THE COINFLIPPER'S DILEMMA

by

Steven E. Landsburg

University of Rochester

1. Alice's Dilemma.

Bob has challenged Alice to a coin-flipping contest. If she accepts, they'll each flip a fair coin repeatedly until it turns up tails, earning a score equal to the number of times the coin turns up heads. (Thus if Alice flips $HHHT$ her score is 3.) The high scorer wins, and collects a prize of $\$4^n$ (or, if you prefer, the utility equivalent of $\$4^n$) from the loser, where n is the loser's score.

Thus if Alice flips a heads and Bob flips b heads, she'll receive a payment of

$$P(a,b) = \begin{cases} -4^a & \text{if } a < b \\ 0 & \text{if } a = b \\ 4^b & \text{if } a > b \end{cases} \tag{1}$$

Alice has commissioned two economists to advise her on whether to accept the challenge.

Economist One observes that conditional on Bob's score b taking on the particular value b_0, Alice's expected return is

$$E_a(P(a,b_0)) = \sum_{a=0}^{\infty} \frac{1}{2^{a+1}} P(a,b_0) = 1/2 \tag{2}$$

(A priori, we might have expected this expression to depend on b_0, but it turns out not to.) Thus, no matter what score Bob earns, Alice's expected return is positive. Therefore she should play.

Economist Two observes that conditional on Alice's score a taking on a particular value a_0, her expected return is

$$E_b(P(a_0,b)) = \sum_{b=0}^{\infty} \frac{1}{2^{b+1}} P(a_0,b) = -1/2 \tag{3}$$

Many thanks to Paulo Barelli, Hari Govindan and Asen Kochov for enlightening conversations.

Thus, no matter what score Alice earns, her expected return is negative. Therefore she should not play.

Who's right?

2. The Economists Make Their Cases.

Economist One elaborates thus:

> Look. Suppose you had a perfectly clairvoyant friend who could predict Bob's score with certainty (but isn't allowed to reveal it to you). That friend, knowing Bob's score to be, say, 3 (or maybe 0 or 7 or 12) would use equation (2) to calculate your expected gain and would surely urge you to play. How can it make sense to ignore the advice of a benevolent friend who has better information than you?
>
> Or if you prefer, look at it this way: Suppose Bob flips first. As soon as you learn his score, you know you're going to want to play this game, and you're going to be sorry if you failed to accept the challenge. Surely a policy you know you're going to regret is a bad policy. That proves that if Bob flips first, you should surely accept the challenge. But at the same time, it clearly doesn't matter who flips first, so you should accept the challenge in any event.
>
> To put this yet another way, you can view Bob's score as a *state of the world* over which you have no control. From that point of view, playing is a *dominant strategy* — in every state of the world, it beats not-playing. Any good game theorist will tell you that when you've got a dominant strategy, you should surely use it.

This makes good sense to Alice. It's true she has no clairvoyant friends, but it seems equally true that clairvoyant friends give good advice, and that *if* she had a clairvoyant friend, this is the advice she'd get.

Unfortunately, Economist Two counters thus:

> Ah, yes — the imaginary clairvoyant friend trick. Let's run with that. Suppoe your clairvoyant friend can predict *your* score with certainty. That friend, knowing your score to be, say, 3 (or maybe 0 or 7 or 12)

> would use equation (3) to calculate your expected gain and would surely urge you not to play. How, indeed, can it make sense to ignore the advice of a benevolent friend who has better information than you?
>
> Or if you prefer, look at it this way: Suppose you flip first (or any time at all before Bob's score is revealed). Then as soon as you learn you're own score, you're going to wish you'd never agreed to play this game — *and you know that in advance*. Surely a policy you're sure to regret is a bad policy. That proves that if you flip first, you should surely reject the challenge. But as my esteemed colleague Economist One has already observed, it clearly doesn't matter who flips first. So you should reject the challenge in any event.
>
> And as for that "dominant strategy" stuff, why don't we try viewing *your* score as the state of the world? In that case, not-playing beats playing in every state of the world, so not-playing is the dominant strategy. I agree with Economist One that if you've got a dominant strategy, you should use it. That's why I think you shouldn't play.

