Entropy Methods for Communication and Coordination in Games

Olivier Gossner

PSE and LSE

GAMES, Udine, Sept.2009

Information Theory and Game Theory

Information Theory (Shannon 1948)

The notion of entropy allows to measure:

- The randomness of a random variable
- The information the observation of a r.v. provides on another r.v.

Questions from Game Theory

- How can players communicate effectively?
- How can they coordinate effectively?

Connexions

- Natural applications of information theory to game theory
- New information theory questions arising from game theory

Roadmap

1. Entropy: a crash course or a quick reminder
2. Optimal use of communication resources
3. Secret coordination and inducing beliefs
4. Open questions

Entropy: Measure of randomness

Definition

X, Y pair of discrete random variables

$$H(X) = -\sum_x P(x)\log P(x), \quad \log = \log_2, 0\log 0 = 0$$

$$H(X|Y) = -\sum_y P(y)\sum_x P(x|y)\log P(x|y)$$

$H(X)$ and $H(X|Y)$ depend on X, Y through their distributions only.

Additivity

$$H(X, Y) = H(Y) + H(X|Y)$$

Entropy on $[0, 1]$

X is a binary random variable, $X \sim (p, 1-p)$.

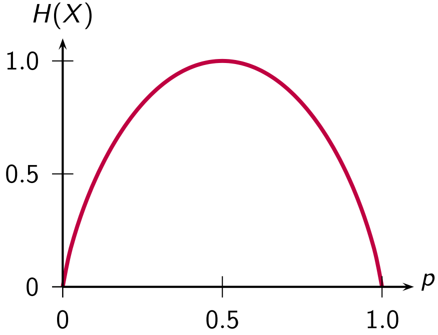

Entropy and counting

Let $X_1, \dots X_n$ i.i.d.$\sim \mu$, with values in a finite set X.

Asymptotic Equipartition Property

$\tilde{x} = (x_1, \dots, x_n)$ is μ-typical if its empirical frequency is (close to) μ.

1. For n large, the set of typical sequences has probability close to 1.
2. For a typical sequence $\tilde{x}$

$$\log P(\tilde{x}) = \sum_k \log \mu(x_k) \sim -nH(\mu)$$

Consequence

There are (approximately) $2^{nH(\mu)}$ typical sequences. $\tilde{x}$ can be described by a binary string of length $nH(\mu)$.

Roadmap

Puzzle

BRAIN TEASERS

Challenging

main

Two (2) friends enter our infamous casino to play roulette. Being conservative they play only the colors red or black. Seeing their behavior the casino proposes a more interesting game to them:

The casino has a very long printout of the roulette results for red/black over the last year. The numbers 0, 00 etc. have been deleted from the printout. One of the friends will get a copy of the printout but cannot tell the other verbally or by gestures the contents of the list. Each of them bets separately on the color at each step of the printout. They win a round only if both bet correctly. After every round the croupier opens both envelopes for everyone to see, reveals the casino printout of this round and declares if they won or the casino won. The bet is 1:2. For each win (both bet correctly) they get $1; for each incorrect bet (one or the other or both bet incorrectly) they lose $2.
Will you accept the casino's offer?

The answer is yes, of course, although it does not make sense at first. **However, we leave you with the task of devising a strategy.**

Model (GHN, Econ 06)

3 players, with action sets I, J, K.

Player 1 is nature

Chooses $x_1, x_2, \ldots$ i.i.d. $\sim \mu \in \Delta(I)$

Player 2 is the prophet

Knows the states of nature in advance, strategy $y = (y_t)_t$

$$y_t : I^\infty \times J^{t-1} \times K^{t-1} \to J$$

Player 3 is the follower

Observes past moves, strategy $z = (z_t)_t$

$$z_t : I^{t-1} \times J^{t-1} \times K^{t-1} \to K$$

Induced distributions

The strategies μ, y, z induce P on $(I \times J \times K)^\infty$.

Expected Empirical Distribution

With $P^t \in \Delta(I \times J \times K)$ the t-th stage marginal of P, let

$$Q^n = \frac{1}{n} \sum_1^n P^t$$

be the expected empirical distribution up to stage t.

What are the possible Q^n?

Information constraint

For $Q \in \Delta(I \times J \times K)$, denote

$$H_Q(i,j|k) = H(X,Y|Z) \text{ where } (X,Y,Z) \sim Q$$

Theorem 1

For every μ, y, z, n:

1. The marginal of Q^n on I is μ.
2. Information constraint

$$H_{Q^n}(i,j \mid k) \geq H(\mu)$$

The amount of information per stage received by the follower given his own action is at least $H(\mu)$.

