From *SIAM News*, Volume 39, Number 3, April 2006

The Kelly Criterion: Fallacy or Kuhnian Paradigm Shift Waiting to Happen?

Fortune's Formula: The Untold Story of the Scientific Betting System That Beat the Casinos and Wall Street. *By William Poundstone, Hill and Wang, New York, 2005, 400 pages, $27.00.*

BOOK REVIEW

By James Case

William Poundstone is the author of nine previous books, including *Carl Sagan: A Life in the Cosmos* and *Prisoner's Dilem-ma*. Despite—or possibly because of—his lack of scientific credentials, Poundstone seems to be drawn to scientifically controversial subjects. The formula of the title—which is intended to parlay modest sums of money into veritable fortunes, with minimal waste of time, by means of an appropriate sequence of wagers—is certainly controversial. Despite years of success at racetracks, in casinos, and on Wall Street, it has been denounced as "a fallacy" by some of the 20th century's most celebrated economists. Claude Shannon and John Kelly discovered the formula at Bell Labs, in 1956.

Athough Poundstone seems to be convinced that the formula is the product of a genuine collaboration, the original paper [2] on the subject was written by Kelly alone. An earlier version, which spoke openly of bookies and insider information, had been declared unacceptable by AT&T management, which was never keen to advertise the fact that bookies long represented an embarrassingly large fraction of the firm's customer base. As a result, Shannon found himself serving as an "anonymous" in-house referee, helping Kelly to prepare a suitably sanitized version of what both considered joint research, for publication in the firm's own *Bell System Technical Journal*.

The published version of the paper concerned a bettor endowed with a private channel of communication that delivers a steady stream of track-related information unavailable to rival bettors. Although the information transmitted need not be 100% accurate, it must be reliable enough to give the recipient an exploitable edge. Kelly's paper demonstrated that, just as information can be transmitted over a noisy communication channel, at or near the so-called "channel capacity," with negligible chance of error, so a bettor can compound his or her net worth at a certain maximum rate with virtually no risk of ruin. Kelly went on to express that fact in the form of an equation: $G_{\max} = R$, where G denotes the rate of growth (logarithmic time derivative) of one's net worth and R refers to the rate at which (Shannon) information is being transmitted over the private channel.

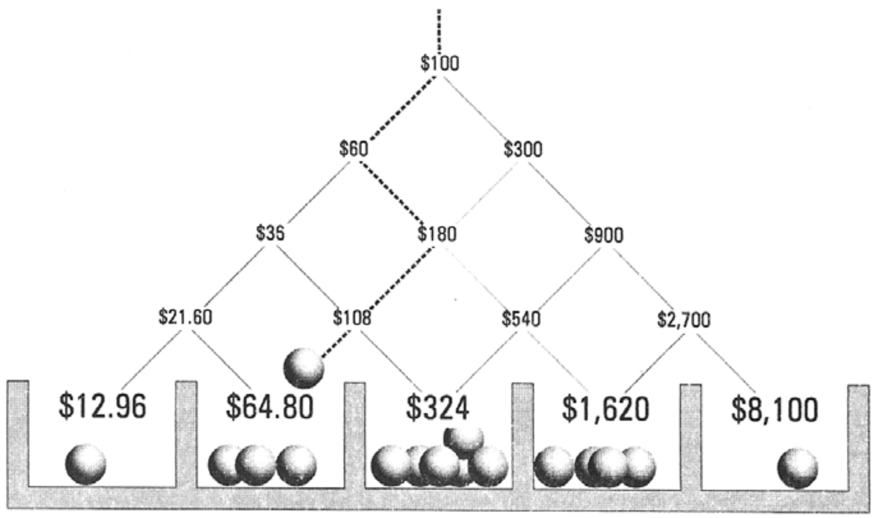

Figure 1. *Pascal's triangle. From* Fortune's Formula.

The key decision to be made for each successive wager concerns the fraction f of one's current bankroll to place at risk. Poundstone expresses Kelly's answer in the form f = edge/odds, where "edge" refers to the expected gain from the current bet and "odds" refers to the multiple of the amount of the bet that the winning bettor stands to receive. Poundstone illustrates the use of the formula with an example based on a pinball machine, based in turn on Pascal's triangle. The essential features are shown in Figure 1.

Starting with a $100 bankroll, the bettor is to toss a fair coin four times, getting back six times the amount of his bet each time he throws a head, and nothing each time he throws a tail. The expected gain on this highly favorable bet is a princely 200%—players forge ahead $5 with every win and fall back $1 with every (equally likely) loss. Meanwhile, the payoff (or tote board) odds are 5 to 1, because the successful bettor gets back five times the amount of the bet, in addition to the bet itself. Hence, f = edge/odds = 2/5 = 40% in the present case; the Kelly rule instructs the bettor to wager 40% of his bankroll at each opportunity. By so doing, he stands 1 chance in 16 of ending up with a mere $12.96 in his bankroll, and an equal chance of possessing $8100. The intermediate possibilities are $64.80 and $1620, each with 4 chances in 16, and $324, with 6 chances in 16. The arithmetic and geometric means of these 16 possible bankrolls are $1049.76 and $324, respectively. The Kelly rule for deciding on the fraction f of a bankroll to bet is to choose the one that maximizes the geometric mean of those 16 possibilities.

