The Kelly Betting System for Favorable Games.

A Simple Example. Suppose that each day you are offered a gamble with probability 2/3 of winning and probability 1/3 of losing. You may bet any positive amount you like, provided you have it. The amount you bet is either doubled or lost, independently each day. This gamble is only offered for 20 days. The question is: how much should you bet on each day?

Let X_0 denote your initial fortune, and let X_k denote your fortune after the kth gamble. On day k, you may bet b_k provided

$$0 \le b_k \le X_{k-1} \tag{1}$$

If you win on day k your fortune increases to $X_{k-1}+b_k$; if you lose it decreases to $X_{k-1}-b_k$. So,

$$X_k = \begin{cases} X_{k-1} + b_k & \text{with probability } 2/3 \\ X_{k-1} - b_k & \text{with probability } 1/3 \end{cases} \tag{2}$$

It is an easy problem to find, by backward induction, the betting system that maximizes $\mathrm{E}(X_{20}|X_0)$, your expected fortune after the 20 days have passed. It is to bet all you have each day. On the last day, given you have fortune X_{19}, to maximize $\mathrm{E}(X_{20}|X_{19})$ you bet it all, $b_{20} = X_{19}$, and your expected final fortune is $(2/3)2X_{19} + (1/3)0 = (4/3)X_{19}$. Since $\mathrm{E}(X_{20}|X_{18}) = \mathrm{E}(\mathrm{E}(X_{20}|X_{19})|X_{18}) = (4/3)\mathrm{E}(X_{19}|X_{18}) = (4/3)^2X_{18}$, the analysis by backward induction shows that $\mathrm{E}(X_{20}|X_0) = (4/3)^{20}X_0 = 315.34X_0$.

However, there is something disturbing about this solution. The distribution of X_{20} using the strategy of betting everything each time is

$$X_{20} = \begin{cases} 2^{20}X_0 & \text{with probability } (2/3)^{20} = .000030 \\ 0 & \text{with probability } .999970 \end{cases} \tag{3}$$

In other words, you will end up destitute with high probability. This is silly. You are playing a favorable game. You can control your destiny. You should want your fortune to tend to infinity if you play infinitely long. Instead you fortune tends to zero almost surely.

Proportional Betting Systems. To avoid this fate, consider using a proportional betting system. Such a system dictates that you bet a fixed proportion, call it π, of your fortune at each stage, $0 < \pi < 1$. This requires that we consider money infinitely divisible. Under such a system, your fortune will tend to infinity almost surely. There remains the problem of which value of π to use.

Let us search for a π that will make our fortune tend to infinity at the fastest rate. If you bet the same proportion π at each stage, then after n stages if you have won Z_n times and lost $n - Z_n$ times, you fortune will be

$$X_n = (1+\pi)^{Z_n}(1-\pi)^{n-Z_n}X_0, \tag{4}$$

where Z_n has the binomial distribution, $\mathcal{B}(n, 2/3)$. The fortune has increased by the factor

$$\begin{aligned}(1+\pi)^{Z_n}(1-\pi)^{n-Z_n} &= \exp\{Z_n \log(1+\pi) + (n - Z_n)\log(1-\pi)\} \\ &= \exp\{n[(Z_n/n)\log(1+\pi) + (1-(Z_n/n))\log(1-\pi)]\} \\ &\simeq \exp\{n[(2/3)\log(1+\pi) + (1/3)\log(1-\pi)]\}\end{aligned} \tag{5}$$

We can see that $(1+\pi)^{Z_n}(1-\pi)^{n-Z_n}$ tends to infinity almost surely exponentially fast. The rate of convergence is defined as the limit of $(1/n)\log((1+\pi)^{Z_n}(1-\pi)^{n-Z_n})$, which here is

$$(2/3)\log(1+\pi) + (1/3)\log(1-\pi). \tag{6}$$

The value of π that maximizes the rate of convergence is found by setting the derivative of (6) to zero and solving for π. It is $\pi = 1/3$. Therefore the optimal rate of convergence of X_n to infinity is achieved by wagering one-third of your fortune at each stage. The optimal rate is

$$(2/3)\log(4/3) + (1/3)\log(2/3) = .0566\cdots \tag{7}$$

This means that if two bettors compete, one using $\pi = 1/3$ and the other using $\pi \neq 1/3$, the fortune of the former will eventually be greater than that of the latter, and stay greater from some stage on.

The Kelly Betting System and Log Utility. Consider the same example now but with 2/3 replaced by an arbitrary probability of winning, p, and let us find the proportional betting system that maximizes the rate at which X_n tends to infinity. Of course if $p < 1/2$, you cannot achieve having your fortune tend to infinity. You might as well bet proportion 0 at each stage.