Proof of Theorem 1

$$\begin{aligned} H_{P^t}(i,j \mid k) &= H(x_t, y_t \mid z_t) = H(z_t, y_t, z_t \mid z_t) \\ &\geq H(x_t, x_t, z_t \mid \ldots x_{t-1}, y_{t-1}, z_{t-1}) \end{aligned}$$

By concavity and additivity:

$$\begin{aligned} H_{Q^n}(i,j \mid k) &\geq \frac{1}{n}\sum_{1}^{n} H_{P^t}(i,j \mid k) \\ &\geq \frac{1}{n}\sum_{1}^{n} H(x_t, y_t, z_t \mid \ldots x_{t-1}, y_{t-1}, z_{t-1}) \\ &= \frac{1}{n} H(x_1, y_1, z_1 \ldots x_n, y_n, z_n) \\ &\geq H(\mu) \end{aligned}$$

Converse Result

How much is Theorem 1 saying about achievable distributions?

Theorem 2

For Q such that:

- The marginal of Q^n on I is μ,
- Q satisfies the IC,

there exists strategies of the follower and the prophet s.t. $Q^n \to Q$.

Given μ, the set of limit distributions Q that can be achieved by strategies of the prophet and the follower is the set of distributions having marginal μ on I that satisfy the information constraint.

Proof of Theorem 2

We construct strategies on blocks of length N: in each block, the prophet tells the follower which actions to play during the next block.

Action plans

The set of action plans is $AP^N \subset I^N$ such that, for each μ-typical sequence $\tilde{x} \in I^N$ there exists $\tilde{z} \in AP^N$ such that $(\tilde{x}, \tilde{z})$ is Q-typical.

There exist a set of action plans of size $\sim 2^{N(H_Q(i)-H_Q(i|k))}$.
Proof by the random method (Alon Spencer).

Message sets

For each typical sequence in $(\tilde{x}, \tilde{z})$, the prophet chooses a sequence of actions $\tilde{y}$ s.t.$(\tilde{x}, \tilde{y}, \tilde{z})$ is Q-typical.

The size of a message set is $\sim 2^{NH_Q(j|i,k)}$.
Proof by counting.

Theorem 2: Completion of the Proof

The construction of strategies for the prophet and the follower is possible if there are as least as many messages as action plans:

$$\begin{aligned} H_Q(j|i,k) &\geq H_Q(i) - H_Q(i|k) \\ H_Q(i|k) + H_Q(j|i,k) &\geq H_Q(i) \\ H_Q(i,j|k) &\geq H_Q(i) \end{aligned}$$

That is, the Information Constraint.

Application: brain teaser

$I = J = K = \{0,1\}$, μ is uniform, $g(i,j,k) = 1_{i=j=k}$.

We seek

$$v^*(\mu) = \max_Q \mathbf{E}_Q g(i,j,k)$$

under the constraints

$$Q(i{=}0) = \frac{1}{2}$$

$$IC: \quad H_Q(i,j \mid k) \geq H(\mu)$$

Solution

The value $v^*(\frac{1}{2})$ is approximately 0.81. We also know the optimal Q, and corresponding strategies of the prophet and follower.

Extension: Malevolent nature

Assume that, when devising strategies, the team doesn't know the distribution of moves by nature, but considers the worst possible case. The maximal payoff the team can guarantee is:

$$\underline{v} = \max_{y,z} \min_{x} \lim \frac{1}{T} \sum_{t=1}^{T} g(x_t, y_t, z_t)$$

Theorem

$$\underline{v} = \min_{\mu} v^*(\mu)$$

The zero-sum game between the team and nature has a value, and nature has an optimal strategy which is i.i.d.

Roadmap

The Duenna Game

Two lovers (1 and 2) seek to meet without the lady's duenna (3).
B="garden bridge", C= "woodcutter's cottage".
Payoffs to 1 and 2 are (- payoff to 3)

	C	B
C	1	0
B	0	0

B

	C	B
C	0	0
B	0	1

C

The game is repeated. 1 and 2 observe all actions. Consider the cases

1. 3 observes all actions
2. 3 observes neither 1 nor 2's actions
3. 3 observes 1's actions, but not 2's
4. 3 observes all actions, but 1,2 observe a common additional signal at each stage

For notational simplicity we focus on the 4th case.

