Betting rules and information theory

Giulio Bottazzi

LEM and CAFED
Scuola Superiore Sant'Anna

September, 2013

Outline

Simple betting in favorable games

The Central Limit Theorem

Optimal rules

The Game

Consider a simple game of chance, like tossing a coin or rolling some dice, in which you bet an amount x and if you win you get $2x$ and if you loose you loose the entire bet.

Let W_0 be the initial wealth so that if a fraction α is bet after the game you have

$$W_1 = \begin{cases} (1-\alpha)W_0 + 2\alpha W_0 = (1+\alpha)W_0 & \text{if you win} \\ (1-\alpha)W_0 & \text{if you loss} \end{cases}$$

How much would you like to bet?

Classification of games

A game is *favorable* if the expected value of what you get is higher than what you bet.

A game is *fair* if the expected value of what you get is equal to what you bet.

A game is *unfair* if the expected value of what you get is lower than what you bet.

wealth maximization

Let the probability to win be p. The expected wealth after participation into the game can be computed

$$\mathrm{E}[W_1] = (1-\alpha)W_0 + p2\alpha W_0 = (1-\alpha+2\alpha p)W_0$$

If $p = 2/3$ then $\mathrm{E}[W_1] = (1+\alpha/3)W_0$ and to maximize the expected wealth one must chose $\alpha = 1$, so that $\mathrm{E}[W_1] = 4/3W_0$.

The choice $\alpha = 1$ is made each time $2p-1 > 0$, that is for any favorable game.

Now suppose that the game is repeated T times. If one plays with $\alpha = 1$ after T games it will have an expected wealth of

$$\mathrm{E}[W_T] = \left(\frac{4}{3}\right)^T W_0$$

Now suppose that the game is repeated T times. If one plays with $\alpha = 1$ after T games it will have an expected wealth of

$$\mathrm{E}[W_T] = \left(\frac{4}{3}\right)^T W_0$$

but the probability to have nothing is equal to

$$\text{Prob}\{W_t = 0\} = 1 - \left(\frac{2}{3}\right)^T$$

Now suppose that the game is repeated T times. If one plays with $\alpha = 1$ after T games it will have an expected wealth of

$$\mathrm{E}[W_T] = \left(\frac{4}{3}\right)^T W_0$$

but the probability to have nothing is equal to

$$\text{Prob}\{W_t = 0\} = 1 - \left(\frac{2}{3}\right)^T$$

With $T = 10$ the expected wealth is $\mathrm{E}[W_T] \sim 17\, W_0$ but the probability to have nothing is around 98%.

Consider the evolution of wealth in a sequence of repeated games

$$W_{t+1} = \begin{cases} (1+\alpha)W_t & \text{if } \omega_t = 1 \\ (1-\alpha)W_t & \text{if } \omega_t = 0 \end{cases}$$

where ω_t is a random variable taking value 1 for a win and 0 for a loss. The sequence of random variables $\omega = (\omega_1, \ldots, \omega_T)$ is a *random process*.

Consider the evolution of wealth in a sequence of repeated games

$$W_{t+1} = \begin{cases} (1+\alpha)W_t & \text{if } \omega_t = 1 \\ (1-\alpha)W_t & \text{if } \omega_t = 0 \end{cases}$$

where ω_t is a random variable taking value 1 for a win and 0 for a loss. The sequence of random variables $\omega = (\omega_1, \ldots, \omega_T)$ is a *random process*.

The final level of wealth depends on the number of 1's and 0's in the sequence ω.

The probability to have t wins in a sequence of length T is distributed according to a binomial

$$\text{Prob}\{\text{wins} = t\} = \binom{T}{t} p^t \, (1-p)^{T-t}$$

The probability to have t wins in a sequence of length T is distributed according to a binomial

$$\text{Prob}\{\text{wins} = t\} = \binom{T}{t} p^t \, (1-p)^{T-t}$$

Example: $t = 2$ and $T = 3$. The number of sequence having 2 wins is 3 (the ways or "combinations" of selecting 2 elements in a set of 3) and the probability to obtain each sequence is $p^2 \, (1-p)$.

