Game Theoretic Probability Notation and Definitions

Monday 30$^{\text{th}}$ May, 2016

1 Protocols

We consider perfect information games between two players, Skeptic and World. World controls Forecaster (if present) and Reality. The available moves are specified by a *protocol* (that we vary). A specific protocol that we will consider in detail in the second lecture is

> Fair-coin Game
> $\mathcal{K}_0 = 1$.
> **for** $n = 1, 2, \ldots$ **do**
>   Skeptic chooses a number of tickets $M_n \in \mathbb{R}$.
>   Reality announces outcome $x_n \in \{-1, +1\}$.
> $\mathcal{K}_n = \mathcal{K}_{n-1} + M_n x_n$.
> **end for**

A more general protocol (with outcomes in the set $\mathcal{X}$) looks like this

> Generic Game
> $\mathcal{K}_0 = 1$.
> **for** $n = 1, 2, \ldots$ **do**
>   Forecaster announces gambles $\mathcal{G}_n \subseteq [\mathcal{X} \to \mathbb{R}]$.
>   Skeptic chooses a gamble $g_n \in \mathcal{G}_n$.
>   Reality announces outcome $x_n \in \mathcal{X}$.
> $\mathcal{K}_n = \mathcal{K}_{n-1} + g_n(x_n)$.
> **end for**

2 Concepts

It will be useful to think about Skeptic and World as alternating moves. To do so, we group the move for Reality and the subsequent move for Forecaster:

$$\underbrace{\text{Forecaster}}_{\text{World}}, \text{Skeptic}, \underbrace{\text{Reality}, \text{Forecaster}}_{\text{World}}, \text{Skeptic}, \underbrace{\text{Reality}, \text{Forecaster}}_{\text{World}}, \ldots$$

A *path* is a complete sequence of moves by World (finite or infinite). The *sample space* Ω is the set of possible paths. A finite prefix of a path is called a *situation*.

(If Forecaster is present, a situation alternates moves by Forecaster and Reality, and ends with one more Forecaster move). We denote the set of situations by $\Omega^\diamond$. The set $\Omega^\diamond$ is a tree (in the descriptive set theory sense), rooted at the empty situation $\Box$. This tree is called *finite horizon* if path length is bounded, *terminating* if all paths are finite and *infinite horizon* if all paths are infinite.

A subset $E \subseteq \Omega$ of the sample space is called an *event*. A function $\mathcal{S} : \Omega \to \mathbb{R}$ is called a *variable*. A function $\mathcal{S} : \Omega^\diamond \to \mathbb{R}$ is called a *process*. We may represent a process $\mathcal{S}$ by a sequence of variables $\mathcal{S}_0, \mathcal{S}_1, \dots$ as follows. Writing $\xi^n = \langle \xi_1, \dots, \xi_n \rangle$ for the n-length prefix of ξ, we take $\mathcal{S}_n(\xi) = \mathcal{S}(\xi^n)$ to be the value after n rounds.

Gambles A *gamble* is a payoff depending on World's next move. The protocol specifies in each situation the set of gambles available to Skeptic. We always assume this set is a convex cone (closed under nonnegative linear combinations). In a *symmetric* protocol the set of gambles is a linear space (closed under arbitrary linear combination).

Strategy and Capital A *strategy $\mathcal{P}$ for Skeptic* chooses a gamble in each situation. We write $\mathcal{K}^\mathcal{P}(t)$ for Skeptic's capital process in situation $t \in \Omega^\diamond$ when he follows $\mathcal{P}$ starting from capital 0. This defines $\mathcal{K}^\mathcal{P}(\xi)$ for finite paths $\xi \in \Omega$. For infinite paths $\xi \in \Omega$ we (cautiously) extend $\mathcal{K}^\mathcal{P}(\xi) = \liminf_{n\to\infty} \mathcal{K}^\mathcal{P}(\xi^n)$. The capital $\mathcal{K}^\mathcal{P}$ is both a process and a variable. The capital process following $\mathcal{P}$ starting from stake α is $\alpha + \mathcal{K}^\mathcal{P}$.

Pricing For a variable $x : \Omega \to \mathbb{R}$, we write

$$\mathcal{K}^\mathcal{P} \geq x \qquad \text{if} \qquad \mathcal{K}^\mathcal{P}(\xi) \geq x(\xi) \text{ for all paths } \xi \in \Omega.$$

The *upper price* and *lower price* for x are

$$\begin{aligned} \overline{\mathbb{E}}[x] &\coloneqq \inf\left\{\alpha \in \mathbb{R} \middle| \exists \mathcal{P} : \mathcal{K}^\mathcal{P} \geq x - \alpha\right\} \\ \underline{\mathbb{E}}[x] &\coloneqq \sup\left\{\alpha \in \mathbb{R} \middle| \exists \mathcal{P} : \mathcal{K}^\mathcal{P} \geq \alpha - x\right\} \end{aligned}$$

If $\overline{\mathbb{E}}[x] = \underline{\mathbb{E}}[x]$ we call their common value the *price* of x, denoted $\mathbb{E}[x]$, and define the *upper variance* and *lower variance* of x by

$$\overline{\mathbb{V}}[x] \coloneqq \overline{\mathbb{E}}\left[(x - \mathbb{E}[x])^2\right] \qquad\qquad \underline{\mathbb{V}}[x] \coloneqq \underline{\mathbb{E}}\left[(x - \mathbb{E}[x])^2\right]$$

Now if $\overline{\mathbb{V}}[x] = \underline{\mathbb{V}}[x]$ we call their common value the *variance* of x, denoted $\mathbb{V}[x]$. We may analogously price any variable starting in situation t, resulting in $\overline{\mathbb{E}}_t[x], \underline{\mathbb{E}}_t[x], \overline{\mathbb{V}}_t[x], \underline{\mathbb{V}}_t[x]$.