Model (GT, MOR 06, 08)

3 players

- Players 1 and 2 form a team, actions sets A^1, A^2.
- Players 3 plays a stage best response, she is an observer

Signals

At each stage t, players 1 and 2 observe the same signal $s_t \in S$, where $s_1 \dots s_t$ i.i.d.$\sim \mu$. Player 3 does not observe this signal.

(Behavioral) Strategy for $i = 1, 2$

Wlog. players 1 and 2 ignore 3's moves: $\sigma^i = (\sigma^i_t)_t$ where

$$\sigma^i_t \colon (A^1 \times A^2 \times S)^{t-1} \to \Delta(A^i)$$

Beliefs

Let P be the induced probability on $(S \times A^1 \times A^2)^\infty$. At each stage t, the belief of player 3 on player 1 and 2's actions is

$$p_t = P(a_t^1, a_t^2 | a_1^1, a_1^2, \dots a_{t-1}^1, a_{t-1}^2)$$

Question

Given a probability distribution Q on $\Delta(A_1 \times A_2)$, can we construct strategies of the team such that p_t is (most of the time, close to) Q? In this case, say that Q is implementable with μ.

Secret Entropy

Definition

The secret entropy of Q is

$$SH(Q) = \min_{k,a^1,a^2} H(a_1,a_2) - H(a_1,a_2|k)$$
$$= \min_{k,a^1,a^2} H(k) - H(k|a^1,a^2)$$

where the minimum is taken over all random variables k, a^1, a^2 such that a^1, a^2 are independent given k and $(a^1,a^2) \sim Q$.

$SH(Q)$ is the part of randomness in (a_1,a_2) that cannot be explained by independence. It is also the minimal information gained on an unobserved phenomenon explaining the correlation between a^1 and a^2 (Okham's razor).

- $SH(Q) = 0$ if and only if a^1, a^2 are independent.
- $SH(Q) = H(a)$ if $a_1 = a_2 = a$.

Characterization

Theorem

Q is implementable with μ if and only if

$$H(\mu) \geq SH(Q)$$

Proof:

- "Only if": Additivity, Concavity.
- "If": Coding, recycling.

Application: the Duenna game

Consider a game between a team (1,2) and 3, where the team observes at each stage a common signal of distribution μ. The payoff to the team is $g(a_1, a_2, a_3)$. We want to characterize

$$v = \max_{\sigma^1,\sigma^2} \min_{\sigma^3} \lim \frac{1}{T} \sum_{t=1}^{T} g(a_t^1, a_t^2, a_t^3)$$

Given a distribution Q of actions for $1, 2$, let

$$\pi(Q) = \min_{a_3} \mathbf{E}_Q g(a_1, a_2, a_3)$$

Theorem

$$v = \max_{p,Q,Q'} p\pi(Q) + (1-p)\pi(Q')$$

where

$$H(\mu) \geq pSH(Q) + (1-p)SH(Q')$$

Generalizations

The approach generalizes to the following situations:

1. Player 3 receives a signal on 1,2's actions, where this signal is known to player 1 and 2 and does not depend on 3's action (as in the initial examples)
2. The signal of player 3 is controlled by player 3's action. In this case, a "vector game" is played between the team and player 3 on the following dimensions
   - Payoffs, where the team uses optimally the past secret signals
   - Signals, as the amount of "secret correlation" to the team at each stage is controlled by both teams

   The problem is tackled using Blackwell's approachability theory.

Roadmap

The secret entropy

The secret entropy is an interesting measure of correlation (Statistics) between two (or more) random variables.

Computation

GLT (Math Prog B, 09) show how to compute $SH(Q)$ if A^1 and A^2 both have cardinality 2. The computation of $SH(Q)$ is an open problem for general sets A^1, A^2.

Extracting public information from correlated signals

Assume that players 1, 2 receive correlated signals $(s^1, s^2) \sim \mu$.

Open Questions

What are the implementable distributions with μ?

- Is Q is implementable if and only if $SH(\mu) \geq SH(Q)$?
- What is the maximum entropy of an implementable distribution Q for which $Q(a^1 = a^2) = 1$?
- Is an appropriate measure of the correlation provided by μ unidimensional, or multidimensional?

Conclusion: G and H

Applications to Game Theory

Entropy is a natural tool to study of communication and coordination, but also learning, etc.

New questions

New questions arise from these fields of applications, answering them potentially opens the scope of applications of Information Theory, including and beyond Game Theory.

Thank you!