10 Contests: a mixed strategy equilibrium

10a The framework

As far as I know, all mathematical methods used till now in the theory of economic games depend heavily on monotonicity of strategies (assumed or proven). The inescapable end of monotone equilibria (see Sect. 9) implies a pressing need for different mathematical methods. In the face of the difficulty, it could be wise first to retreat to simpler games, as simple as possible among games with many players, a single winner, correlated signals and entry cost. Participation becomes more relevant than bidding. The following *contest game* is probably the best choice.

We consider a symmetric game of n players. The signal space is $\mathbb{R}$ (as before), but the action space consists of only two points $0, 1$; the action 0 means quitting, the action 1 means participating. If no one participates, all get nothing. Otherwise each participant pays an entry cost c, and the participant possessing the highest signal is the winner, he gets a price C.[1] Parameters c, C (and n) are common knowledge. Naturally, $0 < c < C < \infty$. Thus, the profit is $\mathbf{\Pi} = \mathbf{G} - \mathbf{L}$,

(10a1)
$$\mathbf{G}(a_1, s_1; \dots; a_n, s_n) = \begin{cases} C & \text{if } a_1 = 1 \text{ and } s_1 = \max_{a_k=1} s_k, \\ 0 & \text{otherwise;} \end{cases}$$
$$\mathbf{L}(a_1, s_1; \dots; a_n, s_n) = ca_1 = \begin{cases} c & \text{if } a_1 = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Correlated signals $S_1, \dots, S_n$ are described (as before, see 8a) by Θ, P_Θ and $(P_{S|\theta})$. Each signal has its marginal (one-dimensional) distribution $P_S = \int P_{S|\theta} \, dP_\Theta$.

A strategy, being a probability distribution on two lines $\mathcal{S} \times \mathcal{A} = \mathbb{R} \times \{0, 1\} = (\mathbb{R} \times \{0\}) \cup (\mathbb{R} \times \{1\})$ with the given projection P_S onto $\mathbb{R}$, may be described by a function $\mathbf{p}$ such that

$$\mathbf{p}(s) = \mathbb{P}\left(A = 1 \,\middle|\, S = s \right) ;$$

that is, $\mathbf{p}(s)$ is the participation probability for a player possessing signal s. A pure strategy corresponds to the case where $\mathbf{p}(\cdot)$ takes on two values only, 0 and 1. In general, $0 \le \mathbf{p}(s) \le 1$. Nothing like monotonicity, continuity etc. is assumed; $\mathbf{p}$ is just a function, defined almost everywhere w.r.t. P_S and taking on values on $[0, 1]$.[2] A monotone (increasing) strategy is a threshold strategy: $\mathbf{p}(s) = \mathbf{1}_{(t,\infty)}(s)$.

We restrict ourselves to *symmetric* equilibria.

[1] Ties will be excluded by assuming nonatomic distribution of signals.

[2] Of course, the function must be measurable w.r.t. P_S. Functions coinciding P_S-almost everywhere describe the same strategy.

10b A simple counterexample

The winning probability $\mathbf{p}^{\text{win}}(s) = \mathbb{P}\,\big(\,G_1 = C \,\big|\, A_1 = 1\,\big)$ depends on the strategy (of others); in any case,

$$\mathbf{p}^{\text{win}}(s) \ge \underbrace{\mathbb{P}\,\big(\,S_1 = \max(S_1,\dots,S_n) \,\big|\, S_1 = s\,\big)}_{h(s)}\,; \tag{10b1}$$

the maximal signal need not participate, but if it participates, it wins.

- If $\mathbf{p}^{\text{win}}(s) < c/C$ then the action 0 is optimal for s, while the action 1 is not.
- If $\mathbf{p}^{\text{win}}(s) = c/C$ then both actions 0 and 1 are optimal for s.
- If $\mathbf{p}^{\text{win}}(s) > c/C$ then the action 1 is optimal for s, while the action 0 is not.

