## Aaron Brown

Limit Bets

What is the optimal strategy for the Haghani–Dewey coin-flip experiment?

L ong-Term Capital Management co-founder Victor Haghani, with Richard Dewey, performed a fascinating experiment detailed in "Rational decision-making under uncertainty: Observed betting patterns on a biased coin."[1] Sixty one subjects were recruited among financial professionals and students of economics or finance. They were given $25 and offered the opportunity to bet at even payout on coin flips with 60 percent chance of heads. They could play for half an hour, which meant a maximum of about 300 flips. There was a maximum payout amount, which would be revealed if the subject tried to place a bet that would cause the balance to exceed the maximum if the bet won. The maximum used was $250.

To put it mildly, participants did not do well. Only 21 percent reached the $250 maximum, less than the 28 percent who went broke. The overall average terminal payout was $91. A simple strategy of betting $5 each time gives a 76.59 percent chance of reaching $250, $211.72 average payout, and only 13.17 percent chance of losing everything. The Kelly strategy of betting 20 percent of bankroll each time gives a 94.21 percent chance of reaching $250, $238.04 average payout, and only 0.02 percent chance of going broke (theoretical Kelly never goes broke, but I assumed bets had to be in integer numbers of cents, and that the subject always bet at least $0.01; also, I never bet such that a win would put the bankroll over $250).

I leave it to others to discuss whether the poor showing results from failures in financial education, or unshakeable perverse behavioral tendencies (even among trained professionals), or something else. I'm interested in what the optimal strategy is for this experiment.
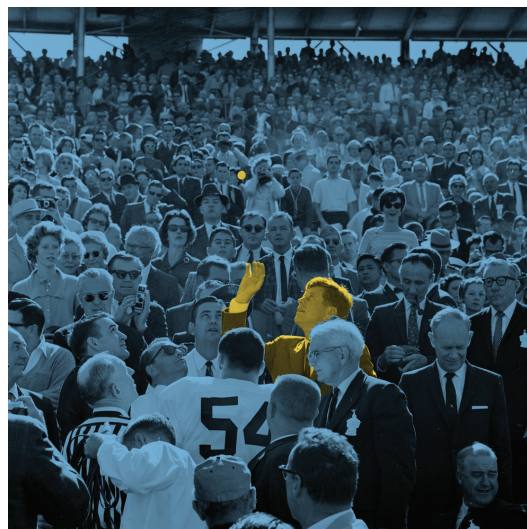
Parameter certainty

Let's begin with a simplified version. Suppose you know the maximum number of flips and the maximum payout. In this case, asking for the optimal bet amount is the wrong way to think about the problem. The right way is to realize that for N flips, you have 2^N possible outcomes, and the only constraint on choosing them is that they average to your initial bankroll. Note that this is a simple average by number of outcomes, not an average weighted by probability. Clearly, you want to assign the highest-probability paths to the maximum outcome, and the lowest-probability paths to a zero outcome (you may need one intermediate-probability path with intermediate results to balance things exactly). Interestingly, this solution does not depend at all on the probability of winning each coin flip, as long as it is greater than 50 percent.

To make this clear with a simple example, suppose you have $6 and five flips, with a maximum payout of $32. There are 32 possible paths, so you want the six highest-probability paths (those with five or four heads) to result in a payout of $32, and the other 26 paths (those with three or fewer heads) to result in zero payout. The simple average payout over paths (ignoring probability) is $6

(six $32 paths and 26 $0 paths), so the solution is feasible. Any other feasible solution will result in a lower expected payout.

If you started with $1 instead of $6, you would assign a $32 payout to the five-head path and a zero payout to all other paths. If you started with $16, you would assign a $32 payout to all paths with five, four, or three heads. These are simple solutions because there is a unique assignment of paths to outcomes that maximizes expected payout. If you started with $7, it's slightly more complicated. You would assign a $32 payout to all five and four-head paths, but you have $1 of "extra" wealth to use on three-head paths. You could assign $32 to one of the three-head paths, or $3.20 to each of the 10 three-head paths, or an infinite number of other combinations.

