The Journal of Portfolio Management

INVITED EDITORIAL COMMENT

Rational Decision Making under Uncertainty: *Observed Betting Patterns on a Biased Coin*

VICTOR HAGHANI AND RICHARD DEWEY

You're invited to a talk by a former hedge fund manager who was a partner at a fund that famously flopped about 20 years ago. You turn up, hoping to hear some valuable insights or at least some entertaining tales, but instead you are offered a stake of $25 to take out your laptop and bet on the flip of a coin for 30 minutes. You're told that the coin is biased with a 60% probability of coming up heads, and you can bet as much as you like on heads or tails on each flip. You will be given a check for whatever amount is in your account at the end of the half hour.[1]

That's it. Would you feel it was worth your time to play—or would you walk out? How would you play the game? What heuristic or mental tool kit would you employ? These questions led us to conduct this exact experiment. By having participants engage in an activity as simple as flipping a coin, we can easily isolate and observe the betting strategy and its evolution. This simple game also turns out to have similar properties to investing in the stock market, as well as implications for finance and economics education.

We'll describe the experiment, how our subjects played the game, and the conclusions we draw from the experiment.

THE EXPERIMENT

Our coin-flipping experiment was played by 61 subjects, in groups of 2–15, in the quiet setting of office conference rooms or university classrooms. The experiment began when subjects were directed to a URL that contained a purpose-built application for placing bets on the flip of a simulated coin. Prior to beginning, participants read a detailed description of the game, which included a clear statement, in bold, indicating that the simulated coin had a 60% chance of coming up heads and a 40% chance of coming up tails. Participants were given $25 of starting capital, and it was explained both in text and verbally that they would be paid, by check, the amount of their ending balance subject to a maximum payout. The maximum payout would be revealed if and when subjects placed a bet that if successful would make their balance greater than or equal to the cap. We set the cap at $250, ten times the initial stake. Participants were told that they could play the game for 30 minutes; and if they accepted the $25 stake, they had to remain in the room for that amount of time.[2] Participants could place a wager of any amount in their account, in increments of $0.01, and they could bet on heads or tails.

The sample was largely composed of college-age students in economics and finance and young professionals at financial firms. We had 14 analyst and associate-level employees of two leading asset management firms. We expected that these participants would be well prepared to play a simple game with a defined positive expected value.

OPTIMAL STRATEGY

Before continuing with a description of what an optimal strategy might look like, we ask you to take a few moments to consider what you would do if given the opportunity to play this game. Once you read on, you'll be afflicted with the curse of knowledge, making it difficult for you to appreciate the perspective of our subjects encountering this game for the first time. So, if you want to take a moment to think about your strategy, this is the time to do it.

Although it's true that the expected value of the game goes up the higher the fraction the player bets, the outcomes become so skewed that a player who exhibits risk aversion will find an optimal betting fraction well below 100%. If players bet a very high fraction, they risk losing so much money on a bad run that they would not be able to recover; and if they bet too little, they will not be making the most of what is a finite opportunity to place bets at favorable odds.

The strategy with the highest expected value is to bet 100% on heads on every flip. This maximizes the expected return on your bankroll of each flip at 20% (0.6 − 0.4). However, after just 20 flips, the probability that you have not gone bust is one in a million (i.e., $\frac{1}{2}^{20}$). Even betting 50% of your bankroll would be too aggressive for most people. Consider that if you bet 50% of your bankroll on heads for 100 flips and you get exactly the expected 60 heads and 40 tails, you will have lost 97% of your wealth—the result of making 50% 60 times and losing 50% 40 times ($1.5^{60} * 0.5^{40} = 0.033$). This is because heads followed by tails, or vice versa, results in a 25% loss of your bankroll ($1.5 * 0.5 = 0.75$). The odds themselves play a role in the optimal fraction to bet; the more favorable the odds, the higher a fraction one ought to bet.