But suppose $p \geq 1/2$. The same analysis as above yields the rate of convergence to be (in place of (6))

$$p\log(1+\pi) + (1-p)\log(1-\pi). \tag{8}$$

This is easily seen to be maximized at $\pi = 2p-1$. This gives us what is known as the Kelly betting system; see Kelly (1956). If, for an even money bet, the probability of success is p, bet proportion $\pi(p)$ of your fortune, where

$$\pi(p) = \begin{cases} 0 & \text{if } p \leq 1/2 \\ 2p-1 & \text{if } p > 1/2. \end{cases} \tag{9}$$

The Kelly Betting Proportion.

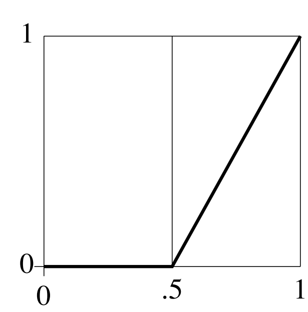

This rule also has an interpretation as the optimal rule of an investor who has the logarithm of the fortune as his utility function and wishes to maximize the expectation of the utility of his fortune at each stage. Given the fortune at stage $k-1$ and $b_k = \pi_k X_{k-1}$ with $0 \le \pi_k \le 1$, the expectation of $\log(X_k)$ given X_{k-1} is

$$\begin{aligned} \mathrm{E}(\log(X_k)|X_{k-1}) &= p\log(X_{k-1}+b_k) + (1-p)\log(X_{k-1}-b_k) \\ &= p\log(X_{k-1}(1+\pi_k)) + (1-p)\log(X_{k-1}(1-\pi_k)) \\ &= \log(X_{k-1}) + [p\log(1+\pi_k) + (1-p)\log(1-\pi_k)] \end{aligned} \tag{10}$$

This is maximized by the same value of π that maximizes (8). This gives a small sample justification of the use of the Kelly betting system.

This betting system may also be used if the win probabilities change from stage to stage. Thus, if there are n stages and the win probability at stage i is p_i for $i = 1, \ldots, n$, the Kelly betting system at each stage uses the myopic rule of maximizing the expected log, one stage ahead. Thus at stage k, you bet proportion $\pi(p_k)$ of your fortune. The asymptotic justification of the Kelly Betting System described above has a generalization that holds in this situation also. See Breiman (1961).

A General Investment Model with Log Utility. A striking fact is that this myopic rule is globally optimal for maximizing $\mathrm{E}(\log(X_n))$ under quite general conditions. One might think that if some future p_i are close to 1, one should pass up marginal p_i and wait for those that really make a difference. You do not have to look ahead; you don't even need to know the future p_i to act at stage k. Here is a quite general model that shows this. It allows information to be gathered at each stage that may alter the bettor's view of his future winning chances.

Play takes place in stages. The bettor starts with initial fortune X_0. At the beginning of the kth stage, the bettor chooses an amount b_k to bet and his fortune goes up or down by that amount depending on whether he wins or loses. We assume

$$0 \le b_k \le X_{k-1} \quad \text{and} \quad X_k = X_{k-1} + b_k Y_k \qquad \text{for } k = 1, \ldots, n \tag{11}$$

where Y_k represents the random variable which is plus one if he wins the kth bet and minus one if he loses it. Choice of b_k may depend on the information gathered through the previous $k-1$ stages. Let $\boldsymbol{Z}_0$ denote the information known to the decision maker before he makes the first bet ($\boldsymbol{Z}_0$ contains X_0), and let $\boldsymbol{Z}_k$ denote the random vector representing the information received through stage k. We assume that $\boldsymbol{Z}_k$ contains $\boldsymbol{Z}_{k-1}$ (he remembers all information gathered at previous stages) and that Y_k is contained in $\boldsymbol{Z}_k$ (he learns whether or not he wins in the kth stage). The joint distribution of the $\boldsymbol{Z}_k$ is competely arbitrary, possibly dependent, but known to the decision maker.

Theorem. *The betting system that maximizes* $\mathrm{E}(\log(X_n)|\boldsymbol{Z}_0)$ *is the Kelly betting system with* $b_k = \pi(p_k)X_{k-1}$, *where* $p_k = \mathrm{P}(Y_k = 1|\boldsymbol{Z}_{k-1})$.