How do you convert the desired outcomes to bets? The principle is simple. You pair up outcomes and average them. Any set of path pairings results in a feasible betting strategy, but we have to choose carefully if we want the $32 outcomes to be the highest-probability paths. If you start with $6, that means the paths with four or more heads. The chance of getting one of those paths, with 60 percent win probability

Figure 1: Betting paths, five flips, $6 initial wealth, and $32 cap
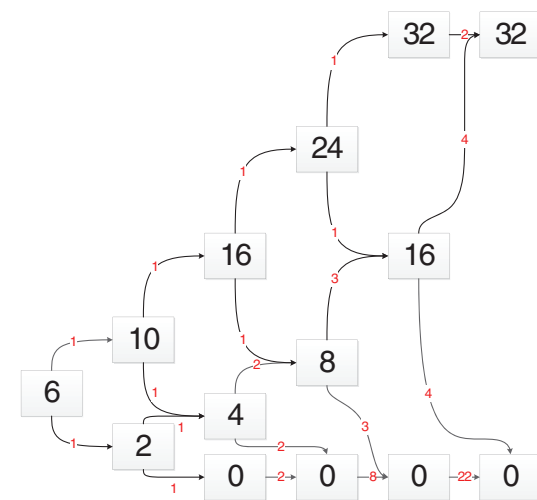
each flip, is 33.70 percent, so our expected payout is $0.3370 \times \$32 = \$10.78$. The betting pattern is the same for any win probability greater than 50 percent, although of course the expected payout changes.

The key is to pair outcomes such that we always pair paths that differ by exactly one in their number of heads.

Figure 1 shows the paths starting from $6 and going to either $0 or $32. The small red numbers show the number of outcomes, taking each branch. For example, you start with $6 and bet $4, creating one set of paths that go to $10 and one set that go to $2.

and so on until the last flip when you bet $16 four times out of 16. The constant average bet is a consequence of the paths recombining, which in turn is a consequence of wanting the final outcome to depend only on the number of heads, not the order in which the heads were realized.

Exploiting this regularity, I can solve for the optimal initial bet by computing the average last-round bet, without drawing the entire tree. I know I want the five-head and one four-head path to result in $32, so if the first four flips are all heads, the last bet is $0. I want to pair the remaining

299 flips, you want to bet $125, so you end up with $250 if you win and zero if you lose. Your average bet is $1.91\% \times \$125 = \$2.38$, and therefore this is your initial bet as well.

If you had instead started with $28.16, you would have wanted to win $250 on any path with 161 or more heads, meaning you would want to bet $125 if you had exactly 160 heads on the first 299 flips, which happens on 2.21 percent of the paths, for an average bet of $2.76. If you start with $25 there is a range of equally good initial bets, but linear interpolation between the lower ($23.01) and higher ($28.16) starting amounts always gives a bet within the optimal range. That would recommend a bet of $2.53 for a starting wealth of $25.

The key is to pair outcomes such that we always pair paths that differ by exactly one in their number of heads

You bet $6 if you win, and $2 if you lose. If you trace it through, there are 32 total paths and the five-head path and the five four-head paths all go to $32; the 26 other paths all go to zero.

Pathbreaking
I constructed this tree backwards. I knew I wanted six $32 outcomes and 26 $0 outcomes after the fifth flip. I knew that the six $32 outcomes included the five-head path and five four-head paths. Since I always pair paths that are exactly one head different, the five-head path gets paired with one of the four-head paths, and the other four four-head paths pair with four of the three-head paths.

The five-head path paired with a four-head path both end up at $32, so they must have $32 and four heads after four flips and bet zero on the fifth flip. The four four-head paths paired with three-head paths must have $16 and three heads after the fourth flip and bet $16, so the winners are four-head paths that go to $32 and the losers are three-head paths that go to $0. All other paths go to zero, so they can be paired up in any manner, it doesn't matter. Continuing to work backward gets me to the initial $6 wealth, and $4 bet.

But you don't have to go through all of this work. Notice that the average bet each round is $4, that is you bet $4 on the first flip, on the second flip you bet $6 half the time and $2 half the time,

four-head paths with three-head paths, so if three of the first four flips are heads, I want to have $32 if the fifth flip is a head (resulting in a total of four heads on the path) and $0 if the fifth flip is a tail (resulting in a total of three heads on the path). If zero, one, or two of the first four flips are heads, I want to have zero whether the last flip is heads or tails. That means I want to bet $16 on four out of the 16 possible paths for the first four flips (the four paths with exactly three heads), for an average bet of $4. Therefore, I bet $4 on the first flip.