In case this doesn't leave Alice sufficiently confused, Economist Three has just arrived and makes this observation:

> For goodness's sake, this is a zero-sum game, so if the game is good for you then it's bad for Bob. But at the same time, it's a perfectly symmetric game, so if the game is bad for Bob, then it's bad for you. In summary, if the game is good for you then it's bad for you, and by the same argument, if the game is bad for you then it's good for you. The only possible conclusion is that it doesn't matter whether you play or not. Pardon the expresssion, but you might as well flip a coin.

Alice believes that each economist has done an excellent job of explaining why his own argument is right. Unfortunately, none of them has even attempted to explain why the other arguments are wrong.

3. The Source of the Problem.

While Economist One has calculated Alice's expected return conditional on Bob's score, and Economist Two has calculated Alice's expected return conditional on Alice's score, it occurs to Alice that she might gain some insight by calculating her return *unconditionally*.

That is, Alice wants to calculate the value of

$$\sum_{a,b} \frac{1}{2^{a+b+2}} P(a,b) \tag{4}$$

where P is the payoff function defined in (1) and (a,b) runs over all possible pairs of scores (i.e. all possible pairs of non-negative integers.) Unfortunately (4) does not converge. Worse yet, the sum of the positive terms diverges to $+\infty$ while the sum of the negative terms diverges to $-\infty$, so it appears that (4) offers no guidance at all.

Indeed, if the sum (4) were absolutely convergent then the paradox could never have arisen in the first place, because then Fubini's theorem would allow us to interchange the order of summation and write:

$$\sum_{a=0}^{\infty} \frac{1}{2^{a+1}} \sum_{b=0}^{\infty} \frac{1}{2^{b+1}} P(a,b) = \sum_{b=0}^{\infty} \frac{1}{2^{b+1}} \sum_{a=0}^{\infty} \frac{1}{2^{a+1}} P(a,b)$$

which, in the presence of (2) and (3), simplifies to

$$-\frac{1}{2} = \frac{1}{2}$$

Thus if (4) were absolutely convergent, then (2) and (3) could not simultaneously hold.

This might seem to suggest a resolution, namely: Economists should not allow themselves to contemplate payoff functions that violate the hypotheses of Fubini's theorem. But where does this leave Alice, who knows nothing of Fubini's theorem but still has a decision to make?

4. Repeated Plays.

What can Alice expect if she plays this game repeatedly? We'll consider two scenarios.

Scenario One: Suppose first that Bob flips once, generating a score b that Alice repeatedly tries to beat by flipping coins to generate a new score a every day.

In this case, (2) tells us that Alice is playing a game with positive expected value, so the Law of Large Numbers is on her side — if she plays long enough she can be confident of coming out ahead. She might need to be pretty patient though. Although her expected return is 1/2, the variance around that expected return is a whopping

$$\sigma^2 = \frac{4}{7}8^b - \frac{9}{28}$$

where b is Bob's score. Thus if Bob earns a score of, say $b = 4$, Alice finds herself playing a game with expected value 1/2, a standard deviation over 48, and negative outcomes 31 times as likely as positives. It turns out that in order to have even a 50% chance of coming out ahead, she'll have to play at least 69 times — and this number increases extremely rapidly with b.

Scenario Two: Suppose instead that Bob and Alice each flip new scores independently each day. Because this is a symmetric zero-sum game, the distribution of Alice's returns must be symmetric around zero. Alice might therefore dare to hope for some version of the Law of Large Numbers, protecting her from large losses if she plays long enough. Alas, this hope is dashed by the main lemma in Section 3 of [F], from which we can extract the following:

Let A_n be Alice's average return after n plays of the game. Then for small ϵ, the expression

$$Prob(|A_n| < \epsilon)$$

does **not** *approach* 1 *as n gets large.*

In fact, with a bit more work, one can invoke the results of [L] and prove that things are even worse for Alice:

For large M and large N, the expression

$$Prob(|A_n| > M)$$

is approximately equal to K/M where $K \approx 2/3$ is a constant that does not depend on n or M.