If you're a professional gambler, chances are you've heard of the Kelly criterion, a formula published in 1955 by John Kelly, a brilliant if somewhat eccentric researcher working at Bell Labs. The formula provides an optimal betting strategy for maximizing the rate of growth of wealth in games with favorable odds, a tool that would appear a good fit for this problem. Dr. Kelly's paper (see Kelly [2012]) built upon work first done by Daniel Bernoulli, who resolved the St. Petersburg Paradox—a lottery with an infinite expected payout—by introducing a utility function that the lottery player seeks to maximize. Bernoulli's work catalyzed the development of utility theory and laid the groundwork for many aspects of modern finance and behavioral economics.

Dr. Kelly's paper and the eponymous formula caught the attention of gamblers and investors. It was further developed and applied to casino games and financial markets by Ed Thorp in a series of papers and popular books, most notably *Beat the Dealer* and *Beat the Market* (see Thorp [1969]). Following Kelly and Thorp's initial work, many others, including Murray Gell-Mann have further developed the theoretical foundations, while notable investors such as Warren Buffett, Bill Gross, and James Simons have all reportedly made use of the Kelly formula.

The basic idea of the Kelly formula is that a player who wants to maximize the rate of growth of his wealth should bet a constant fraction of his wealth on each flip of the coin, defined by the function $2 * p - 1$, where p is the probability of winning. The formula implicitly assumes the gambler has log utility.[3]

In our game, the Kelly criterion would tell the subject to bet 20% ($2 * 0.6 - 1$) of his account on heads on each flip. So, the first bet would be $5 (20% of $25) on heads. If the subject won, he'd bet $6 on heads (20% of $30), but if he lost, he'd bet $4 on heads (20% of $20), and so on.

FINDINGS: HOW WELL DID OUR PLAYERS PLAY?

Our subjects did not do very well. Suboptimal betting came in all shapes and sizes: overbetting, underbetting, erratic betting, and betting on tails were just some of the ways a majority of players squandered their chance to take home $250 for 30 minutes play.

Exhibit 1 shows summary statistics for the outcomes of play across all subjects. Only 21% of participants reached the maximum payout of $250,[4] well below the 95% that should have reached it given a simple

Exhibit 1
Subject Performance

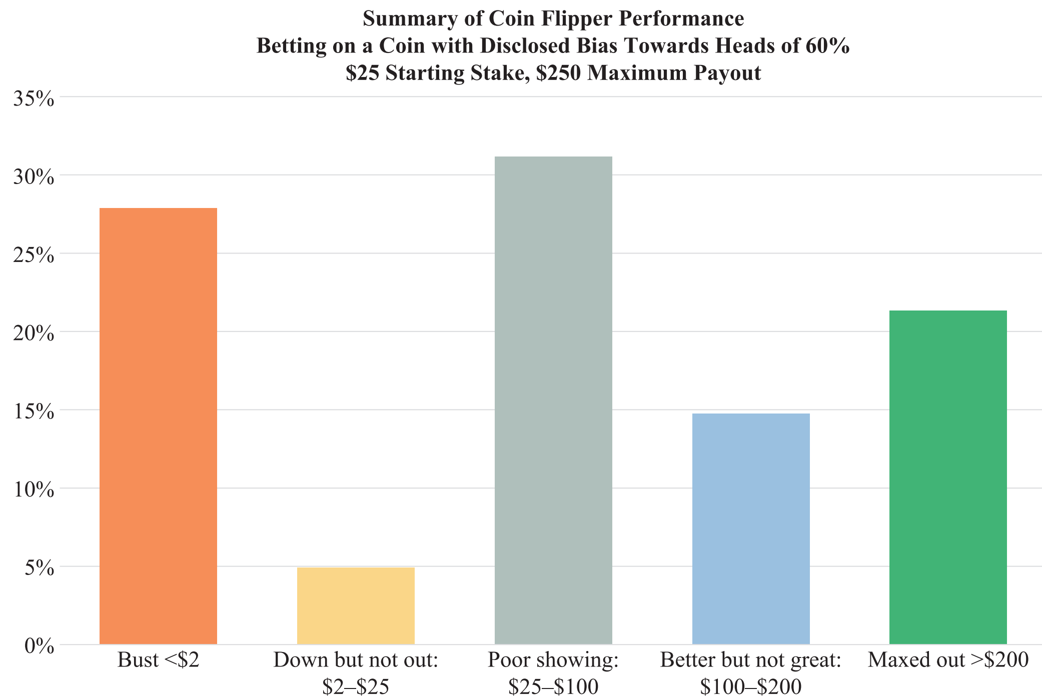

constant percentage betting strategy of anywhere from 10% to 20%.[5]

One-third of the participants wound up with less money in their account than they started with. More astounding still is the fact that 28% of participants went bust and received no payout. That a game of flipping coins with an ex ante 60/40 winning probability produced so many subjects who lost everything is remarkable.