With 300 flips and a $250 maximum payout, as in the actual experiment, you have 2.04×10^{90} possible outcomes. If you start with $25, you want 2.04×10^{89} of the paths to result in $250, and the rest to result in zero. You can accomplish this by assigning every path with 162 or more heads to $250, plus 38.58 percent of the paths with 161 heads. There are multiple ways of accomplishing this, as mentioned above, so it's simpler and more instructive to consider the case where you start with $23.01 and want to end up with $250 on exactly those paths with 162 or more heads. In that case, you bet zero if you have 162 or more heads on the first 299 flips (you should already have $250 at that point) or if you have 160 or fewer heads on the first 299 flips (you should already be down to zero at that point). On the 1.91 percent of paths that have exactly 161 heads on the first

Keep it simple
Unfortunately, this solution does not give much insight into what's going on. That, in turn, makes it of limited use in practical situations in which parameters are not known exactly and any analytic result must be blended with other considerations. We have an exact solution to a textbook problem rather than a general bet-sizing principle of broad utility.

A more illuminating approach begins by considering the case where initial wealth (W) is half of the maximum payout (M). If the number of remaining flips (n) is odd, we want to attain the maximum payout on all paths with $\frac{n+1}{2}$ or more heads, which means betting $\frac{M}{2}$ on the last flip if there are exactly $\frac{n-1}{2}$ heads on the first $n-1$ flips and zero otherwise. If n is even, there are multiple solutions, so we'll stick with the simpler case. For large n, the probability of getting exactly $\frac{n-1}{2}$ heads on the first $n-1$ flips is approximately $\sqrt{\frac{2}{\pi n}} \approx \frac{0.8}{\sqrt{n}}$, making the optimal initial bet approximately $\frac{0.4M}{\sqrt{n}}$. This is a general principle, if you have n independent equally advantageous bets (n not too small, say > 5) and a maximum outcome of doubling your initial stake, a bet size equal to 0.4 times the maximum stake (or 0.2 times your initial wealth), divided by the square root of the number of potential bets, is about the right size.

How about for W not equal to $\frac{M}{2}$? Here we take advantage of our knowledge that the optimal bet is zero at $W = 0$ and $W = M$. It happens that the

Figure 2: Expected value of terminal payout vs. initial bet, 300 flips, $25 initial wealth, $250 cap, 60 percent win rate
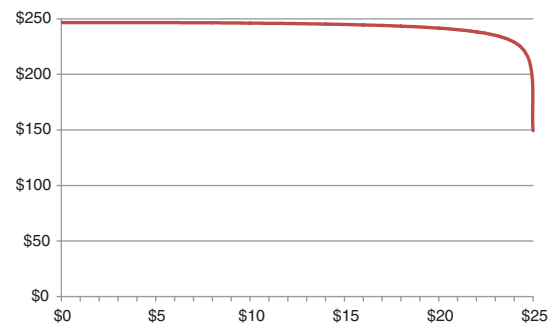
optimal bet is near quadratic in wealth, suggesting a near-optimal bet of $\frac{4\frac{W}{M}\left(1-\frac{ZWZ}{M}\right)0.4M}{\sqrt{n}} = \frac{1.6\left(1-\frac{W}{M}\right)}{\sqrt{n}}W$. For $W = \$25$, $M = \$250$, and $n = 300$, this may be a bit beyond what most people can do easily in their heads, but anyone should be able to figure that the optimal bet is 7.2 percent of wealth at $n = 400$ and

Using Kelly we got an expected payout of $238.04, even though it makes wildly different bets

9 percent of wealth at $n = 256$, so a bet between 8 percent and 9 percent of wealth ($2.00 and $2.25) is probably reasonable for $n = 300$. If you had only 75 flips instead of 300, you'd double your bet. If M shrank to $100, you'd multiply your bet by $\frac{\frac{3}{4}}{\frac{9}{10}} = \frac{5}{6}$. These kinds of quick approximate calculations are more generally useful in practical risk-taking situations than exact complicated formulae.

You don't give up much expected value working with the approximation. For $W = \$25$, $M = \$250$, and $n = 300$, Figure 2 shows the expected value of terminal payout after the first bet for any bet from $0 to $25. You don't give up much value unless you make a really large bet. In fact, any initial bet from $0 to $3.99 has an expected value that rounds to the same cent, and you have to bet more than half your wealth to reduce your expected value by $1.00 or more.