The probability to have t wins in a sequence of length T is distributed according to a binomial

$$\text{Prob}\{\text{wins} = t\} = \binom{T}{t} p^t (1-p)^{T-t}$$

Example: $t = 2$ and $T = 3$. The number of sequence having 2 wins is 3 (the ways or "combinations" of selecting 2 elements in a set of 3) and the probability to obtain each sequence is $p^2 (1-p)$.

In general how does the distribution of t looks like?

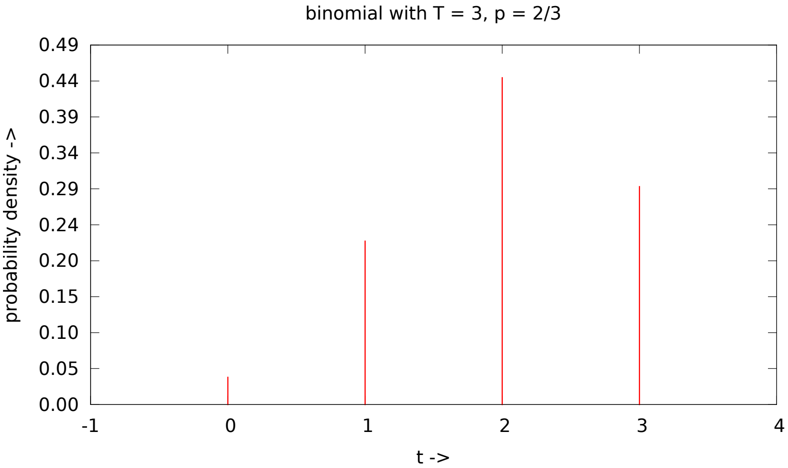

For $T = 3$ the possible values of t range from 0 to 3. Notice that Prob$[\text{wins} = 3]$ is larger than Prob$[\text{wins} = 0]$ because the probability of a win is larger than $1/2$.

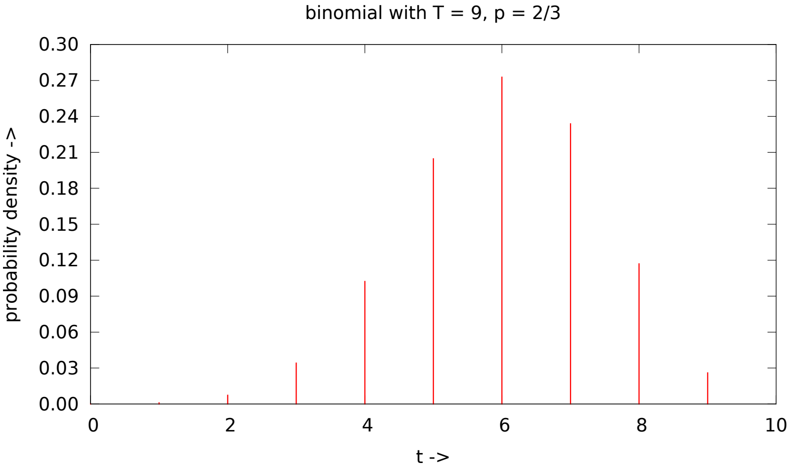

For $T = 9$ the possible values are from 0 to 9. The distribution is peaked around the mean and median value $p * T = 2/3 * 9 = 6$.

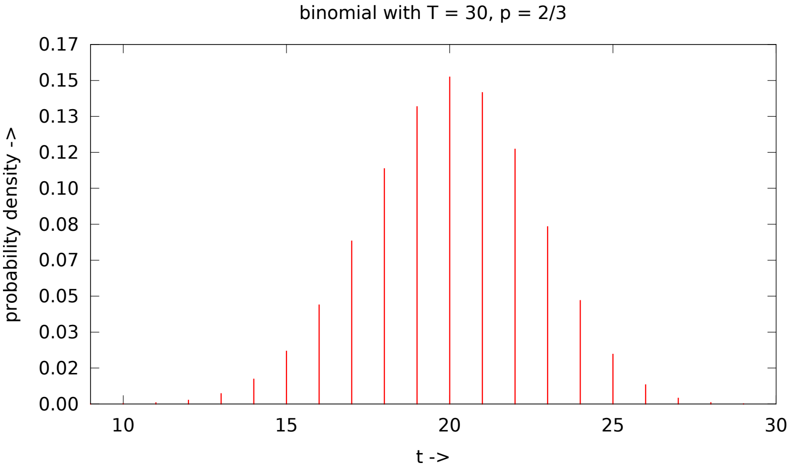

For $T = 30$ the distribution has already assumed a bell shape around $p * T = 2/3 * 30 = 20$.

