A GAME BETWEEN TWO BAYESIANS FOR ESTIMATING THE MEAN OF A GAUSSIAN DISTRIBUTION

PIERCESARE SECCHI

Two Bayesian players are engaged in a multi-stage competition where the final goal for each of them is to estimate the mean μ of a Normal distribution $\mathcal{N}$ with variance equal to 1, minimizing the total costs due to sampling and variance of the posterior distribution of μ.

1. The game

The competition is played in stages. At every stage, each player can select one among two possible actions: he can either take a sample from $\mathcal{N}$ or he can estimate μ. If he chooses to sample, he pays a cost $c > 0$ and a sample of size one from $\mathcal{N}$ is generated and made public; i.e. both players observe it. If the player chooses to estimate μ, he computes the mean of his posterior distribution for μ and he pays a cost equal to the variance of the posterior distribution of μ. Whenever a player estimates μ, the game is over for him in the sense that from that stage on he cannot observe samples from $\mathcal{N}$ nor make new estimates for μ and his cost at every following stage of the game is 0. Both players assume that successive observations sampled from $\mathcal{N}$ are conditionally independent given μ.

Everything stated in the previous paragraph is known and agreed by the two players; moreover each player knows his prior distribution for μ as well as the prior of his opponent. For the sake of avoiding philosophical intricacies, I will assume that both players have the same prior distribution for μ and this is a Normal with mean μ_0 and variance σ_0^2.

Let m be the strategy that prescribes to a player to take a sample in the first m stages of the competition and then to estimate μ at the next stage. Set $\mathcal{M} = \{m : m = 0, 1, 2, ...\}$: in fact, the strategy m is identified with the integer m defining it. I am constraining the two players to use strategies in $\mathcal{M}$ for controlling the competition. When $(n, m) \in \mathcal{M} \times \mathcal{M}$ is the profile of strategies used by the players, i.e. player 1 uses strategy n while player 2 uses strategy m, the total cost for player 1 is

$$L_1(n,m) = nc + \frac{\sigma_0^2}{1+(n+n\wedge m)\sigma_0^2}$$

where $n \wedge m$ indicates the minimum between the integers n and m. The total cost for player 2 is $L_2(n,m) = L_1(m,n)$. A Nash equilibrium (in pure strategies) for the game is a profile $(n^*,m^*) \in \mathcal{M} \times \mathcal{M}$ such that

$$L_1(n^*,m^*) \leq L_1(n,m^*)$$

for all strategies $n \in \mathcal{M}$ of player 1, while

$$L_2(n^*,m^*) \leq L_2(n^*,m)$$

for all strategies $m \in \mathcal{M}$ of player 2. The aim of the following section is to characterize the set of Nash equilibria for the game described above.

2. Nash equilibria

For $i = 1,2$, set g_i to be the best response map for player i; hence, for $n,m \in \mathcal{M}$,

$$g_1(m) = \{n^* \in \mathcal{M} : L_1(n^*,m) \leq L_1(t,m) \text{ for all } \; t \in \mathcal{M}\}$$

and

$$g_2(n) = \{m^* \in \mathcal{M} : L_2(n,m^*) \leq L_2(n,t) \text{ for all } t \in \mathcal{M}\}.$$

In order to describe the maps g_i, note that there exist a unique integer $\bar{n} \geq 0$ such that, for all $n \leq \overline{n}$,

$$L_1(n,n) \geq L_1(\bar{n},\bar{n})$$

while, for all $n > \overline{n}$,

$$L_1(n,n) > L_1(\bar{n},\bar{n}).$$

Let

$$\underline{n} = \lceil \frac{1}{2\sigma_0^2}(\frac{\sigma_0^2}{\sqrt{c}} - 1)\rceil \wedge \overline{n}.$$

where, for any real number x, $\lceil x \rceil$ represents the smallest integer larger than or equal to x. Now observe that, for $m \in \mathcal{M}$, $g_1(m) = g_2(m)$; moreover a few computations show that:

(i) if $m \geq \overline{n}$,

$$g_1(m) = \begin{cases} \{\overline{n}, \overline{n}-1\}, & \text{when } L_1(\bar{n}-1,\bar{n}-1) = L_1(\bar{n},\bar{n}), \\ \\ \{\overline{n}\}, & \text{otherwise;} \end{cases}$$

(ii) if $\underline{n} \leq m < \overline{n}$, $g_1(m) = \{m\}$;

(iii) if $m < \underline{n}$ and $n \in g_1(m)$, then $n > m$.

