Gamble with time

Probabilistic thinking has gone awry and stands in need of correction.

Date: circa 1950. Topic: information theory. Location: AT&T Bell Laboratories. Ask any physicist, computer scientist or electrical engineer to add another item to this list and they will probably say "Person: Claude Shannon", naming the engineer who famously laid the foundations of information theory in his landmark paper, 'A Mathematical Theory of Communication', published in 1948 in *The Bell System Technical Journal*.

But if Shannon is the most obvious response, it isn't the only one that makes sense. In 1956, a physicist colleague of Shannon's from Texas, John Larry Kelly, published a lesser-known but equally profound paper looking at how a gambler — facing a series of risky bets — could optimize his winnings in the long run and avoid ruin along the way. Kelly offered a concise answer: the gambler should at each stage wager a specific fraction of his or her current wealth, the fraction determined by the odds and potential winnings.

Today, more than half a century later, Kelly's solution — now known as the Kelly Criterion — finds wide use in finance as a tool for guiding investments over time. But the deeper meaning of Kelly's perspective, and its relation to other ideas about optimal behaviour in the presence of risk, remain controversial. This is clear from the polarized responses to the recent work of physicist Ole Peters, who has re-visited Kelly's thinking — and applied it to resolve a centuries' old paradox (*Philos. Trans. R. Soc. A* **369,** 4913–4931; 2011).

Consider how much you might be willing to pay to play a lottery based on a coin flip. If the first flip is heads, you win $1. If tails, you flip again. Heads on the second toss and you win $2, otherwise you flip again, with heads on the third toss giving $4 and so on. The lottery pays out 2^n dollars if the first head comes up on the nth roll. An easy calculation shows that the expected payout of the lottery is actually infinite, as the size of the payout grows just as fast as its likelihood decreases.

There's nothing paradoxical in this, of course, but what has seemed contradictory — since the eighteenth century, when Nicolas Bernoulli proposed the puzzle — is that no sensible person would pay much to play this game, despite the infinite expected pay-off. Real people do not find this lottery appealing and generally offer less than $10 or so to play. Maximizing expected return is, in this case, just not what people do.

Of course, this is still only strange if you believe for some reason that people should act to maximize expected return — a notion first proposed by Pierre de Fermat and often taken for granted in economics. However, this is where Peters, inspired by Kelly's perspective, suggests that probabilistic thinking has gone awry and stands in need of correction. The way to make sense of the 'paradox', he argues, is to think in terms of time.

After all, the familiar calculation based on expected value actually entails supposing that the gamble plays out simultaneously in several parallel worlds, one for each possible outcome. The result of the calculation is influenced by every one of these no matter how unlikely. This is the essence of probability, of course, yet it clearly introduces an artificial element into the situation. An alternative way to treat the problem — Kelly's way — is instead to calculate the expected pay-off from a string of such wagers actually playing out in real time, as a real person would experience if trying to learn how to play the game by trial and error.

Mathematically, this way of thinking leads Peters to consider the time average of the growth rate (log return) of the wealth of a player who begins with wealth W and plays the gamble over N periods, in the limit as N goes to infinity. A simple calculation leads to a formula for this growth rate that gives more sensible guidance. The rate is positive when the cost C of playing is sufficiently low, relative to a player's wealth, and negative when C becomes too high. Hence, how much you ought to be willing to pay depends on your initial wealth, as this determines how much you can afford to lose before going broke. Plug in real numbers, and the results predict fairly well what real people feel about the bet.

This aspect — the dependence of the calculation on the gambler's initial wealth — doesn't figure in the usual ensemble average in any way. Coincidentally, this result is identical to a solution to the paradox, originally proposed by Daniel Bernoulli — that people don't care about the absolute pay-off, but about the logarithm of the pay-off. But Bernoulli offered that solution without any fundamental justification; it emerges more naturally if one simply imagines playing in time, not in parallel worlds.

Some people argue that this in-time perspective still doesn't really solve Bernoulli's paradox (known more usually as the St. Petersburg paradox). After all, the original question asks what to do in one play, not in repeated play, which is required to calculate the time average. To my mind, this objection doesn't really hold. After all, any gamble has to be situated in time. You only care about winning or losing a gamble because you intend to go on living afterwards, and can profit from the extra wealth, facing future challenges in a more secure position. Psychologically, it's more or less impossible to consider any gamble as happening outside of time, because we live in time (and living in time is part of what makes us averse to risk, because we actually have to live with the consequences).

There is one other aspect to Peters's resurrection of Kelly's way of thinking that holds special interest for physicists. We're familiar with the notion of ergodicity — the property of the dynamics of a system that makes a time average equal to an ensemble average. This is a neat trick, hugely useful, as ensemble averages are typically much easier to calculate. The assumption of ergodicity lies at the basis of statistical mechanics.

And it is the failure of this assumption that distinguishes time and ensemble averages in gambling problems also. Economists have long relied on the equality of ensemble and time averages, assuming that the probabilities they deal with often have this feature. But the multiplicative growth process involved in any situation of repeated gambles is necessarily not ergodic. Go broke at one time step, and you are permanently out of the game, stuck at wealth = 0, a situation never captured by the ensemble average, which assumes continued exploration of the space of outcomes.

It's curious that a Bell Labs physicist hit on this idea so long ago, and we still can't quite fathom its full implications. ❐

MARK BUCHANAN