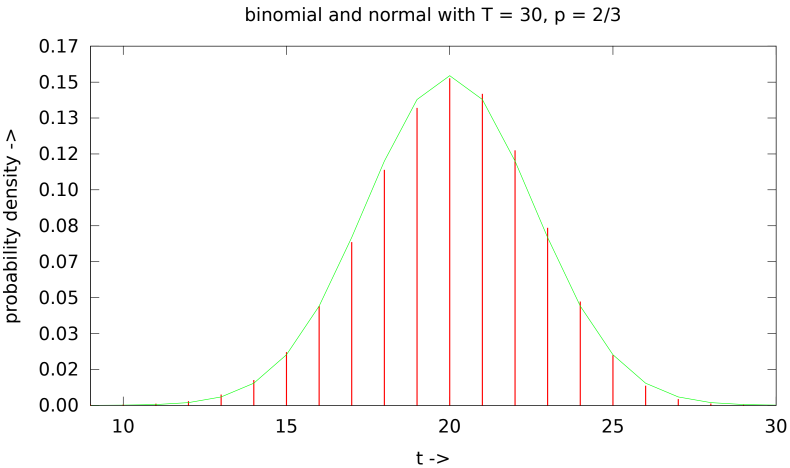

The distribution is well approximated by a normal distribution with mean $\mu = pT$ and standard deviation $\sigma = \sqrt{p(1-p)T}$. In this case the mean is 20 and the standard deviation 2.58. This is the Central Limit Theorem at work.

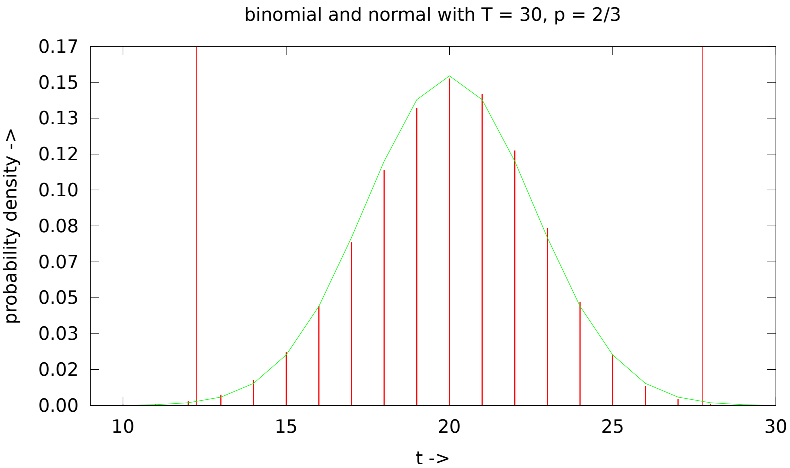

Moreover, due to the normal approximations, we can conclude that in the 99.73% of sequences the number of wins t will be in the interval