Indicate with $\mathcal{G}_i$ the graph of the map g_i; hence

$$\mathcal{G}_1 = \{(n,m) \in \mathcal{M} \times \mathcal{M} : n \in g_1(m)\}$$

and

$$\mathcal{G}_2 = \{(n,m) \in \mathcal{M} \times \mathcal{M} : m \in g_2(n)\}.$$

A profile $(n,m) \in \mathcal{M} \times \mathcal{M}$ is a Nash equilibrium for the game if and only if $(n,m) \in \mathcal{G}_1 \bigcap \mathcal{G}_2$. Because of (i)-(iii), the set of Nash equilbria (in pure strategies) is therefore:

$$N_1 = \{(n,n) \in \mathcal{M} \times \mathcal{M} : \underline{n} \leq n \leq \overline{n}\}.$$

Note that these equilibria are Pareto-ordered in the sense that, if (n,n) and (m,m) are equilibria in N_1 and $m > n$, then $L_1(n,n) = L_2(n,n) \geq L_1(m,m) = L_2(m,m)$. Hence both players prefer the equilibrium $(\overline{n}, \overline{n})$.

Example 1. Take $\sigma_0^2 = 1$ and $c = 1/8$. Then $N_1 = \{(1,1),(2,2)\}$ and both players prefer the equilibrium $(2,2)$ to the equilibrium $(1,1)$.

It is interesting to observe that if there was only one player with prior distribution for μ equal to a Normal with mean μ_0 and variance σ_0^2, the optimum sample size for estimating μ minimizing the player's total cost

$$L(n) = cn + \frac{\sigma_0^2}{1 + n\sigma_0^2}$$

would be an integer n^* close to the real number

$$\frac{1}{\sigma_0^2}(\frac{\sigma_0^2}{\sqrt{c}} - 1).$$

The total cost $L(n^*)$ is approximately

$$\frac{c}{\sigma_0^2}(\frac{\sigma_0^2}{\sqrt{c}} - 1) + \sqrt{c}$$

and this quantity is greater than or equal to $L_1(n,n)$ for all Nash equlibria $(n,n) \in N_1$.

3. The "pass" option

A different version of the previous game is obtained by allowing each player to use the action "pass", in addition to the actions "sample" and "estimate". When, at a given stage of the game, a player who is not yet out of the game, i.e. who has not yet estimated μ, chooses the "pass" action, he pays rc, with $r \in [0,1]$, and moves to the next stage, getting however the right to observe the sample possibly taken and paid by the other player, since all observations generated from $\mathcal{N}$ are public.

If $r = 1$, the profiles in N_1 are Nash equilibria also in the new game since "pass" can never be more advantageous than "sample" when the two actions have the same cost. However, for $r < 1$ it is not always the case that a profile in N_1 is an equilibrium in the new game.

Example 2. Again, assume that $\sigma_0^2 = 1$ and $c = 1/8$. For $1/3 \leq r < 3/5$, $(2, 2)$ is not a Nash equilibrium; in fact, if $m = 2$ is the strategy of player 2 (sample $\mathcal{N}$ for the first two stages of the game, then estimate), the total cost $L_1(n, 2)$ of player 1 is minimized by taking a sample in the first stage of the game, passing in the second stage and estimating μ in the third. However, when player 1 selects this last strategy, then player 2's cost is minimized by choosing the strategy $m = 1$. Indeed, the profile $(1, 1)$ where both players take a sample at the first stage of the game and estimate μ at the second stage is a Nash equilibrium in the new game for all $r \in [0, 1]$. Note that the "pass" option is not used by any player in the profile $(1, 1)$; nevertheless, introducing this possibility with a sufficiently low cost has ruined the equilibrium $(2, 2)$ even though this profile still entails a lower total cost than $(1, 1)$ for both players.

The profile $(2, 2)$ is not an equilibrium even for $r < 1/3$: interestingly, in this case two 'asymmetric' Nash equilibria in $\mathcal{M} \times \mathcal{M}$ appear along with $(1, 1)$. They are $(0, 2)$ and $(2, 0)$.

4. A concluding remark

The games described in these pages are examples of negative rewards stochastic games with two players and unbounded payoffs. When the state space of the game is countable, the action sets for the players are finite and payoffs are bounded, the existence of Nash equlibria in mixed strategies for a negative rewards stochastic game with n players is proved in Secchi and Sudderth (2002).

References

Secchi, P. and Sudderth, W.D. (2002). N-Person Stochastic Games with Upper Semi-Continuous Payoffs, *International Journal of Game Theory*, 30 (4), 453-478.

E-mail address: secchi@mate.polimi.it

Dipartimento di Matematica "F. Brioschi", Politecnico di Milano, Piazza Leonardo da Vinci, 32, 20123 Milano - Italy