disclose fully or partially this information to the players.

Formally, a *realization* of $G^{\mathbf{p}}$ is a subset of participants $S \subseteq N$. A *truthful signaling scheme* maps every realization S of $G^{\mathbf{p}}$ to a signal $\mathcal{T}(S) \subseteq 2^N$, s.t. $S \in \mathcal{T}(S)$. That is, the possible sets of participants must include the real set. We assume that signals are *disjoint*, and thus partition 2^N into equivalence classes. Denote by $\hat{\mathcal{T}} = \{\mathcal{T}(S) : S \subseteq N\}$ the set of possible signals under scheme $\mathcal{T}$. The two extreme signaling schemes are the *complete information* scheme $\mathcal{T}^*$, where $\mathcal{T}^*(S) = \{S\}$ for all S; and the *no information* scheme $\mathcal{T}^0$, where $\mathcal{T}^0(S) = 2^N$ for all S.

A *threshold signaling scheme* is characterized by a number $t \leq n$. It is a binary signal, which only indicates whether there are at least t participants. Note that the no information scheme $\mathcal{T}^0$ is a threshold scheme for $t = 0$.

While the broadcasted signal does not change the game $G^{\mathbf{p}}$ (i.e. the costs), it may affect the behavior of the agents. We assume that agents are perfectly rational, and given a signal $\mathcal{T}(S)$ they are able to compute the posterior participation probabilities and play accordingly.

Example. Suppose that $n = 4, p = \frac{1}{2}$. With no information, the probabilities that there will be $0, 1, 2, 3$ and 4 active players are $1/16, 4/16, 6/16, 4/16, 1/16$, respectively. Consider the threshold signaling scheme $\mathcal{T}$ with $t = 2$. Now, suppose agents get the high signal $[\![|S| > 2]\!]$. Then they know that they are now playing the game $G^{\mathbf{p}'}$, where $p'(|S| = 3) = 4/5, p'(|S| = 4) = 1/5$ (and zero probability for smaller sets).

The PNEs in the game $G^{\mathbf{p}'}$ may be different than those in G^p, and also different for the low signal $[\![|S| \leq 2]\!]$.

Equilibrium and social costs under signaling. The game $(G, \mathbf{p}, \mathcal{T})$ is a Bayesian game, where w.p. $p(S)$ the set S is realized, and agents get the signal $\mathcal{T}(S)$. A pure strategy in this game is a mapping from any possible signal $T \in \hat{\mathcal{T}}$ to a pure action $A_i(T) \in \mathcal{F}_i$. Since $\mathcal{T}(S)$ is a public signal, a (pure) Nash equilibrium in $(G, \mathbf{p}, \mathcal{T})$ is a profile of strategies, where the action of each player depends on the broadcasted signal. Applying the fact that for every signal $T \in \hat{\mathcal{T}}$ a PNE exists, we can show:

Proposition 1 *For any game $(G, \mathbf{p}, \mathcal{T})$, there exists a PNE.*

The Value of Ignorance. Revealing full or partial information on the number (or identities) of participants may change the outcome of the game, and thus the welfare of players. In order to measure the effect of using a particular signaling scheme $\mathcal{T}$, we compare to the outcome under full information. The formal definition requires some lengthy notation and is omitted, as are most proofs.

Informally, the *Value of Ignorance* (VoI) in a game is the ratio between the worst welfare in PNE under $\mathcal{T}^*$ and the worst welfare in PNE under $\mathcal{T}^0$. We are interested in upper and lower bounds on the VoI, as well as bound on the value of hiding some of the information. The "ignorance" we measure can be thought of as a special case of *Bayesian ignorance* (Alon et al. 2012), dealing with agents' participation rather than their perceived costs.

Results

Revealing information typically increases welfare, and can do so by an arbitrary factor. While it is not very surprising that revealing information can improve welfare, we are more interested in the other direction, i.e. in cases where *hiding* information will benefit the players.

The power of hiding information

We start by setting an upper bound on the Value of Ignorance in RSG (the example above shows that the only lower bound is 0).

Proposition 2 *For any RSG G and any distribution $\mathbf{p}$, the value of (full or partial) ignorance is at most n.*

We next show that the bound in proposition 2 is tight, up to a constant factor.

Example: a modified El-Farol game. The example we construct is composed of two resources, named B (for *bar*) and H (for *home*), both with (weakly) increasing costs. Assume the costs for B are $c_B = (0, 0, \dots, 0, 1, M)$, (that is, $n - 2$ times 0, one 1, and then M), the costs for H are always 2. Assume that $p(N) = 1/n^3$ (i.e., all agents participate w.p. $1/n^3$), and with the remaining probability exactly one random agent does not participate, i.e. $p(N \setminus \{i\}) = (1 - p(N))/n$ for all $i \in N$.

Intuitively, 1 is the cost incurred on party-goers when the bar (B) it at its full capacity $n - 1$, whereas 2 is the "cost" of staying at home. $c_B(n) = M$ is the high penalty incurred by a police intervention when congestion exceeds capacity. We set M high enough so that players will avoid even the slightest chance of exceeding bar's capacity.

It can be easily verified that in the no information case, there is a unique PNE up to permutations of agents, where a single agent selects H and all other agents select B. Thus without information the expected social cost experienced by players is roughly 3. On the other hand, under full information, whenever there are $n - 1$ participants (i.e. almost always), they will all select B, yielding a social cost of $n - O(1)$. Thus, the value of ignorance is $\Theta(n)$.

Note that in the example above we allowed correlated participation probabilities. A a similar example (with a somewhat more complicated analysis) can be constructed under i.i.d. participation, yielding a VoI of $\Theta(n/\log n)$.

The power of partial revelation

Our last result states that there are scenarios where hiding all information is not advised, but sending a proper threshold signal allows the agents to reach an almost optimal outcome.

Proposition 3 *For every n, there is an RSG G with two resources and i.i.d. participation probability p, s.t.:*

(a) The value of ignorance is negligible (i.e. hiding all information is bad);

(b) There is a threshold scheme $\mathcal{T}$ that improves upon the full information scheme $\mathcal{T}^*$ by a factor of $\Omega(n/\log n)$.

Current research

We are exploring ways to extend our bounds to more general congestion games, as well as improving the bounds for restricted classes of cost functions.

The finding that partial information can increase welfare highlights the algorithmic question of *design*. That is, given a game G and participation probabilities $\mathbf{p}$, design a signaling scheme $\mathcal{T}$ that yields the highest welfare.

References

Alon, N.; Emek, Y.; Feldman, M.; and Tennenholtz, M. 2012. Bayesian ignorance. *Theoretical Computer Science* 452:1–11.

Arthur, W. B. 1994. Inductive reasoning and bounded rationality. *The American economic review* 84(2):406–411.

Chakrabarti, B. K. 2007. Kolkata restaurant problem as a generalised el farol bar problem. In *Econophysics of Markets and Business Networks*. Springer. 239–246.

Challet, D.; Marsili, M.; and Zhang, Y.-C. 2001. Minority games: interacting agents in financial markets. *Quantitative Finance* 1:168–176.

Meir, R.; Tennenholtz, M.; Bachrach, Y.; and Key, P. 2012. Congestion games with agent failures. In *Proceedings of 26th AAAI*, 1401–1407.

Meir, R. 2013. A note on potential functions, congestion games, and the shapley value. In *CoopMAS @ AAMAS'13*.

Rosenthal, R. 1973. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1):65–67.