In particular, there is no M for which $Prob(|A_n| > M)$ tends to zero as n gets large. Thus Alice cannot use repeated plays to reduce the probability of, say, a $1000 average net loss. Indeed, playing twice as many games renders a \$1000 average net loss just as likely but twice as painful.[1]

5. Resolution, Part I.

To resolve Alice's dilemma, we must first be explicit about what's at stake. Do the payoffs in (1) denote dollars, or do they denote units of utility?

In this section, we'll assume the payoffs are denominated in dollars. Thus the arguments of Economist One and Economist Two are valid only if Alice is an expected *value* (as opposed to expected utility) maximizer.

But why should Alice be an expected value maximizer in the first place? There can be two good reasons to maximize expected value. The first assumes repeated play and appeals to the Law of Large Numbers. But in this case, even if we assume repeated play — with Bob and Alice flipping independently each time — we've seen that the Law of Large Numbers not only fails, but fails in the strong sense that repeated play actually *increases* the probability of a given net total loss. So we can dispose of that reason.

The second good reason to maximize expected value is that one is really maximizing expected utility, and the amounts at stake are sufficiently small that changes in expected utility are well approximated by changes in expected value. This reason applies only if the amounts at stake are small, which they are arguably not, and only if Alice maximizes expected utility, which I will argue in the next section is not a viable assumption.

Thus we've eliminated both of the good reasons for Alice to maximize expected value and therefore rendered both economists' arguments invalid when the payoffs are monetary.

6. Resolution, Part II.

Suppose now that the payoffs in (1) denote units of utility. Then the economists' arguments rest on the assumption that Alice is an expected utility maximizer. But why should we believe such a thing?

The usual answer is that we envision an agent choosing among some set of *lotteries*,

---

[1] Although the result above is stated "for large n", computer simulations strongly suggest that the distribution of A_n looks nearly identical for all values of n from 5 to $5,000,000$. See the appendix for some data.

where a lottery is a probability distribution over some set C. We assume the agent has some preference ordering, we make some assumptions about the properties of that preference ordering, and then we *prove* that the agent is an expected utility maximizer.

There are, of course, innumerable versions of such *representation theorems*, each with its own technical assumptions about the set C, the set $\mathcal{C}$ of allowable probability distributions on which the preference ordering is defined, and the technical properties of the preference ordering. For example, the original vonNeumann-Morgenstern expected utility theorem assumes that each distribution in $\mathcal{C}$ has finite support.

Unless one of those theorems applies, we have no reason to believe Alice is an expected utility maximizer — and therefore no reason to be swayed by the arguments of Economists One and Two.

So in order to take these arguments seriously, we need a theorem, and before we can have a theorem, we need some hypotheses. For the Economists' arguments to work, these hypotheses would have to include at least the following assumptions:

- The set C includes the zero payoff (which Alice can earn by declining to play).
- The set C includes all possible payoffs $P(a, b)$.
- For each fixed b_0, the probability distribution that assigns probability $1/2^{a+1}$ to the outcome (a, b_0) is in $\mathcal{C}$.
- For each fixed a_0, the probability distribution that assigns $1/2^{b+1}$ to the outcome (a_0, b) is in $\mathcal{C}$.

But there can be no representation theorem with these hypotheses, because if there were, it would imply the contradictory conclusions of Economists One and Two.

If we want a representation theorem, then, we have to prohibit Alice from having preferences over some of the lotteries we've considered. This seems a quite unsatisfactory solution, because all of those lotteries are easily implemented as long as a fair coin is available — and it's easy to imagine asking Alice to choose between any two of them.

That leaves the option of acknowledging that we have no representation theorem, hence no reason to believe that Alice is an expected utility maximizer, hence no reason to lend any credence to the arguments of either Economist.

What, then, should Alice do? Should she or should she not accept Bob's challenge?

The answer, of course, is that she should choose whatever she prefers! Presumably she can figure that out for herself. If she can't, no expected utility calculation can help — and we shouldn't expect it to.

Appendix

Let A_n be Alice's average payoff if she accepts Bob's challenge n times. The results of Section 4 say that for large n and large M, the probability that $|A_n| > M$ is approximately constant. The question remains "How large is large?". Computer simulations suggest that $n = 5$ is plenty large, in the sense that the distribution of A_5 appears indistinguishable from the distribution of $A_{5,000,000}$.