According to Kelly

Using Kelly we got an expected payout of $238.04, even though it makes wildly different bets (for example, it bets $25 with $125, which is much higher than optimal early on, and too small near the end). On the one hand, this is further evidence that we don't need precise strategy calculations. On the other hand, it undersells Kelly. We didn't tell Kelly about either the limit on flips or payout, so it wastes a lot of expected value trying to get payouts above $250, and it doesn't react when the number of remaining flips gets small.

A real Kelly bettor will take these factors into account. One simple way to do that is to choose the bet proportion r that maximizes the chance of getting to $250 by the last flip. One way to state the general Kelly principle is to maximize the long-term median result. In this case, unless you are as perverse as the actual bettors in the experiment, the median outcome will also be the maximum outcome, $250. A reasonable way to adjust the

Kelly solution is to minimize the percentile of the distribution that gets the maximum outcome. I think this is in the spirit of Kelly because it is working in percentiles of long-term outcome, rather than expected return or expected utility.

The exact adjusted Kelly bet requires solving for r such that $\frac{-\ln\left(\frac{W}{M}\right)}{n} = \frac{(1-r)\ln(1-r)+(1+r)\ln(1+r)}{2}$. This isn't hard to solve numerically, and a good approxima-

tion is $r = \sqrt{\frac{-2\ln\left(\frac{W}{M}\right)}{n}}$, which has similarities to the approximation to the exact solution, $\frac{1.6\left(1-\frac{W}{M}\right)}{\sqrt{n}}$. Both have the square root of n in the denominator, the same value at $\frac{W}{M} = 1$, and similar derivatives (−2 vs. −1.6) at that point. Adjusted Kelly gives bets that are somewhat larger than optimal (since the way we set up adjusted Kelly, it assumes the bet proportion doesn't change, so it ignores the option of increasing bet sizes for future flips). However, even with the simplifying assumption and approxima-

tion, this adjusted Kelly gives an expected payout of $245.52, only 0.4 percent below the exact optimal solution.

The point is that it's worth doing math in order to get to near the correct bet size, and to determine a reasonable way to change bets as results become known, but the math doesn't have to be perfect or the calculations precise to get almost all the value. Math helps you understand the question and focus on what matters, the answer is secondary. Many real risk-taking situations have aspects that are captured by the Kelly assumptions, and other aspects that are captured by the exact solution for the capped problem in which you trace out specific paths to accomplish your goal. Knowing both of these solutions and being able to blend them in reasonable ways is more useful in most practical situations than coming up with new exact models.
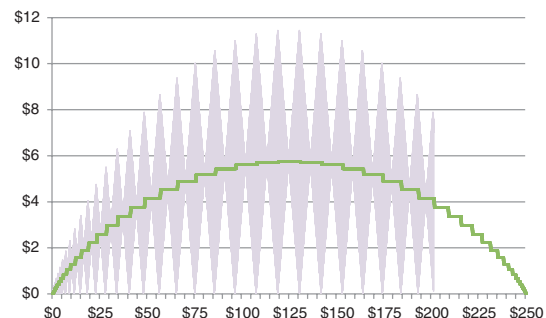
In between the fixed points

Another way to demonstrate that exact solutions may not be necessary is to consider something I have postponed discussing in the prior examples. When I gave the optimal solution for starting with $6 and having five flips to get to a maximum of $32, I picked the $6 carefully. The optimal solution in that case is fully determined, since you want to get to $32 as long as you get four or five heads. That's six paths.

But what if you start with $2? Now you can choose two paths, one of which should be HHHHH, but the other of which can be any of THHHH, HTHHH, HHTHH, HHHTH, or HHHHT. In the first case, you start by betting $0, and then bet all your money each time thereafter. In the last case, you bet all your money on the first four flips and zero on the fifth flip.

You don't even have to pick just one of the five four-head paths. You could bet $0.50 on the first flip. If you lose, bet all your money for the next four flips to end up with $24 if you get THHHH. If you win, bet all your money on the next three flips, then $8 on the last flip. You end up with $32 for HHHHH, $24 for THHHH, and $8 for HHHHT. This is the same expected value as getting $32 for one of the four-head outcomes. So any legal initial bet, from zero to $2, can be optimal.