Therefore in equilibrium $p(S) = 1$ whenever $h(S) > c/C$. The *hope function* $h(\cdot)$ is determined by the signal distribution only,

$$h(s) = \mathbb{E}\,\big(\mathbb{P}\,\big(\,S \le s \,\big|\, \Theta\,\big)\,\big)^n = \int_\Theta F^n_{S|\theta}(s)\,dP_\Theta(\theta)\,.$$

It is natural to think that $h(s)$ increases continuously from 0 to 1. However, in general it can be much worse.

10b2. Example. Let Θ run over integers: $\mathbb{P}\,\big(\,\Theta = m\,\big) > 0$ for all $m = \dots,-2,-1,0,1,2,\dots$ and $\sum_{m=-\infty}^{\infty} \mathbb{P}\,\big(\,\Theta = m\,\big) = 1$. Further, let $P_{S|\Theta=m} = \mathrm{U}(m, m+1)$ for each integer m. Then $h(\cdot)$ is a discontinuous *periodic* function. Namely, $h(m+\alpha) = \alpha^{n-1}$ for all $\alpha \in (0,1)$ and integer m.

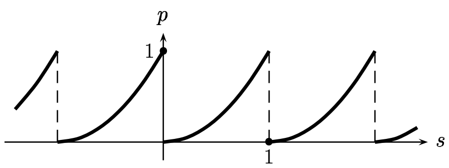

10b3. Exercise. There is no monotone equilibrium for such a game (even for $m = 2$). Prove it.

In fact, such signals are affiliated, which appears to be rather irrelevant to participation.

10c Normally correlated signals

Recall, among games with many players, a single winner, correlated signals and entry cost, we want to choose the simplest game. A contest game should fit, unless the distribution of signals is bizarre (which happens in 10b2). There is a wide choice of specific one-dimensional distributions (uniform, normal, exponential, gamma, beta etc), however, the choice of specific multidimensional distributions is scanty.[3] The multinormal distribution

[3]Do you know of any specific joint distribution of $S_1,\dots,S_n$ such that $S_1 \sim \mathrm{U}(0,1)$, ... , $S_n \sim \mathrm{U}(0,1)$? Two evident examples are (a) independent $S_1,\dots,S_n$, and (b) $S_1 = \dots = S_n$, but these two are degenerate. What about a non-degenerate correlation between these $\mathrm{U}(0,1)$? Maybe you can *invent* such a distribution rather than find a *canonical* one.

is distinguished. It is uniquely determined by expectations, variances and correlations. Expectations $\mathbb{E} S_1, \dots, \mathbb{E} S_n$ are equal by symmetry; the same for variances and correlations. Thus, irrespective of n, we have only 3 parameters μ, σ, ρ: $\mathbb{E} S_k = \mu$, $\operatorname{Var}(S_k) = \sigma^2$, $\operatorname{Cov}(S_k, S_l) = \rho\sigma^2$ for $k \ne l$. Moreover, the contest game is insensitive to a linear transformation of all signals, thus we may choose $\mu = 0$, $\sigma = 1$; now ρ is the sole parameter. So,

$$\text{(10c1)} \qquad \begin{gathered} \mathbb{E} S_k = 0\,; \quad \mathbb{E} S_k^2 = 1\,; \quad \mathbb{E}(S_k S_l) = \rho \quad \text{for } k \ne l\,; \\ S_1, \dots, S_n \text{ are multinormal.} \end{gathered}$$

A representation via θ is well-known, canonical and simple:

$$\text{(10c2)} \qquad \begin{aligned} & S_k = \sqrt{\rho}\,\Theta + \sqrt{1-\rho}\,\xi_k\,, \\ & \Theta \sim \mathrm{N}(0,1)\,, \\ & \xi_1 \sim \mathrm{N}(0,1), \dots, \xi_n \sim \mathrm{N}(0,1)\,, \\ & \Theta, \xi_1, \dots, \xi_n \text{ are independent.} \end{aligned}$$