$$[\mu - 3\sigma, \mu + 3\sigma] \ .$$

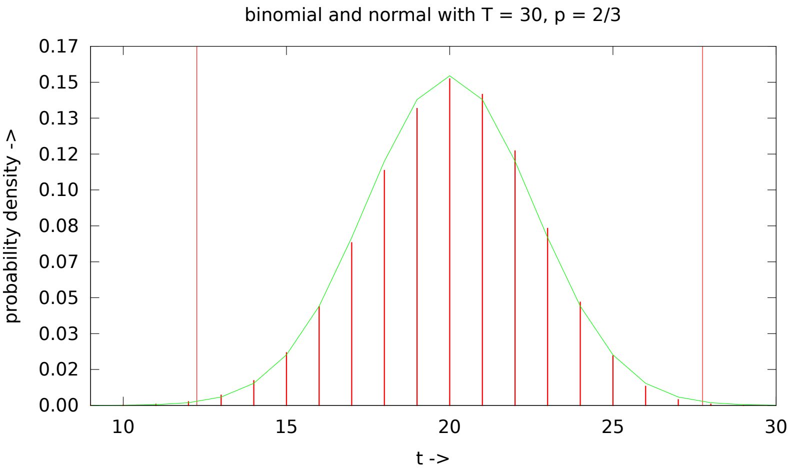

Different betting strategies work best in different cases: $\alpha = 0$ is the best if $t = 0$ and $\alpha = 1$ is the best if $t = T$. These are however unlikely events. It is more effective to have a strategy that works better around the modal value of the distribution of t.

The most probable level of wealth after T bets for a given α is $W_T^{\text{modal}} = (1+\alpha)^{pT}(1-\alpha)^{(1-p)T}W_0$. It grows exponentially with T.

The most probable level of wealth after T bets for a given α is $W_T^{\text{modal}} = (1+\alpha)^{pT}(1-\alpha)^{(1-p)T}W_0$. It grows exponentially with T.

Substituting the minimal and maximal value of t for the interval derived before, in the 99.73% of sequences the wealth at time T satisfies

$$\frac{1}{T}\log\frac{W_T}{W_T^{\text{modal}}} \in \left[\frac{3\sqrt{p(1-p)}}{\sqrt{T}}\log\frac{1-\alpha}{1+\alpha}, \frac{3\sqrt{p(1-p)}}{\sqrt{T}}\log\frac{1+\alpha}{1-\alpha}\right] .$$

The most probable level of wealth after T bets for a given α is $W_T^{\text{modal}} = (1+\alpha)^{pT}(1-\alpha)^{(1-p)T}W_0$. It grows exponentially with T.

Substituting the minimal and maximal value of t for the interval derived before, in the 99.73% of sequences the wealth at time T satisfies

$$\frac{1}{T}\log\frac{W_T}{W_T^{\text{modal}}} \in \left[\frac{3\sqrt{p(1-p)}}{\sqrt{T}}\log\frac{1-\alpha}{1+\alpha}, \frac{3\sqrt{p(1-p)}}{\sqrt{T}}\log\frac{1+\alpha}{1-\alpha}\right] .$$

When T becomes large the interval is reduced to 0 so that, in probability

$$\lim_{T\to\infty} W_t = W_T^{\text{modal}}$$

which is just the Law of Large Numbers.

The best betting rule in the long run is the rule which better performs when the number of wins is equal to the modal value

$$\alpha^* = \arg\max \left\{ (1+\alpha)^{pT} (1-\alpha)^{(1-p)T} \right\}$$

removing the unnecessary T and taking the log

$$\alpha^* = \arg\max \left\{ p \log(1+\alpha) + (1-p) \log(1-\alpha) \right\}$$

The best betting rule in the long run is the rule which better performs when the number of wins is equal to the modal value

$$\alpha^* = \arg\max \left\{ (1+\alpha)^{pT} (1-\alpha)^{(1-p)T} \right\}$$

removing the unnecessary T and taking the log

$$\alpha^* = \arg\max \{ p \log(1+\alpha) + (1-p) \log(1-\alpha) \}$$

Notice that the argument in the previous expression is just the expected log growth rate $E[\log(W_{t+1}/W_t)]$, thus α^* is the log-optimal rule.

In our case taking the derivative w.r.t. α one has

$$\frac{p}{1-\alpha^*} - \frac{1-p}{1-\alpha^*} = 0$$

whose solution reads $\alpha^* = 2p - 1$.

In our case taking the derivative w.r.t. α one has

$$\frac{p}{1-\alpha^*} - \frac{1-p}{1-\alpha^*} = 0$$

whose solution reads $\alpha^* = 2p - 1$.

According to the log-optimal rule: if the game is *favorable* (p>1/2) bet a finite amount of money; if the game is *fair* (p=1/2) bet nothing; if the game is *unfair* one wants to bet a negative amount, that is take the bet and commit herself to pay the possible payoff.