Of the 51% of the sample who did not reach the maximum but did not go bust, the average ending bankroll was $75. While this was a tripling of their initial $25 stake, it still represents a very suboptimal outcome given the opportunity presented. The average payout across all subjects was $91, letting the authors off the hook relative to the $250 per person they'd have had to pay out had all the subjects played well. In aggregate, our subjects wagered on 7,253 coin flips, 59.6% of which came up heads.

Only 5 of our 61 financially sophisticated students and young investment professionals reported that they had ever heard of the Kelly criterion. Interestingly, having heard of Kelly did not seem to help two of them: one barely managed to double his stake, and the other one only broke even after about 100 flips. In post-experiment interviews, we found that the notion of betting a constant proportion of wealth seemed to be a surprisingly unintuitive approach to playing this game. Our results do not offer any indication that participants were converging to optimal play over time, as evidenced by suboptimal betting of similar magnitude throughout the game.

How subjects played the game in the absence of employing Kelly was illuminating. Of the 61 subjects, 18 subjects bet their entire bankroll on one flip, which increased the probability of ruin from close to 0% using Kelly to 40% if their all-in flip was on heads, or 60% if they bet it all on tails, which amazingly some of them did. Hemingway's [2006] description from *The Sun Also Rises* seems apt: "How did you go bankrupt? Gradually, and then suddenly."

The average bet size across all subjects was 15% of the bankroll, so participants bet less, on average, than the Kelly criterion fraction, which would make sense in the presence of a maximum payout that would be within reach. However, this apparent conservatism was completely undone by participants generally being very erratic with their fractional betting patterns—betting too small and then too big. Betting patterns and post-experiment interviews revealed that quite a few participants felt that some sort of doubling down, or martingale betting strategy, was optimal, wherein the gambler increases the size of his wagers after losses. Another approach followed by a number of subjects was to make small and constant wagers, apparently trying to reduce the probability of ruin and maximize the probability of ending up a winner.

We observed 41 subjects (67%) betting on tails at some point during the experiment. Betting on tails once or twice could potentially be attributed to curiosity about the game, but 29 players (48%) bet on tails more than five times in the game. Some of these subjects may have questioned whether the coin truly had a 60% bias toward heads, but that hypothesis is not supported by our finding that the 13 subjects who bet on tails more than 25% of the time were more likely to make that bet immediately after a string of heads. This leads us to believe that some combination of the illusion of control, law of small-numbers bias, gamblers fallacy, or hot-hand fallacy was at work. After the game concluded, we asked participants a series of questions, including whether they believed the coin actually had a 60% bias toward heads. Of those who answered that question, 75% believed that was the case. You can see all the pre- and post-trial questions by playing the game (we won't pay you though) here: coinflipbet.herokuapp.com.

HOW MUCH SHOULD YOU BE WILLING TO PAY TO PLAY?

Not only did most of our subjects play poorly, they also failed to appreciate the value of the opportunity to play the game. If we had offered the game with no cap, this experiment could have become very, very expensive for your authors.[6] If we assume that a player with agile fingers can put down a bet every 6 seconds, 300 bets would be allowed in the 30 minutes of play.[7] The expected gain of each flip, betting the Kelly fraction, is 4%, and so the expected value of 300 flips is $\$25 * (1 + 0.04)^{300} = \$3,220,637$![8]