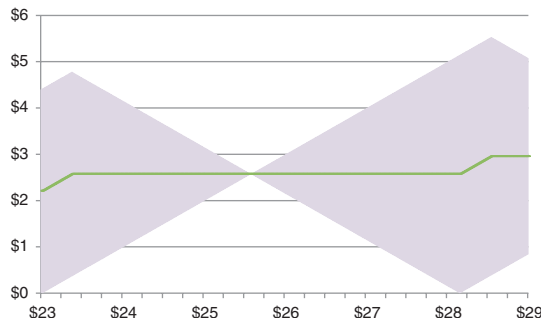**Figure 3: Optimal initial bet vs. initial wealth, 300 flips, $250 cap**

Figure 4: Optimal initial bet vs. initial wealth, 300 flips, $250 cap, close up

Figure 5: Betting paths with five flips, $6 initial wealth, and $32 expected cap

Figure 6: Optimal initial bet vs. initial wealth, five flips, $32 cap and expected cap

If $W = M \begin{pmatrix} n \\ k \end{pmatrix} 2^{-n}$ for some k, then there is a unique optimal bet, but in between those points there is a range of bets that are equally good. If W is at one of the special points, it is possible to get to M with all paths with no more than k tails, but only if you make the precise correct bet. In between the k and $k + 1$ special points, you can get to M with all paths that have no more than k tails, plus some paths with $k + 1$ tails, and there is a range of bets you can select that give you different sets of those $k + 1$-tail paths. At the mid-point between special points, you can bet anything from zero to about the sum of the optimal bets on the two special points.

Figure 3 shows the optimal initial bet for $n = 300$, $M = \$250$, with W on the horizontal axis. Any bet in the shaded purple region can be part of an optimal strategy, the green line is the mid-point of the optimal bet range. However, although there are precise optimal bets for certain values of W, it obviously can't make much difference if you bet quite a bit more or less, since those larger or smaller amounts are exact optimal bets for nearby levels of W. Figure 4 shows a close up for $W = \$29$ to $\$29$. Any bet from $1 to $4 will likely give close to the maximum expected payout for any W in this range.

Unknown M

So far we've assumed that M is known for certain, but in the actual experiment participants did not know it, although they were told that a cap on payouts did exist. So let's assume that participants have a prior probability distribution on M. This actually simplifies the problem over known M, as long as we assume the actual M is revealed only at the end of the experiment and participants are paid the mini-

mum of their final wealth or M. In the actual experiment, participants were told M as soon as they made a bet that would have put them above M if it won, and their bet amount was scaled back to $M-W$.

Let $F(x)$ be the prior probability that the cap is less than x, and $f(x)$ be its derivative. The expected payout if your terminal wealth is W is $\int_0^W xf(x)\,dx + W \int_W^\infty f(x)\,dx$. The derivative of this with respect to W is $Wf(W) - Wf(W) + \int_W^\infty f(x)\,dx = 1 - F(W)$. We maximize expected value by assigning outcomes to paths such that the probability of the path times $F(\text{outcome})$ is constant, for all paths with nonzero outcome.

This is particularly simple if $F(x) = 0$ for $x < S$, where S is the initial endowment and $F(x) = 1 - e^{-a(x-s)}$ for $x > s$ and some a. For easy comparisons to other parts, I'm going to write $a = \frac{1}{M-s}$ so the expected value of the cap equals M. That makes $F(x) = 1 - e^{-(x-s)/(M-s)}$ and we want the probability of each path that ends at terminal wealth W to be inversely proportional to $e^{-(W-s)/(M-s)}$. Since each path has 1.5 times the probability of the path with one fewer head, we want it to result in an outcome $\ln(1.5) \times (M-s)$ higher payout.