A protocol is *coherent* in situation t if $\overline{\mathbb{E}}_t[0] = \underline{\mathbb{E}}_t[0] = 0$ and $\overline{\mathbb{E}}_t[x] \geq \underline{\mathbb{E}}_t[x]$ for all variables x. A protocol is *coherent* if it is coherent in all situations.

Probability For any event $E \subseteq \Omega$, we denote its indicator function (a variable) by

$$\mathbf{1}_E(\xi) \;=\; \begin{cases} 1 & \text{if } \xi \in E \\ 0 & \text{o.w.} \end{cases}$$

Now *upper probability* and *lower probability* of E are

$$\overline{\mathbb{P}}(E) \;\coloneqq\; \overline{\mathbb{E}}[\mathbf{1}_E] \qquad \underline{\mathbb{P}}(E) \;\coloneqq\; \underline{\mathbb{E}}[\mathbf{1}_E].$$

Forcing Fix an event $E \subseteq \Omega$. A strategy $\mathcal{P}$ for Skeptic *forces* E if

$$1 + \mathcal{K}^{\mathcal{P}}(t) \geq 0$$

for every situation $t \in \Omega^\diamond$ and

$$\lim_{n\to\infty} \mathcal{K}^{\mathcal{P}}(\xi^n) = \infty$$

for every path $\xi \notin E$.

Terminology If Skeptic can force E we say that

- E is *practically impossible* and happens *almost never*,
- E^c is *practically certain* and happens *almost surely*.

3 Kolmogorov's Axioms: Probability Measures

Consider a sample space Ω, a set $\mathcal{F}$ of subsets of Ω and a function $\mathbb{P} : \mathcal{F} \to \mathbb{R}$.

1. $\mathcal{F}$ is a field of sets (closed under union, intersection and difference).
2. $\Omega \in \mathcal{F}$.
3. $\mathbb{P}(E) \geq 0$ for each $E \in \mathcal{F}$.
4. $\mathbb{P}(\Omega) = 1$.
5. If $E, F \in \mathcal{F}$ are disjoint then $\mathbb{P}(E \cup F) = \mathbb{P}(E) + \mathbb{P}(F)$.
6. For any decreasing sequence $E_1 \supseteq E_2 \supseteq \cdots$ of events with empty intersection $\bigcap_{n=1}^{\infty} E_n = \emptyset$, $\lim_{n\to\infty} \mathbb{P}(E_n) = 0$.

Axioms 3-6 render $\mathbb{P}$ a *probability measure* on algebra $\mathcal{F}$. It can be extended to the smallest σ-algebra (closed under countable union and intersection) containing $\mathcal{F}$, and the extension will be *countably additive*

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} E_n\right) \;=\; \sum_{n=1}^{\infty} \mathbb{P}(E_n)$$

for every disjoint sequence $E_1, E_2, \ldots$ in σ-algebra. Conditional probability is *defined*,

$$\mathbb{P}(F|E) \;:=\; \frac{\mathbb{P}(E \cap F)}{\mathbb{P}(E)}$$

A *random variable* is a function $x : \Omega \to \mathbb{R}$ that is *measurable*, meaning that $\{\xi \in \Omega \mid x(\xi) \in I\} \in \mathcal{F}$ for each interval I. The expectation $\mathbb{E}[x]$ is the integral of x w.r.t. the measure $\mathbb{P}$.

4 GTP Martingales

Let us write v_i and x_i for the moves by Forecaster and Reality in round i. A process $\mathcal{S}$ is *predictable* if $\mathcal{S}(v_1 x_1 \ldots v_n x_n)$ does not depend on Reality's move x_n. A strategy for Skeptic is a predictable process.

A capital process (with any starting value) in a symmetric protocol is called a *martingale*. A *supermartingale* is a process of the form $\mathcal{T} = \mathcal{S} - \mathcal{B}$, where $\mathcal{S}$ is a capital process and $\mathcal{B}$ is an increasing process. We say that $\mathcal{S}$ is a *bounding process* for $\mathcal{T}$. A process $\mathcal{U} = \mathcal{T} + \mathcal{A}$ is called a *semimartingale* if $\mathcal{T}$ is a supermartingale and $\mathcal{A}$ is an increasing predictable process. We say that $\mathcal{A}$ is a compensator for $\mathcal{U}$.

5 Limits and Asymptotic notation

Fix a sequence $x_1, x_2, \ldots$. The following two limits are always defined (they may be infinite):

$$\begin{aligned} \limsup_{n\to\infty} x_n &= \inf_n \sup_{m \geq n} x_n \\ \liminf_{n\to\infty} x_n &= \sup_n \inf_{m \geq n} x_n \end{aligned}$$

If $\limsup_{n\to\infty} x_n = \liminf_{n\to\infty} x_n$ then their common value is called the limit, denoted $\lim_{n\to\infty} x_n$.

In addition to $x_1, x_2, \ldots$, consider a positive sequence $y_1, y_2, \ldots$. We say that x is *big-O* of y, denoted

$$x_n \;=\; O(y_n),$$

if $\exists n \geq 1 \exists c > 0 \forall m \geq n : x_m \leq c y_m$. Intuitively x_n grows as most as fast as y_n. We say that x is *little-o* of y, denoted

$$x_n \;=\; o(y_n),$$

if $\forall c > 0 \exists n \geq 1 \forall m \geq n : x_m \leq c y_m$. Here the intuition is that x_n grows strictly slower than y_n. If x is positive too, we have $\lim_{n\to\infty} \frac{x_n}{y_n} = 0$.

Some authors prefer to replace the equals signs in the above two displays by "is", "$\leq$" or "$\in$".