Given that the expected value of the uncapped game is about $3 million, how much should a person be willing to pay to play this game, assuming that he or she believes that the person offering the game has enough money to meet all possible payouts?[9] Just as is the case with the St. Petersburg paradox, in which players are generally unwilling to pay more than $10 to play a game with an infinite expected value, in our game too, players should only be willing to pay a fraction of the $3 million expected value of the game.[10] For example, if we assume that our gambler has log utility (which the Kelly solution implies) and has de minimis investable wealth, then she should be willing to pay about $10,000 to play the game (the dollar equivalent of the expected utility)—a small fraction of $3 million, but still a very large absolute amount of money in light of the $25 starting stake.[11]

With a capped payout (the game we actually offered), a simple (but not strictly optimal) strategy would incorporate an estimate of the maximum payout. If the subject rightly assumed we wouldn't be offering a cap of more than $1,000 per player, then a reasonable heuristic would be to bet a constant proportion of one's bank using a fraction less than the Kelly criterion; and if and when the cap is discovered, reduce the betting fraction further, depending on the betting time remaining to glide in safely to the maximum payout. For example, betting 10% or 15% of one's account may have been a sound starting strategy.

We ran simulations on the probability of hitting the cap if the subject bet a fixed proportion of wealth of 10%, 15%, and 20% and stopped if the cap was exceeded with a successful bet. We found there to be a 95% probability that the subjects would reach the $250 cap following any of those constant-proportion betting strategies, and so the expected value of the game as it was presented (with the $250 cap) would be just under $240. However, if they bet 5% or 40% of their bank on each flip, the probability of exceeding the cap goes down to about 70%.

Following the initial circulation of our results in a working paper, a number of researchers wrote to us describing work inspired by our experiment. Aaron Brown [2016] solves the optimal betting strategy problem for a game in which the maximum number of flips and maximum payout are known. Vladimir Ragulin explored the certainty-equivalent value of the game under known parameters, different utility functions, and with players who have wealth outside the game, making them less risk averse in terms of how they would play the game. He found that, for example, a player with outside wealth of $1 million and the risk aversion of a Kelly bettor (log utility) should be willing to pay up to $214,000 to play the game for 300 flips and a cap on the payout of $1 million.[12] We wonder whether we could find many players who would be willing to pay such a sum so far in excess of the roughly $10,000 value of the game ignoring outside wealth? Finally, Arjun Viswanathan [2016] provides an analytical solution for a more general case of the game incorporating external wealth.

SIMILARITIES TO INVESTING IN THE STOCK MARKET

This experiment has significant similarities to investing in the stock market. For example, the real return of U.S. equities over the past 50 years was a bit over 5% with an annual standard deviation of about 15%, giving a return/risk ratio of about 0.33. Many market observers believe that the prospective return/risk ratio of the stock market is well below its historical average and closer to that of our coin flip opportunity, which has a return/risk ratio of 0.2.[13]

Of course, there are significant differences—from the binary versus continuous nature of outcomes to the question of risk versus uncertainty when investing in the stock market, where no one can tell you the distribution from which you will draw outcomes. Furthermore, most investors believe that the stock market is not a successive set of independent flips of a coin, but that there are elements of mean reversion and trending in stock market behavior; and, of course, outlier events happen with much higher probability than would evolve from a series of coin flips.[14]

Most people we discussed this with felt that there is a fundamental difference between flipping a coin 300 times in 30 minutes and investing in the stock market where we have to wait 30 years to get 30 flips of the coin. In fact, if stocks follow a random walk, with both return and the risk we care about—variance—growing proportionately with time, then horizon should not affect our betting strategy, although it does affect how highly we value the opportunity to play.[15]

After the experiment, we discussed Kelly and optimal betting strategies with our subjects. We were left with the feeling that they would play the game more effectively if given another chance. We wonder whether any long-lasting impact could be had on investor behavior through similar discussions of sensible approaches to stock market investing. Perhaps investing in the stock market is much more nuanced and complex than betting on a biased coin—or perhaps it's easy to stick to a sound, albeit boring, strategy for 30 minutes but impossible to maintain that discipline for 30 weeks, months, or years.