Wealth trajectories associated to different α's will diverge. Set $w_0 = 1$ and imagine to play different sequences of games of length T with a given α and report the wealth levels $\log(W_t)$ on a graph. Then repeat the experiment with a different α.

When the length of the sequence increases, the performances of the two strategies becomes more differentiated.

On a longer time horizon the divergence is clear.

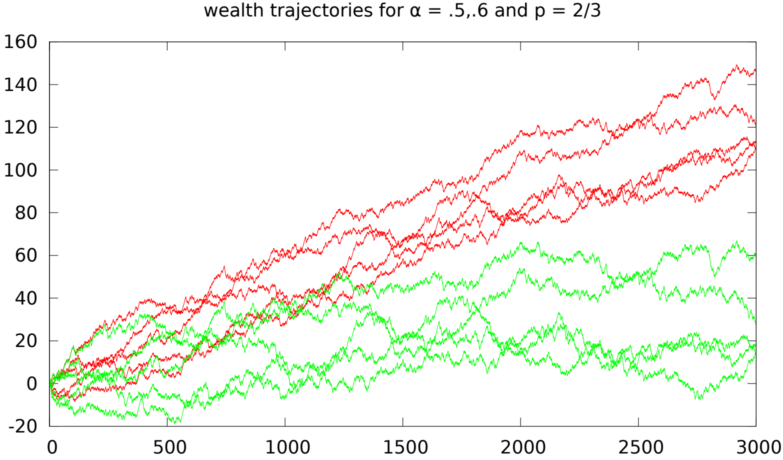

expected wealth ratios

Starting with the same wealth at time t, consider the wealth at time $t+1$ obtained with two different rules α and α'

$$W_{t+1}^{\alpha} = \begin{cases} (1+\alpha)W_t & \text{if win} \\ (1-\alpha)W_t & \text{if loss} \end{cases} \quad W_{t+1}^{\alpha'} = \begin{cases} (1+\alpha')W_t & \text{if win} \\ (1-\alpha')W_t & \text{if loss} \end{cases} .$$

then the expected ratio at time $t+1$ is

$$\mathrm{E}\left[\frac{W_{t+1}^{\alpha}}{W_{t+1}^{\alpha'}}\right] = p\frac{1+\alpha}{1+\alpha'} + (1-p)\frac{1-\alpha}{1-\alpha'}$$

ratio-optimal rule

The question is whether there exists a betting rule α' such that for any α it is

$$\mathrm{E}\left[\frac{W_{t+1}^{\alpha}}{W_{t+1}^{\alpha'}}\right] \leq 1 \ .$$

ratio-optimal rule

The question is whether there exists a betting rule α' such that for any α it is

$$\mathrm{E}\left[\frac{W_{t+1}^{\alpha}}{W_{t+1}^{\alpha'}}\right] \leq 1 \ .$$

By rearranging terms

$$\mathrm{E}\left[\frac{W_{t+1}^{\alpha}}{W_{t+1}^{\alpha'}}\right] = \frac{p}{1+\alpha'} + \frac{1-p}{1-\alpha'} + \alpha\left(\frac{p}{1+\alpha'} - \frac{1-p}{1-\alpha'}\right)$$

ratio-optimal rule

The question is whether there exists a betting rule α' such that for any α it is

$$\mathrm{E}\left[\frac{W^{\alpha}_{t+1}}{W^{\alpha'}_{t+1}}\right] \leq 1 \; .$$

By rearranging terms

$$\mathrm{E}\left[\frac{W^{\alpha}_{t+1}}{W^{\alpha'}_{t+1}}\right] = \frac{p}{1+\alpha'} + \frac{1-p}{1-\alpha'} + \alpha\left(\frac{p}{1+\alpha'} - \frac{1-p}{1-\alpha'}\right)$$

This can be made greater than 1 for some α if the expression inside parenthesis is different from zero. Thus the only solution is $\alpha' = 2p - 1 = \alpha^*$. ratio-optimal = log-optimal

absolute vs. relative

The log-optimality is related to the betting rules that *beats* all other rules: the probability to have a lower wealth than anyone else using a different rule tend to zero when the number of games increases.

absolute vs. relative

The log-optimality is related to the betting rules that *beats* all other rules: the probability to have a lower wealth than anyone else using a different rule tend to zero when the number of games increases.

