

Lecture 8: Kelly's horse race

Biology 429
Carl Bergstrom

February 4, 2008

Sources: This lecture follows Cover and Thomas Chapter 6. The original work is due to J. L. Kelly (1956). As usual, some of the text and equations are taken directly from those sources.

In today's lecture we treat a subject that is not usually included in information theory texts: the relationship between information theory and gambling. Tangential as the subject may appear at first glance, this topic offers a great opportunity to explore a model of the value of information under geometric compounding or growth, and thereby to further develop our intuitions about Shannon information, mutual information, and so forth—and to start to see how information-theoretic measures may bear on the value of information.

Suppose that you've done something really, really, bad. For your sins, you are condemned to spend eternity at a seedy, smoky, run-down racetrack, betting all of your money on every single horse race. Each race features m horses, with win probabilities of p_1, p_2, \dots, p_m . The bookie (who presumably earned his job by doing something even worse than whatever you did) allows you to be on the winning horse (no bets to place at this track!) at odds o_1, o_2, \dots, o_m for 1.

Even though you are condemned to this horsetrack purgatory, you figure that there is no point in losing money. So you come to grips with your lot as a sempeternal gambler, and start planning your strategy. You note that if you bet all of your money on a single horse, sooner or later you're sure to lose it all and then go through eternity absolutely penniless. So you realize that you have to *hedge your bets*, betting a fraction of your wealth b_1, b_2, \dots, b_m

on each horse (where $\sum_i b_i = 1$ by divine fiat: you have to bet your whole wealth on every race.)

On the back of a strained napkin in the empty cocktail lounge overlooking the East straightaway of the track, you figure out that after n races, your wealth will be

$$S_n = S_0 \prod_{i=1}^n b(X_i) o_i(X_i)$$

where X is the random variable representing which horse wins the race. Thus, you calculate over a flat and rather warm gin and tonic, your wealth will grow asymptotically at the rate $W(n, p)$:

$$S_n/S_0 = 2^{nW(b,p)}$$

where $W(b, p) = E[\log S(X)] = \sum_{k=1}^m p_k \log b_k o_k$.

Proof: Taking the log of your growth in wealth at time n ,

$$\frac{1}{n} \log S_n/S_0 = \frac{1}{n} \sum_{i=1}^n \log S(X_i) \rightarrow E[\log S(X)] = W(b, p)$$

where the convergence is in probability by the weak law of large numbers. Multiplying both sides by n and applying $\exp[\cdot]$, we get $S_n/S_0 = 2^{nW(b,p)}$.

That's great, you think, ordering another drink from the taciturn and chain-smoking bartender. Now I've got an expression for my wealth. Calculus works even in Hell, so I can find a betting strategy that maximizes it. You apply the technique of Lagrange multipliers to guess at a solution, and manage eventually (after a third gin and tonic and several more uneventful races) to prove the following theorem:

Theorem 1 *Proportional gambling is log-optimal. The optimal betting scheme at this track is the "proportional gambling" scheme $b^* = p$ in which you bet on each horse according to its probability of winning, independent of the odds offered. This scheme gives a doubling rate of*

$$W^*(p) = \sum p_i \log o_i - H(p)$$

Proof:

$$\begin{aligned}
W(b, p) &= \sum p_i \log b_i o_i \\
&= \sum p_i \log \left(\frac{b_i}{p_i} p_i o_i \right) \\
&= \sum p_i \log o_i - H(p) - D(p||b) \\
&\leq \sum p_i \log o_i - H(p)
\end{aligned} \tag{1}$$

We have equality in the last step only if $p = b$, i.e, if the gambling portfolio matches the winning probabilities. This is rather surprising: you ignore the odds entirely and simply match probabilities when placing your bet.

This place is hell for bookies too; the bookie is required to offer fair odds o_i such that $\sum \frac{1}{o_i} = 1$. To maximize his take, the bookie does his best to match the odds that he offers to the probabilities that the horses will actually win, and he comes up with estimates $r_i = 1/o_i$ of these winning probabilities. If he doesn't know the true odds p_i , you realize that you can make some money off of him:

$$\begin{aligned}
W(b, p) &= \sum p_i \log b_i o_i \\
&= \sum p_i \log \left(p_i \frac{b_i}{p_i} \frac{1}{o_i} \right) \\
&= D(p||r) - D(p||b).
\end{aligned}$$

You maximize your winnings off of him if you match probabilities by setting $b = p$ such that $D(p||b) = 0$. The amount that you can make is a function of the relative entropy between his estimates and the actual winning probabilities.

One day, as you sit sipping a badly mixed martini, you make a new friend. Your new friend has some...connections...at the track. In fact, he (usually) knows which horse is going to win each race, before the race has actually been run. Quite useful!

He offers to share some of this information, at a price. You get these sense that this isn't really an offer you can refuse, so you decide to work out the the actual value of his information, and offer him the whole thing. You

promise to get back to him, and grabbing another handful of pretzels and a cocktail napkin, you get to work.

Suppose he tells you his guess of who will win the race; call this Y . The actual winner is X . You now want to find a betting strategy $b(x|y)$ that maximizes your winnings given your... well, “side information” seems a suitable delicate term. The optimal betting strategy will then give you a winnings rate of $W^*(X|Y)$ compared to your previous winnings of $W^*(X)$. How do these compare? As before, the optimal betting strategy $b(x|y)$ will be a proportional betting strategy $b(x|y) = p(x|Y)$. (The same proof holds). Thus, you are able to show that:

Theorem 2 *The increase in W due to side information Y is exactly the mutual information between X and Y .*

Proof: Under optimal betting with side information, you bet $b(x|y) = p(x|y)$ and obtain a winnings rate of

$$W^*(X|Y) = \sum_{x,y} p(x,y) \log o(x)p(x|y).$$

Expanding the two product terms in the log into two separate sums,

$$W^*(X|Y) = \sum_x p(x) \log o(x) + \sum_{x,y} p(x,y) \log p(x|y).$$

By the definition of conditional entropy,

$$W^*(X|Y) = \sum_x p(x) \log o(x) - H(X|Y).$$

You already know that the without side information the optimal betting strategy gives you an earnings rate of

$$W^*(X) = \sum_x p(x) \log o(x) - H(X).$$

Therefore the side information increases your fractional earnings by $H(X) - H(X|Y) = I(X; Y)$. You now know how much to offer (as a percentage cut of your winnings) to your friend for his side information.