Let's start with $n = 5$, $s = \$6$, and $M = \$32$ (remember, M is now the expected value of the cap under the participant's prior distribution, not the actual cap). We want each nonzero outcome to be $10.54 below the outcome of a path with one more head. We can guess that three head outcomes have the lowest nonzero outcome and call it C, so the total of all 32 outcomes will be $16 \times \$0 + 10 \times \$C + 5 \times (\$C + \$10.54) + 1 \times (\$C + \$21.08) = 16 \times \$C + \$73.78$. This must equal $32 \times \$6 = \$192$, so $C = \$7.38$. If we
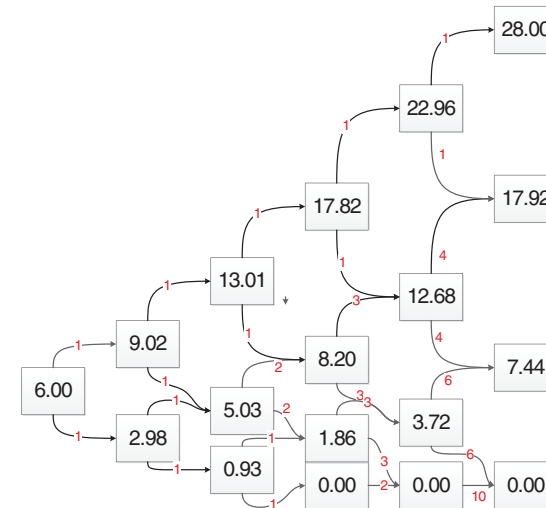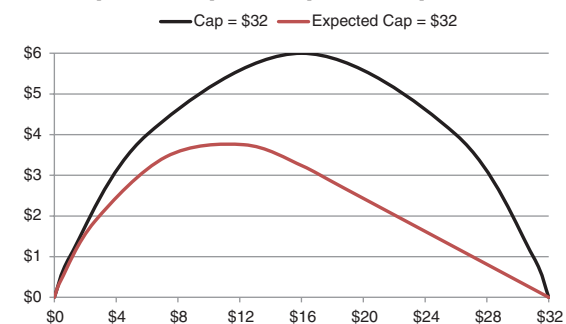
had gotten a $C < \$6$ or $C > \$6 + \$10.54$, we would have had to guess a different number of heads for the lowest nonzero outcome.

Therefore, we want to have $7.38 if we get three heads out of five, $17.92 if we get four heads, and $28.46 if we get five heads. Actually, that doesn't quite work due to rounding. If we instead use $7.44, $17.92, and $28.00, which has nearly identical expected value, we'll get no fractional cents in any of the betting, as shown in Figure 5. Uncertainty costs us $2.00. When we knew the cap was $32, our expected payout was $10.78; under an uncertain cap, with a mean of $32, our expected payout is $8.78.

What's the bet?

We can use the same trick as before to compute the initial bet without deriving the entire tree. If we

Figure 7: Optimal initial bet vs. initial wealth, 300 flips, $250 cap and expected cap

Figure 8: Optimal initial bet vs. initial wealth, with win rate adjusted so known cap is similar to unknown cap
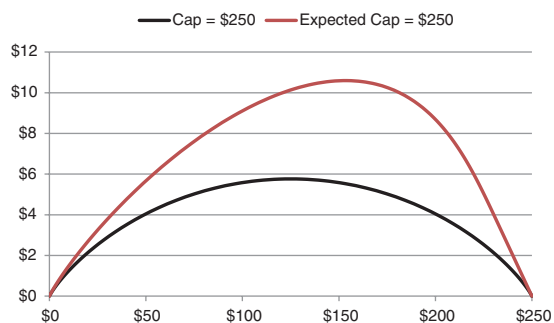
have three or four heads among the first four flips, we want to bet $10.54/2 = $5.27 on the final flip. If we have two heads among the first four flips, we want to bet $7.38/2 = $3.69. That means our average bet is $\frac{5 \times \$5.27 + 6 \times \$3.69}{16} = \$3.03$ (this differs from the $3.02 initial bet in Figure 5 because it does not round bets to the nearest cent).

Figure 6 compares the optimal bet for $n = 5$, $M = \$32$ for the certainty case, where we know the cap is $32, and the uncertainty case, where the expected value of the cap under our prior distribution is $32.
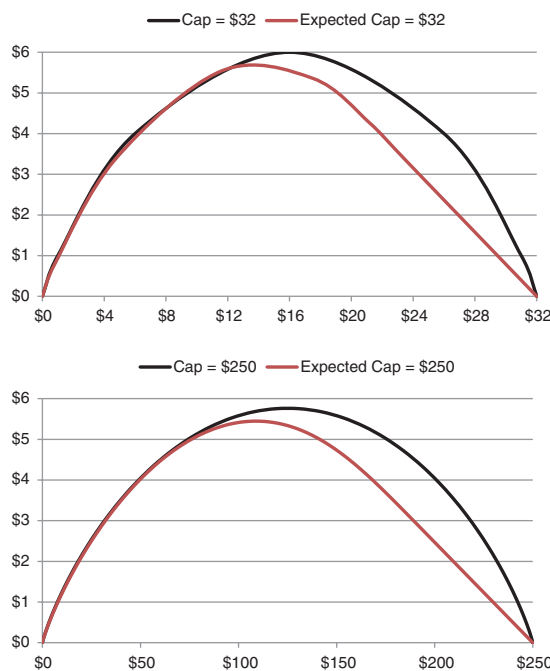
We could do the case with $s = \$25$, $M = \$250$, and $n = 300$ the same way. We want the paths with 163 heads to result in terminal W of $27.60, and for each extra head on the path, we want the outcome to be $91.23 higher, up to a terminal W of $12,526.06 if

we win all 300 flips. We can then determine that our first bet should be $3.25. The optimal initial bet for all levels of W is shown in Figure 7.