The game is completely specified by (10a1), (10c2) provided that n, c, C, ρ are given. However, we introduce also a reserve level r.[4] A player possessing a signal $S < r$ is not allowed to participate. Accordingly, a strategy may be described by a function

$$\mathbf{p} : (r, \infty) \to [0,1]\,.$$

10d Equilibrium

Most results of this subsection are reported without proofs. Proofs will appear elsewhere.[5]

10d1. Lemma. If $\mathbf{p} : (r, \infty) \to [0,1]$ supports a symmetric equilibrium, and $r_1 \in (r, \infty)$, then the restriction of $\mathbf{p}$ onto (r_1, ∞) supports a symmetric equilibrium for the game with the reserve level r_1.

You see, every player depends on higher competitors. (For an auction with no entry cost the situation is just opposite.)

The hope function may be written as

$$h(s) = \int \Phi^{n-1}\big(\sqrt{1-\rho}\,s - \sqrt{\rho}\,u\big)\, d\Phi(u)\,,$$

where Φ is the standard normal cumulative distribution function, $\Phi(x) = \int_{-\infty}^{x} e^{-u^2/2}\, du$. Thus, h is a strictly increasing continuous function on $\mathbb{R}$, and $h(-\infty) = 0$, $h(+\infty) = 1$. It follows that the equation

$$\text{(10d2)} \qquad h(s^{\mathrm{high}}) = \frac{c}{C}$$

has exactly one solution $s^{\mathrm{high}} \in \mathbb{R}$.

[4] You may think that r makes the game more complicated, but we'll see soon that r helps to understand the situation.

[5] Landsberger and Tsirelson, 2001 (in preparation).

10d3. Lemma. (a) If $r \ge s^{\text{high}}$ then there exists one and only one symmetric equilibrium, namely, 'never quit', that is, $\mathbf{p}(s) = 1$ for all $s \in (r, \infty)$.

(b) If $r < s^{\text{high}}$ then the 'never quit' strategy does not support a symmetric equilibrium.

10d4. Lemma. If $r < s^{\text{high}}$ and $\mathbf{p} : (r, \infty) \to [0, 1]$ supports a symmetric equilibrium then

$$\mathbf{p}(s^{\text{high}}-) = \rho$$

(maybe after correcting $\mathbf{p}$ on a negligible set).

Therefore, pure strategy equilibria are excluded for $r < s^{\text{high}}$. Note the (rather unexpected) jump of $\mathbf{p}$ at s^{high}.

10d5. Theorem. There exists $s^{\text{low}} \in [-\infty, s^{\text{high}})$ such that

(a) there exists a function $\mathbf{p} : (s^{\text{low}}, \infty) \to [0, 1]$ such that $\mathbf{p}(\cdot)$ is continuous on $(s^{\text{low}}, s^{\text{high}}]$, and $0 < \mathbf{p}(s) < 1$ for all $s \in (s^{\text{low}}, s^{\text{high}}]$, and $\mathbf{p}(s^{\text{high}}) = \rho$, and $\mathbf{p}(s) = 1$ for all $s \in (s^{\text{high}}, \infty)$, and for every $r \in [s^{\text{low}}, \infty)$ the restriction of $\mathbf{p}(\cdot)$ onto (r, ∞) is an equilibrium strategy;

(b) If $s^{\text{low}} > -\infty$ then the limit $\mathbf{p}(s^{\text{low}}+)$ exists and is equal to 0 or 1;

(c) for every $r \in [s^{\text{low}}, \infty)$, every equilibrium strategy on (r, ∞) is equal to $\mathbf{p}(\cdot)$ almost everywhere on (r, ∞).

10e Some numerics and asymptotics

The equilibrium function $\mathbf{p}$ of Theorem 10d5 is in fact a solution of a nonlinear integral equation. Very probably it has no explicit analytic expression. However, it can be computed numerically.