More generally, B_n denotes the bettor's bankroll after n bets. B_n is obviously a random variable, dependent on the vector Ω_n of outcomes—wins and losses—of the first n bets, as well as the Kelly fraction f. Thus, $B_n = B_n(\Omega_n, f)$. In simple circumstances, such as Pound-stone's pinball machine, Ω_n can be replaced by the random number S_n of wins in the first n bets, because the order of the wins and losses doesn't matter. Then, if $p_n(s) = \text{prob}\{S_n = s\}$,

$$G_n(f) = \Pi_s B_n(s,f) \uparrow p_n(s)$$

is the weighted geometric mean of the value distribution $F_n(v) = \text{prob}\{B_n(S_n,f) \uparrow v\}$.

It turns out—in surprisingly general circumstances—that the same value of f that maximizes $G_n(f)$ also maximizes $r_n(f)$, the expected compound rate of return on investment after n bets, and $E(\log(B_n(f)))$, the expected value of the logarithm of the bettor's bankroll after n bets, not to mention the median of the value distribution F_n. Proofs for the case in which successive bets represent independent Bernoulli trials can be found in [3], which is downloadable. Breiman [1] gave more general proofs in 1961.

Breiman also succeeded in showing that the expected time for current capital $B_n(\varphi)$ to reach any fixed pre-assigned level is asymptotically least with a strategy φ^* that maximizes the expected value of $\log(B_n(\varphi))$. Furthermore, if φ^* is a strategy that does maximize $E(\log(B_n(\varphi)))$, while φ is a strategy "essentially different" from φ^*, then $\lim_n B_n(\varphi^*)/B_n(\varphi) = \infty$ a.s. Stanford probabilist Tom Cover compares the Kelly rule's variety of extremal properties with the way pi keeps coming up in contexts that have nothing to do with circles. "When something keeps turning up like that," he suggested to Poundstone, "it usually means it's fundamental."

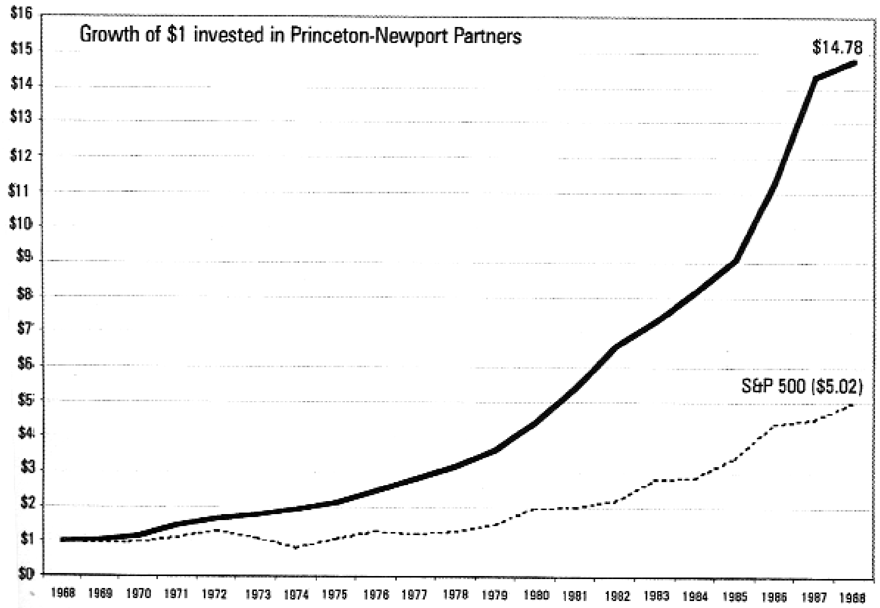

Figure 2. *Performance of the PNP fund. From* Fortune's Formula.

The next chapter in the story of the Kelly criterion concerns Edward Thorp, the first to devise a winning strategy for casino blackjack. He became interested in the game during a Christmas visit to Las Vegas in 1958. On joining the MIT math department the following summer, he realized that the school's new IBM 707 could be just what he needed to realize his plan to devise a winning strategy. After teaching himself Fortran—the new new thing in computer programming at the time—he used it to calculate a card-counting strategy capable of turning a profit at the tables in Vegas. He also concluded that *The Proceedings of the National Academy of Sciences* was the most prestigious journal in which he could hope to publish his work. But for that, he needed a member of the Academy to submit his paper for him.

After learning that Claude Shannon was the only mathematician on the MIT faculty who was also a member of the Academy, he called Shannon's secretary to schedule an appointment. Shannon agreed, in November 1960, that Thorp had significantly advanced the state of the art of blackjack and offered to submit a lightly edited version of his paper to PNAS. He also advised Thorp to read Kelly's paper. Seeing its relevance immediately, Thorp hastened to incorporate a version of the Kelly criterion into the revised version of his blackjack strategy that appeared in book form [4] in 1962.