Proof. By backward induction. At the last stage, we seek b_n to maximize

$$\begin{aligned} \mathrm{E}(\log(X_n)|\boldsymbol{Z}_{n-1}) &= \mathrm{E}(\log(X_{n-1}+b_nY_n)|\boldsymbol{Z}_{n-1}) \\ &= \mathrm{E}(\log(X_{n-1}(1+\pi_nY_n))|\boldsymbol{Z}_{n-1}) \\ &= \log(X_{n-1}) + \mathrm{E}(\log(1+\pi_nY_n)|\boldsymbol{Z}_{n-1}) \end{aligned} \tag{12}$$

since X_{n-1} is contained in $\boldsymbol{Z}_{n-1}$, and where $\pi_n = b_n/X_{n-1}$. Note that

$$\begin{aligned} \mathrm{E}(\log(1+\pi_n Y_n)|\boldsymbol{Z}_{n-1}) &= \mathrm{P}(Y_n = 1|\boldsymbol{Z}_{n-1})\log(1+\pi_n) + \mathrm{P}(Y_n = -1|\boldsymbol{Z}_{n-1})\log(1-\pi_n) \\ &= p_n \log(1+\pi_n) + (1-p_n)\log(1-\pi_n) \end{aligned} \tag{13}$$

We have seen in (8) that the value of π_n that maximizes this is $\pi(p_n)$. Since we know we will use this bet on the last stage, the problem for the next to last stage is to choose $\pi_{n-1} = b_{n-1}/X_{n-2}$ to maximize the expected value of (12) given $\boldsymbol{Z}_{n-2}$. The second term of (12) is independent of π_{n-1}, so this reduces to finding π_{n-1} to maximize $\mathrm{E}(\log(X_{n-1})|\boldsymbol{Z}_{n-2})$. By the same argument, the value of π_{n-1} that maximizes this is $\pi(p_{n-1})$ and this argument can be repeated down to the initial stage. ■

EXAMPLE 1. Suppose that the p_k are chosen i.i.d. from a given distribution $F(p)$ on $[0,1]$ known to you. Each p_k is announced to you just before you bet at stage k. How much should you bet at each stage to maximize $\mathrm{E}\{\log(X_n)\}$?

Here, $\boldsymbol{Z}_0$ contains only X_0, and $\boldsymbol{Z}_{k-1}$ contains the exact values of $p_1, \ldots, p_k$. Thus at stage k, you should bet $\pi(p_k)$. This is a standard example where the Kelly betting system is optimal for maximizing the log of your fortune n steps ahead.

EXAMPLE 2. Suppose the win probability at each trial is p but is unknown to you. Suppose the prior distribution of p is the uniform distribution on $[0,1]$. How much should you bet at each stage to maximize $\mathrm{E}\{\log(X_n)\}$?

Here the only information given to us at each stage is whether we win or lose, so $\boldsymbol{Z}_{k-1}$ contains just the record of wins and losses during the first $k-1$ stages. At each stage $\mathrm{P}(Y=1) = p$ but p is unknown. At stage 0, we have $\mathrm{E}(p|\boldsymbol{Z}_0) = 1/2$, the mean of the uniform distribution, $\mathrm{Beta}(1,1)$; so $p_1 = 1/2$, and $\pi(p_1) = 0$. If one observes S_{k-1} successes during the first $k-1$ trials, the posterior distribution of p is $\mathrm{Beta}(1+S_{k-1}, k-S_{k-1})$, so $p_k = \mathrm{E}(p|\boldsymbol{Z}_{k-1}) = \mathrm{E}(p|S_{k-1}) = (1+S_{k-1})/(k+1)$. At Stage k we bet proportion $\pi(p_k)$ of our current fortune. Thus, if we observe a failure on the first stage, we bet nothing on the second stage since $\pi(1/3) = 0$. However if a success occurs on the first stage, we bet proportion $\pi(2/3) = 1/3$ of our fortune at the second stage.

If we were interested in maximizing the *expected fortune* n steps ahead in this example, the myopic rule of maximizing the expected fortune one step ahead at each stage is not optimal. In fact for $n \geq 2$, betting everything on the first stage is strictly better than betting nothing. You may even be called upon to bet everything on a gamble unfavorable to you.

EXAMPLE 3. Suppose a deck of cards consists of n red cards and n black cards. The cards are put in random order and are turned face up one at a time. At each stage you may bet as much of your fortune as you like on whether the next card will be red or black. There are $2n$ stages. How much should you bet at each stage to maximize $\mathrm{E}\{\log(X_{2n})\}$, and should you bet on red or black?

Here the observations at each stage tell you how many red and black cards are left in the deck. So if at stage k there are r_k red and b_k black cards left in the deck, then if $r_k \geq b_k$

one should bet $\pi(r_k/(r_k+b_k))$ on red, while if $b_k \geq r_k$ one should bet $\pi(b_k/(r_k+b_k))$ on black.