This idea is NOT related to the maximization of some function of the individual wealth, but is instead based on a notion of *relative* performance of one rule with respect to other rules.

Entropy

Given a random variable X taking N values with probabilities $(p_1, \ldots, p_N)$ its *entropy* is defined as

$$S(X) = -\sum_i p_i \log p_i$$

Entropy

Given a random variable X taking N values with probabilities $(p_1, \ldots, p_N)$ its *entropy* is defined as

$$S(X) = -\sum_i p_i \log p_i$$

The entropy (or negative information, or information loss)

- does no depend on the value taken by the random variable, but only on its probability distribution.
- is non-negative
- is maximal when the "information" of the variable is minimal, $p_i = 1/N$
- it is minimal when the random variable is maximally informative, i.e. there is a j such that $p_j = 1$

Relative entropy

Given two random variables X and Y taking values on a set of N events with probabilities $(x_1, \ldots, x_N)$ and $(y_1, \ldots, y_N)$, the *relative entropy* or *Kullback-Leibler divergence* (sometimes distance) is

$$D(X|Y) = \sum_i x_i \log \frac{x_i}{y_i}$$

and measure the "information" missing for the knowledge of X when one knows Y.

Relative entropy

Given two random variables X and Y taking values on a set of N events with probabilities $(x_1, \ldots, x_N)$ and $(y_1, \ldots, y_N)$, the *relative entropy* or *Kullback-Leibler divergence* (sometimes distance) is

$$D(X|Y) = \sum_i x_i \log \frac{x_i}{y_i}$$

and measure the "information" missing for the knowledge of X when one knows Y.

The relative entropy

- is non-negative
- is asymmetric
- is minimal when X and Y coincides

entropy-optimal

The expression for the expected log growth rate can be rewritten as

$$\begin{aligned} E[\log(W_{t+1}/W_t)] &= -p\log\tfrac{p}{(1+\alpha)/2} - (1-p)\log\tfrac{1-p}{(1-\alpha)/2} + \\ &\quad p\log p + (1-p)\log(1-p) + log(2) \end{aligned}$$

entropy-optimal

The expression for the expected log growth rate can be rewritten as

$$E[\log(W_{t+1}/W_t)] \quad = \quad -p\log\frac{p}{(1+\alpha)/2} - (1-p)\log\frac{1-p}{(1-\alpha)/2} + p\log p + (1-p)\log(1-p) + log(2)$$

The betting rule can be seen as a random variable taking values on the two events *win* and *loss* with probabilities proportional to $1+\alpha$ and $1-\alpha$ respectively. Then

$$E[\log(W_{t+1}/W_t)] \sim -D(\{p, 1-p\}|\text{betting rule})$$

the relative entropy of the occurrences (win/loss) *given* the betting rule.

entropy-optimal

The expression for the expected log growth rate can be rewritten as

$$E[\log(W_{t+1}/W_t)] \quad = \quad -p\log\frac{p}{(1+\alpha)/2} - (1-p)\log\frac{1-p}{(1-\alpha)/2} + p\log p + (1-p)\log(1-p) + log(2)$$

The betting rule can be seen as a random variable taking values on the two events *win* and *loss* with probabilities proportional to $1+\alpha$ and $1-\alpha$ respectively. Then

$$E[\log(W_{t+1}/W_t)] \sim -D(\{p, 1-p\}|\text{betting rule})$$

the relative entropy of the occurrences (win/loss) *given* the betting rule.

Maximizing $E[\log(W_{t+1}/W_t)]$ is equivalent to minimize the relative entropy. entropy-optimal = log-optimal.

Summary

The *dominant* betting rule is the log-optimal rule, the Kelly rule

Summary

The *dominant* betting rule is the log-optimal rule, the Kelly rule

On average, no other rules can provide an higher level of wealth in the long run

Summary

The *dominant* betting rule is the log-optimal rule, the Kelly rule

On average, no other rules can provide an higher level of wealth in the long run

Under some well specified conditions, the log-optimal rule is the rule that minimizes the conditional entropy