You will notice that in the five-flip/$32 case, uncertainty made the bets smaller. In the 300-flip/$250 case, uncertainty made the bets larger. In the former case, the uncertainty peak came before the middle (where the initial bet in the certainty case is largest); in the latter case, the uncertainty peak comes after. However, a moment's reflection will recall that the certainty case is the same for all probabilities of winning each flip, while the uncer-

tainty case includes a factor of $ln\left(\frac{p}{1-p}\right)$, where p is the probability of winning each flip. That is, ln(1.5) = 0.4055 for Figures 6 and 7, since the experiment involved coins with a 60 percent chance of heads.

A larger probability of winning would naturally make the bets larger, and while it is not obvious (but can be seen from the equations), it also makes the maximum initial bet happen at a higher W.

Figure 8 shows what happens if we change p so that the initial bets in the uncertainty case are similar to the initial bets in the certainty case. For the five-flip/$32 case, we must raise p from 60 percent to 68.76 percent, while for the 300-flip/$250 case, we must reduce it to 52.47 percent. These values for p come from an empirically derived formula. For a variety of parameter values, it

seems that the initial bet in the uncertainty case is similar to the initial bet in the certainty case if $n\left[0.5 - \sqrt{p\left(1-p\right)}\right] = \frac{1}{\sqrt{30}}$, which has the solution $p = 0.5 + \sqrt{\frac{1}{\sqrt{30}n} - \frac{1}{30n^2}}$. I have not investigated why that works, but there may be some useful intuition there. What's easy to see is why the larger p is, the larger the bets in the uncertainty case. The surprise, if any, is that p does not affect the bet sizes in the certainty case.

We can use the Gaussian assumption to approximate the initial bet in the 300-flip/$250 case. First we have to solve for k, the number of heads out of 300 we need in order to get a nonzero payout. We actually want it in standard deviation terms, $\frac{k-150}{\sqrt{75}} = v$, which must satisfy

$$\ln\left(1.5\right) \times (250 - 25)\left[\sqrt{\frac{75}{2\pi}}e^{-\frac{v^2}{2}} - v(1 - N\left(v\right))\right] = 25.$$

That's a bit messy, but easy to solve numerically with $v = 1.47$ (an approximate closed-form solution is

$$v = \sqrt{-\ln\left(\frac{2\pi s}{(M-s)\times\ln(1.5)\times\left(5\sqrt{3}-1\right)}\right)} = 1.55). \quad 1 - N\left(v\right) = 0.0713$$

is the fraction of paths with nonzero payout, and $\ln(1.5) \times (M-s)(1-N(v)) = 6.50$ should be the difference between winning and losing the first bet, so the first bet should be $3.25, which agrees to the cent with the exact calculation (we get $2.75 with the closed-form approximation above, which is close enough for practical purposes).

Until this point, I've only considered the goal of maximizing expected payout. That's reasonable if we have a cap of $250, as most people (at least most of the financial professionals who played this game) likely have close to linear utility of money from zero to $250. But when we consider larger amounts, that assumption may not be accurate. Fortunately, if we are willing to assume exponential utility (that is, utility equal to $1 - e^{-aW}$ where a is a risk-aversion parameter and W is wealth), then it's easy to incorporate. Since the product of two exponentials is an exponential with the sum of the arguments, we merely have to redefine M above from the expected value of the payoff to the expected utility of the payoff (I'm cheating a little bit, because the cap prior distribution will generally be shifted from the total wealth utility distribution). Making M expected utility gives us an accurate

Making M expected utility gives us an accurate approximate solution for any reasonable smooth utility function

approximate solution for any reasonable smooth utility function.