A practical example of the use of this type of information gathering occurs in the game of blackjack. At the start of this game with a well-shuffled deck, the game is biased against the player, i.e. the probability of win is less than 1/2, so the player bets as small an amount as the rules of the blackjack table allow. As the game progresses, cards are observed that are removed from the deck, and occasionally there occur reduced decks that are favorable to the player. When these situations occur the player will bet more heavily, depending on how favorable the reduced deck is.

Competitive Optimality of Kelly Betting. We have seen that the Kelly betting system has a large sample optimality property. Here is a finite sample optimality property due to Bell and Cover (1980) involving competition between two players to see who can end up with the most capital.

Players I and II start with 1 dollar each. They are going to play a favorable game with probability $p > 1/2$ of doubling your bet and probability $1-p$ of losing it. However, as the first move of the game, they can exchange their dollar for the outcome of any non-negative random variable with expectation 1. Let U, resp. W, denote the fortune of Player I, resp. Player II, after this move. Thus,

$$U \geq 0, \quad W \geq 0, \quad \text{and} \quad \mathrm{E}(U) = \mathrm{E}(W) = 1. \tag{14}$$

Then, not informed of the opponent's fortune, each player chooses an amount to bet on the favorable game. Let $a(U)$, resp. $b(W)$, denote the proportion of I's fortune, resp. II's fortune, that is bet. Let Z denote the indicator outcome of the favorable game, i.e.

$$Z = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability 1-p} \end{cases} \tag{15}$$

Let X, resp. Y, denote the final fortune of Player I, resp. Player II. Then

$$\begin{aligned} X &= U[(1-a(U)) + a(U)2Z] \\ Y &= W[(1-b(W)) + b(W)2Z]. \end{aligned} \tag{16}$$

Player I wins if $X > Y$ and Player II wins if $Y > X$. Player I wants to maximize $\mathrm{P}(X > Y)$ and Player II wants to minimize this.

The game is symmetric so the value if it exists is 1/2, and any strategy optimal for one player is optimal for the other.

Theorem. *The value of the game is 1/2, and an optimal strategy for Player I is choose the distribution of U to be uniform on (0,2),*

$$\mathrm{P}(U \leq u) = \begin{cases} u/2 & \text{if } 0 \leq u \leq 2 \\ 1 & \text{if } u > 2 \end{cases} \tag{17}$$

and $a(u)$ to be Kelly's betting proportion, $a(u) = 2p - 1$, independent of u.

Proof. Suppose Player I uses the strategy of the theorem. We show that $\mathrm{P}(X < Y) \le 1/2$ for all strategies of Player II. Then, from (16)

$$\begin{aligned}
\mathrm{P}(X < Y) &= \mathrm{P}(U[(1-a(U)) + a(U)2Z] < W[(1-b(W)) + b(W)2Z]) \\
&= \mathrm{P}\left(U < \frac{W[(1-b(W)) + b(W)2Z]}{2(1-p) + 2(2p-1)Z}\right) \\
&\le \frac{1}{2}\mathrm{E}\left(\frac{W[(1-b(W)) + b(W)2Z]}{2(1-p) + 2(2p-1)Z}\right) \quad \text{from (17)} \\
&= \frac{1}{2}\mathrm{E}\left(p\frac{W(1+b(W))}{2p} + (1-p)\frac{W(1-b(W))}{2(1-p)}\right) \quad \text{from (15)} \\
&= \frac{1}{2}\mathrm{E}(W) = \frac{1}{2} \quad \text{from (14). } \blacksquare
\end{aligned}$$

Note that this proof requires the players to wager on the same outcome of the favorable game. The theorem is no longer true if they may wager on independent games. For example, take $p = 3/4$, and suppose that Player I uses the strategy of the theorem, and Player II chooses W to be degenerate at 1, and then bets her whole fortune. Player I ends up with a distribution for U that has density 1/2 on the interval [0,1] and density 1/4 on the interval [1,3]. Player II ends up at 0 with probability 1/4 and at 2 with probability 3/4. If she ends up at 0, she loses; if she ends up at 2, she has probability 3/4 of winning. Her overall probability of winning is therefore $(1/4)(0) + (3/4)(3/4) = 9/16$, greater than 1/2.

Exercises.

1. (Bayes Hypothesis Testing.) You know the probability of win, p, is either 1/4 or 3/4. The prior probability of p is $\mathrm{P}(p = 1/4) = \mathrm{P}(p = 3/4) = 1/2$. You never observe p; you only learn whether you win or lose. How much of your fortune should you bet at each stage to maximize $\mathrm{E}\{\log(X_n)\}$?