Thorp made at least a dozen trips to Nevada during the 1960s. After learning to detect cheating dealers and other hazards of casino gambling, he seems to have cleared about $25,000 using his system. He later estimated that he might have earned as much as $300,000 a year had he been able to play full time without interference from management. But that, of course, was not to be. Casino managers soon came to regard card counting as a form of cheating, and to deal accordingly with suspected practitioners. More than one suspected card counter, on declining an invitation to vacate the premises, was summarily beaten up. In any case, Thorp and Shannon soon realized that the stock market represented a much larger casino—one in which $300,000 per annum is considered chump change.

For Shannon, participating in the stock market meant investing his own personal funds and convening regular seminars at MIT on scientific investment. At the rare meetings at which Shannon himself spoke, the meetings had to be moved to one of the largest lecture halls on campus. Poundstone obtained enough of Shannon's personal records to confirm that (starting from next to nothing) his portfolio had grown to more than half a million dollars by 1986, when he reached the age of 70. By his own admission, Shannon had gone through a learning period during which he bought and sold so often that transaction costs ate up the lion's share of his profits. But, having absorbed that lesson, he was subsequently able to increase his net worth at a rate of about 28% per annum. And, although his records are incomplete in the sense that they contain no mention of stocks discarded—a considerable omission for accounting purposes—Poundstone was able to confirm that the estimate is not grossly inaccurate.

For Thorp, in contrast, the record is far more complete. Thorp was quick to yield to requests from friends and relatives that he begin investing their money along with his own. His main vehicle of investment was Princeton–Newport Partners (PNP), a hedge fund with offices in Princeton, New Jersey, and Newport, California (near which Thorp has lived since 1964), between 1969 and 1988. Figure 2 shows the fund's performance during the 19 years of its existence.

It is not only the fact that one dollar invested at the outset grew to $14.78 over the life of the fund (for an APR in excess of 15%, during an era in which the S&P 500 index grew at 8.8%) that arouses the jealousy of money managers everywhere. Nor is it the fact that PNP had to be earning at least 20% on the dollar in order to pay the investors 15%. It is actually the steadiness of the growth that most impresses industry professionals, especially those charged with attracting additional investors to their funds. Nothing is more appealing to investors than rapid, consistent gains. Poundstone does a good job of explaining how PNP's consistency was achieved. That, together with a seemingly endless ability to discover favorable bets lurking in remote corners of the market, have been the most remarkable features of Thorp's lengthy career in money management. The edge/odds formula is worthless without a steady supply of positive edges to exploit.

Thorp is not the only money manager to have used the Kelly criterion successfully. According to Poundstone, Kenneth Griffin's Citadel Investment Group, James Simons's Medallion Fund, and D.E. Shaw and Co. have done so too. Baltimore's legendary William Miller, manager of the Legg Mason Value Trust, is another convert, having written in his 2003 annual report that "The Kelly criterion is integral to the way we manage money." This is significant because Miller's fund is the only SEC-regulated mutual fund—Thorp's various funds and the others mentioned above being unregulated hedge funds—ever to outperform the S&P 500 for 10 consecutive calendar years. Indeed, it has currently done so for 13 consecutive years and seems to be on track for a 14th. Yet Miller suggested to Poundstone that fewer than a tenth of working portfolio managers have ever heard of the Kelly criterion, which—unlike the standard tools of portfolio management—did not arise from the work of Nobel prize-winning economists.

Poundstone is unmistakably eager to publicize the fact that the mainstream approach to portfolio management has a credible rival, about which investors are rarely told. He attributes the apparent suppression of this information to the towering prestige of the Nobel prize-winning economists aligned behind the efficient market hypothesis (EMH) and the orthodox methodology based on it. To Poundstone, and to many of his informants, the Kelly criterion seems to represent a Kuhnian paradigm shift just waiting to happen. What it will take to bring that about remains to be seen.

Poundstone's story goes far beyond the Kelly criterion itself, which, along with the various forms of gambling and investment to which it has been successfully applied, is the focus of this review. At a guess, more than half the text is devoted to the cast of characters involved—however peripherally—in the development and testing of the Kelly/Shannon theory. Many of the scientific names will be familiar to readers of *SIAM News*. Others, including the shady characters who accompanied Thorp on his proof-of-concept trips to Las Vegas, will be entirely unfamiliar. For a book that reads almost like a novel, Poundstone has managed to incorporate a remarkable amount of scientific history and fact.

References

[1] L. Breiman, *Optimal gambling systems for favorable games*, Fourth Berkeley Symp. Prob. and Stat., I (1961), 65–78.
[2] J.L. Kelly, Jr., *A new interpretation of information rate*, Bell System Tech. Jour., (1956), 917–26 .
[3] E.O. Thorp, *The Kelly criterion in blackjack, sports betting, and the stock market*, http://www.bjmath.com/bjmath/thorp/paper.htm, 1997, revised 1998.
[4] E.O. Thorp, *Beat the Dealer: A Winning Strategy for the Game of Twenty-One*, Random House, New York, 1966 (revised version of 1962 edition).

James Case writes from Baltimore, Maryland.