Temporarily unknown *M*

The last extension I will consider briefly is when the participant is initially uncertain about the cap, but is told the cap if she makes a bet that would cause her wealth to exceed the cap if she wins (and at that point, her bet will be reduced to the cap minus her wealth). We've already solved the problem for bets after she learns the cap. Before she learns the cap, her situation is similar to the case with uncertain cap, except that she has some incentive to bet larger if the sum of wealth plus bet is higher than any previous wealth plus bet, so she can gain information, which increases her expected payout because she can play better knowing the cap. How does this affect optimal bet size?

One immediate observation is that she should always bet her entire wealth on the last flip (if this would result in a win taking her over the cap, the bet will be automatically reduced to avoid that). This is the same as the certainty case.

The five-flip case starting with $6 and an expected cap of $32 is tractable by computer search. As you would expect, the early bets are somewhat higher than the permanent uncertainty case considered above, about $4.07 vs. $3.03 for the first flip and $5.89 vs. $3.99 for the second flip if you win the first. For all other cases, you bet your entire wealth.

The 300-flip/$25/$250 expected cap case in the original experiment is too large for convenient computer search. The complexity results need to optimize decisions for any value of wealth, with either any known minimum cap (for example, if you have $25 and bet $5, and the bet is accepted, you know that the cap is at least $30, even if you lose the bet), or any known cap. This creates path dependencies in the solution. Instead of $\frac{300(300+1)}{2} = 45{,}150$ decisions to consider, you have $300 \times 2^{300} = 6 \times 10^{92}$.

I suspect that the early bets are somewhat larger than in the permanent uncertainty case, but by a few percent rather than the 35 to 50 percent of the five-flip case. If you have not discovered the cap, the increases should – over the permanent uncertainty case – get larger as the remaining flips

dwindle, particularly if *W* is low relative to the minimum possible cap.

I close with one of the concluding thoughts in the original paper:

Given that many of our subjects received formal training in finance, we were surprised that the Kelly Criterion was virtually unknown and that they didn't seem to possess the analytical tool-kit to lead them to constant proportion betting as an intuitively appealing heuristic. Without a Kelly-like framework to rely upon, we found that our subjects exhibited a menu of widely documented behavioral biases such as illusion of control, anchoring, over-betting, sunkcost bias,

and gambler's fallacy. We reviewed the syllabi of introductory finance courses and elective class-es focused on trading and asset pricing at five leading business schools in the United States.[2] *Kelly was not mentioned in any of them, either explicitly, or by way of the topic of optimal bet-ting 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?*[3] *If so, it's time to bury the hatchet and move forward.*

I certainly agree with the sentiment here, but I want to focus on the most important aspect of Kelly for this problem. Samuelson would tell bet-tors to determine their utility function of terminal wealth and find the betting strategy that maximiz-es its expected value. I assume most of the study participants could have solved this as a textbook problem, at least for the easiest version (known number of flips and cap). Clearly they didn't do this, or even try to guesstimate an approximate solution to this problem, because a utility theory solution is a sound betting strategy (a Samuelson

partisan might say it gives the optimal betting strategy).

Why didn't participants use what they had presumably learned in the classroom? One pos-sibility is that they didn't believe it. Another is that it was too complicated, especially if they tried to tackle the most difficult version of temporary uncertainty. What's worrying is that either issue is made worse by real-life risk deci-sions. If you don't believe textbook solutions for simple games with small amounts at stake, why would you believe them for far more complicat-ed and high-stakes gambles like how to invest or whether to take a loan? And if you can't come up with even approximate solutions for

Why didn't participants use what they had presumably learned in the classroom? One possibility is that they didn't believe it

fully specified coin-flip problems, you will be completely at sea with complex problems with uncertain parameters and outcomes, includ-ing considerations that cannot be measured in money.

It's not the precise Kelly computation that's important here, it's having simple, reliable general tools to make reasonable decisions. And just as important, actually trusting and using those tools. The former requires additions to the curriculum, the latter requires practical risk-taking experience. The experiment should be a wake-up call to finance teachers. If they don't heed it, I predict they will have to suffer the embarrassment of foolish risk decisions by many of their students.

ENDNOTES
1. papers.ssrn.com/sol3/papers.cfm?abstract_id=2856963
2. MIT, Columbia, Chicago, Stanford, and Wharton.
3. Ziemba, W.T. 2015. Response to Paul A Samuelson letters and papers on the Kelly capital growth invest-ment strategy. *Journal of Portfolio Management* Fall, 153–167.

W