CONCLUSION

The two primary challenges of investing are identifying attractive opportunities and then sizing them. The first might seem more difficult and important than the second, but as some investors can attest (including one of your authors), sizing can be more critical. Getting the sizing decision wrong, particularly when leverage is involved, can lead to disastrous results no matter how great the expertise brought to bear on uncovering investment gems. Nassim Taleb was at least partly referring to the importance of investment sizing when he remarked, "If you gave an investor the next day's news 24 hours in advance, he would go bust in less than a year."

In light of the spectacular losses incurred over the years by some of our brightest minds, perhaps we should not be surprised by the suboptimal betting strategies followed by the vast majority of our subjects, including the fact that 28% of our subjects went bust betting on a coin that they were told was biased to come up heads 60% of the time.

Given that many of our subjects received formal training in finance, we were surprised that the Kelly criterion was virtually unknown among our subjects, nor were they able to bring other tools (e.g., utility theory) to the problem that would also have led them to a heuristic of constant-proportion betting. We found that without a Kelly-like framework to rely upon, our subjects exhibited a menu of widely documented behavioral biases such as illusion of control, anchoring, overbetting, sunk-cost bias, and gambler's fallacy.

We reviewed the syllabi of introductory finance courses and elective classes focused on trading and asset pricing at five leading business schools in the

United States.[16] Kelly was not mentioned in any of them, either explicitly, or by way of the topic of optimal betting strategies in the presence of favorable odds. Could the absence of Kelly be the effect of Paul Samuelson's vocal critique of Kelly in public debate with Ed Thorp and William Ziemba (Ziemba [2015])? If so, it's time to bury the hatchet and move forward.

These results raise important questions. If a high fraction of quantitatively sophisticated, financially trained individuals have so much difficulty in playing a simple game with a biased coin, what should we expect when it comes to the more complex and long-term task of investing one's savings? Given the propensity of our subjects to bet on tails (with 48% betting on tails on more than five flips), is it any surprise that people will pay for patently useless advice, as documented in studies like Powdthavee and Riyanto [2012]? What do the results suggest about the prospects for reducing wealth inequality or ensuring the stability of our financial system?

Our research suggests that there is a significant gap in the education of young finance and economics students when it comes to the practical application of the concepts of utility and risk taking. Our research will be worth many multiples of the $5,574 winnings we paid out to our 61 subjects if it helps encourage educators to fill this void, either through direct instruction or through trial-and-error exercises like our game. As Ed Thorp remarked to us upon reviewing this experiment, "It ought to become part of the basic education of anyone interested in finance or gambling."

ENDNOTES

The authors are grateful for the help of our colleagues, friends, and family members who helped us design and program the experiment, test it out, and then put our findings into this note. In particular, thank you to Samantha McBride, Adam Smith, Jinwoo Baek, Larry Hilibrand, Vlad Ragulin, William Ziemba, Edward O. Thorp, Andrew Morton, Chris Rokos, Jamil Baz, Rob Stavis, Jeff Rosenbluth, James W. White, Jessica Haghani, and Joshua Haghani.

[1]This is subject to a maximum payout that you'll be informed of if you get close.

[2]This was required even if they chose to not play, or did play and went bust or hit the cap.

[3]We present the Kelly criterion as a useful heuristic that a subject could gainfully employ. It may not be the optimal approach for playing the game we presented for several reasons. The Kelly criterion is consistent with the bettor having log utility of wealth, which is a more tolerant level of risk aversion than most people exhibit. On the other hand, the subjects of our experiment likely did not view $25 (or even $250) as the totality of their capital, and so they ought to be less risk averse in their approach to maximizing their harvest from the game. The fact that there is some cap on the amount the subject can win should also modify the optimal strategy.

[4]We define "maxing out" as reaching at least $200 by the end of the game, and we define "going bust" as finishing the game with less than $2 in the player's account.

[5]We calculated this result through Monte Carlo simulation.

[6]In fact, one reason we suspect this experiment was not performed until now is that it is quite an expensive undertaking even with just 60 subjects.

[7]We programmed the coin to be in a flipping mode for about 4 seconds to create some suspense on each flip and also to limit the number of flips to about 300.