2. Here is a problem that combines the type of learning that occurs in Examples 2 and 3. A deck of n cards is formed containing Y winning cards and $n - Y$ losing cards. Cards are turned face up one at a time and at each of the n stages you may bet as much of your fortune as you like that the next card will be a winning card. How much of your fortune should you bet at each stage to maximize $\mathrm{E}\{\log(X_n)\}$, if
(a) the prior distribution of Y is binomial, $\mathcal{B}(n, 1/2)$?
(b) the prior distribution of Y is uniform on the set $\{0, 1, \ldots, n\}$?

References.

Robert M. Bell, Thomas M. Cover (1980) "Competitive Optimality of Logarithmic Investment", *Math. of Operations Research,* **5**, 161-166.

L. Breiman. (1961) Optimal gambling systems for favorable games. *Fourth Berk. Symp. on Math. Stat. and Prob.*, **1**, 65-78.

J. B. Kelly. (1956) A new interpretation of information rate. *Bell System Technical J.*, **35**, 917-926.

Google 'Kelly Betting System' for further references.

Solutions.

1. At stage k, after observing s_k wins and $k - s_k$ losses, the posterior probability of win is $p_{k+1} = 3^{s_k}/(3^{s_k} + 3^{k-s_k})$. So bet $\pi(p_{k+1})$ at stage $k+1$.

2. (a) The prior density is $g(y) = \binom{n}{y}(1/2)^n$. Let Z_k denote the number of wins through stage k. The distribution of Z_k given $Y = y$ is hypergeometric with density $f(z|y) = \binom{y}{z}\binom{n-y}{k-z}/\binom{n}{k}$. The joint distribution of Y and Z is the product, $g(y)f(z|y)$, and the posterior density of Y given Z is proportional to this:

$$g(y|z) \propto \binom{n}{y}\binom{y}{z}\binom{n-y}{k-z} \propto \frac{1}{(y-z)!(n-y-k+z)!}$$

for $y = z, z+1, \ldots, n-k+z$. Let $R = Y - Z_k$ be the number of winning cards remaining after stage k. Change variable from Y to R to find

$$g(r|z) \propto \frac{1}{r!(n-k-r)!} \propto \binom{n-k}{r}(1/2)^{n-k}$$

for $r = 0, 1, \ldots, n-k$. So the posterior distribution of R is binomial $\mathcal{B}(n-k, 1/2)$, independent of Z_k. The expectation of win is $p_{k+1} = \mathrm{E}(R/(n-k)) = 1/2$. So bet nothing.

(b) The prior density for Y is $g(y) = 1/(n+1)$ for $y = 0, 1, \ldots, n$. Let Z denote the number of winning cards observed through stage k. Given $Y = y$, the density of Z is Hypergeometric, $f(z|y) = \binom{y}{z}\binom{n-y}{k-z}/\binom{n}{k}$. The posterior density of Y given $Z = z$ is proportional to the product $g(y)f(z|y)$. Thus,

$$g(y|z) \propto \frac{y!}{(y-z)!} \cdot \frac{(n-y)!}{(n-y-k+z)!}$$

It is preferable to work with the remaining number of winning cards left after k stages, $R = Y - Z$ in place of Y, and the number of cards remaining $n' = n - k$ in place of n. Making this change of variable, we find the posterior density of R given $Z = z$ as

$$g(r|z) \propto \frac{(r+z)!}{r!} \cdot \frac{(n'+k-r-z)!}{(n'-r)!} \quad \text{for } r = 0, 1, \ldots, n'$$

We show this is a beta-binomial distribution. The beta-binomial distribution $\mathcal{BB}(\alpha, \beta, n)$ has density

$$\begin{aligned} h(x) &= \binom{n}{x}\frac{\Gamma(\alpha+\beta)\Gamma(x+\alpha)\Gamma(n-x+\beta)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(n+\alpha+\beta)} \\ &\propto \frac{(x+\alpha-1)!}{x!} \cdot \frac{(n-x+\beta-1)!}{(n-x)!} \quad \text{for } x = 0, 1, \ldots, n \end{aligned}$$

From this we see that the distribution of R given $Z = z$ is $\mathcal{BB}(z+1, k-z+1, n')$. The expectation of $\mathcal{BB}(\alpha, \beta, n)$ is $n\alpha/(\alpha+\beta)$. The probability of win on the next stage given R is R/n'. Its expectation is therefore $p_{k+1} = \mathrm{E}(R/n') = (z+1)/(k+2)$. At stage $k+1$, you should bet proportion $\pi(p_{k+1})$ of your fortune.