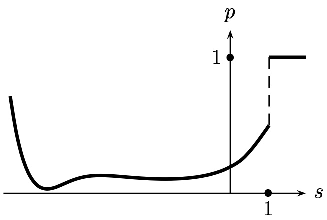

ρ=0.5, n=20, c/C=0.1

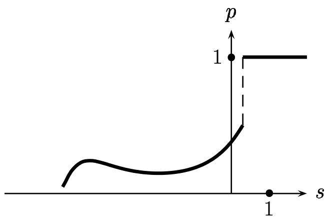

ρ=0.5, n=20, c/C=0.03

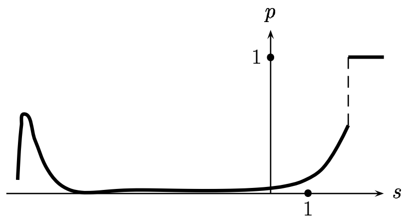

ρ=0.5, n=100, c/C=0.1

Wonders never cease! In addition to the unexpected jump at s^{high} we observe unexpected oscillations on the lower tail. Though, the oscillations are of little economic interest, since the tail is of probability much smaller than $1/n$.

Using some ideas of the theory of large deviations[6] one can get asymptotic formulas for

$$n \to \infty\,, \quad \frac{c}{C} \to 0\,, \quad \frac{c}{C} = \frac{\text{const}}{n^a}\,.$$

Namely,

$$\mathbf{p}(s) \approx \frac{1}{n^{1-(1-\rho)A^2}} \qquad \text{where } A = \tilde{s} + \sqrt{\frac{a\rho}{1-\rho}}\,, \ \tilde{s} = \frac{s}{\sqrt{2 \ln n}}\,.$$

[6]A branch of probability theory.

A better approximation is more complicated:

$$\mathbf{p}(s) \approx \frac{1}{n^{1-(1-\rho)A^2} \big(\sqrt{\ln n}\big)^{B\tilde{s}+\rho-1}} \qquad \text{where } B = \sqrt{\frac{\rho(1-\rho)}{a}}\,.$$

A really good approximation is rather frightening,

$$\mathbf{p}(s) = (1+o(1)) \cdot \frac{1}{n^{1-(1-\rho)A^2} \big(\sqrt{\ln n}\big)^{B\tilde{s}+\rho-1}} \cdot 2\rho\sqrt{\pi(1-\rho)}A \cdot \left(\frac{1}{2\sqrt{\pi a}} \Gamma\Big(1 + \frac{1}{B\tilde{s}+\rho}\Big) \right)^{B\tilde{s}+\rho}$$

(here Γ is the gamma function), but it is much better than nothing, given that we have no explicit expression for $\mathbf{p}(s)$. Compare exact and approximate strategies:

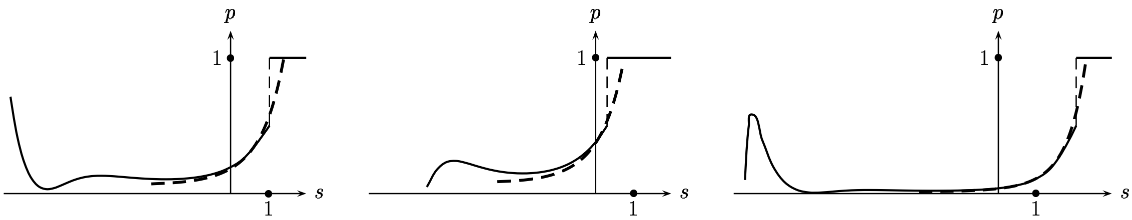

ρ=0.5, n=20, c/C=0.1 ρ=0.5, n=20, c/C=0.03 ρ=0.5, n=100, c/C=0.1

The course is now finished,
but the story is only starting ...