[8]Your opening bet, according to Kelly, would be $5 on heads. The expected gain from that flip would be $1, because there is a 60% chance of winning $5 and a 40% chance of losing $5 (i.e. $0.6 * \$5 - 0.4 * \$5 = \$1$). Your capital in the game at the moment you place that bet would be $25, so the expected return on capital would be 4% ($1/$25). Each successive flip of the coin will have that same 4% expected return up until the cap is encountered, or if the subject gets down to $0.04 or less, at which point players can no longer bet 20% of their accounts because we limit the subject to betting $0.01 or more on each flip. And that's just the expected value. If a subject was very lucky and flipped 210 heads and only 90 tails (admittedly, very unlikely), then we'd have owed him about $2 billion!

[9]Of course, this is not realistic because that would be about $14 trillion trillion ($\$25 * 1.2^{300}$). We suspect that not even the Fed, European Central Bank, and Bank of Japan working together could print that much money.

[10]As with the St. Petersburg paradox, much of the high expected value of our game comes from unlikely, but very big, positive outcomes. The skew can also be seen from the fact that the median of the distribution is so much lower than the mean, which is because if you bet 20% of your account and win, you go up to 1.2 of your wealth; and then if you bet 20% of that and lose, you now wind up at $1.2 * 0.8 = 0.96$, or 4% less than what you had. The median outcome of 180 heads ($0.6 * 300$) and 120 tails would produce an outcome of only $10,504 ($\$25 * 1.2^{180} * 0.8^{120}$), much below the $3,220,637 expected value.

[11]For each flip, the expected utility is $0.6 * \ln(1.2) + 0.4 * \ln(0.8) = 0.0201$, and $\exp(0.0201) = 1.02034$, which means that each flip is giving a dollar-equivalent increase in utility of about 2%; and so for 300 flips, we get $\$25 * 1.02034^{300} = \$10{,}504$. This is also the median of the distribution, as per the previous endnote. If we relax the assumption regarding the player having no outside wealth, the amount he or she should be willing pay can be much higher than $10,000.

[12]Alternatively, an investor who is twice as risk averse as a Kelly bettor should be willing to pay up to $167,000.

[13]More precisely, the ratio is 0.204.

[14]Perhaps more nuanced, our coin flip game generates a distribution in which you make or lose a fixed amount on each flip, whereas many people believe that the stock market has more of a lognormal distribution in which the positive flip outcome is greater than the loss from a negative flip. That is, stocks may be characterized by outcomes of e^d and e^{-d}, whereas our coin flip has $1 + d$ and $1 - d$ for outcomes.

[15]There are a number of assumptions in this statement, including that we display constant relative risk aversion, which is a common but certainly not the only representation of risk aversion among classic and modern behavioral models.

[16]These were MIT, Columbia, Chicago, Stanford, and Wharton.

REFERENCES

Brown, A. "Optimal Betting Strategy with Uncertain Payout and Opportunity Limits." 2016, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2947863.

Hemingway, E. *The Sun Also Rises*. New York, NY: Scribner, 2006.

Kelly, J.L., Jr. "A New Interpretation of Information Rate." In *The Kelly Capital Growth Investment Criterion*, edited by L.C. MacLean, E.O. Thorp, and W.T. Ziemba, pp. 25-34. Singapore: World Scientific, 2012.

Powdthavee, N., and Y.E. Riyanto. "Why Do People Pay for Useless Advice? Implications of Gamblers and Hot-Hand Fallacies in False-Expert Setting." Working paper (IZA DP No. 6557), 2012.

Thorp, E.O. "Optimal Gambling Systems for Favorable Games." *Review of the International Statistical Institute*, Vol. 37, No. 3 (1969), pp. 273-293.

Viswanathan, A. "Generalizing the Kelly Strategy." Working note, 2016, https://arxiv.org/abs/1611.09130.

Ziemba, W.T. "A Response to Professor Paul A. Samuelson's Objections to Kelly Capital Growth Investing." *The Journal of Portfolio Management*, Vol. 42, No. 1 (2015), pp. 153-167.

Victor Haghani *is the founder and CEO of Elm Partners in Wilson, WY, and London, U.K.*
victor@elmfunds.com

Richard Dewey *is an independent researcher in New York, NY.*
richdewey@gmail.com