Figure 1 shows 100 data points for Alice's average (not total!) return over 5 simulated plays of the game (that is, the computer played five times, computed the average, plotted a point, and repeated this 100 times), then for 50, 500, 5000, and so on up to 5,000,000. Except for a few sporadic outliers, it's hard to discern much difference among these distributions. Figure 2 presents the same data on a different scale that makes it easier to discern the details at the cost of excluding the outliers.

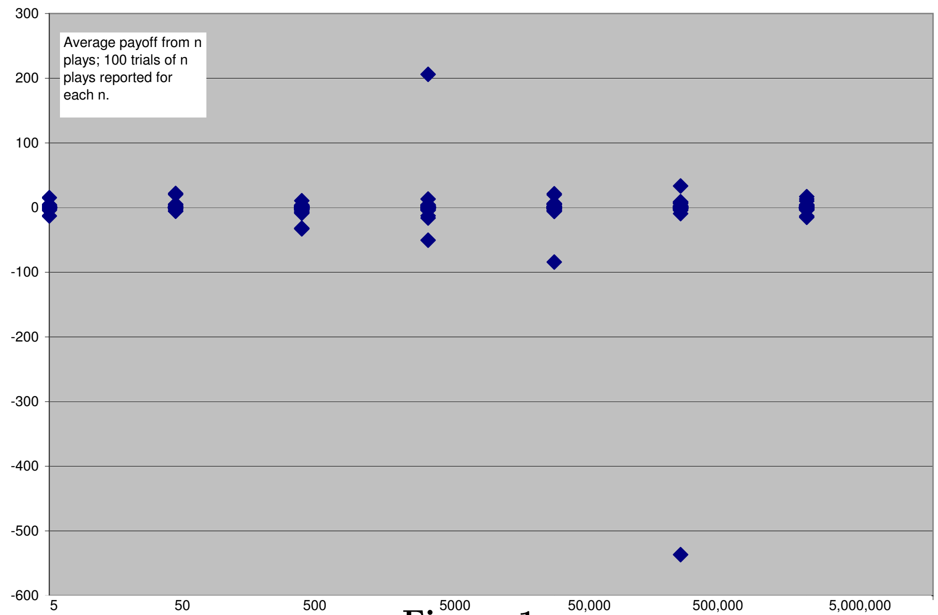

Figure 1

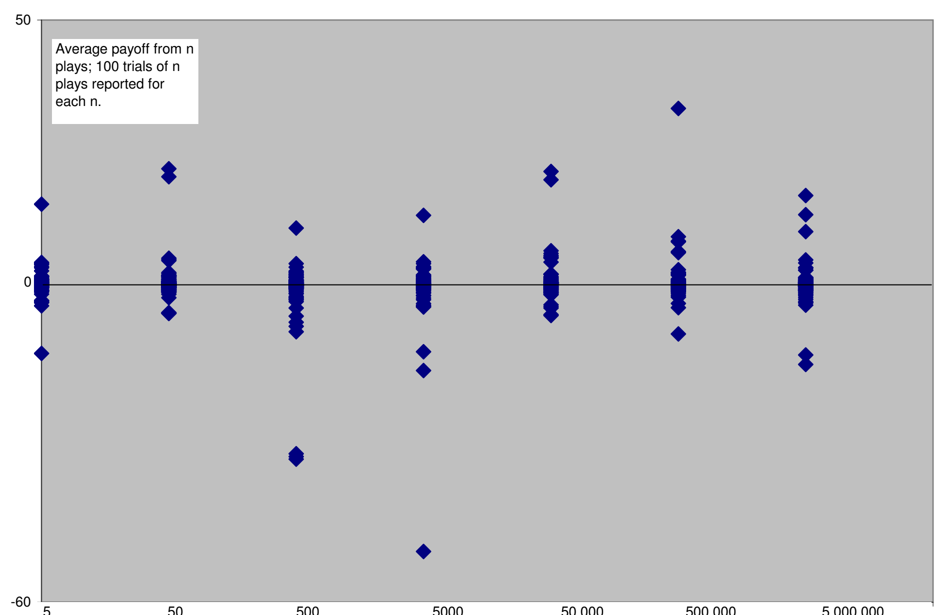

Figure 2

References

[F] W. Feller, "Note on the Law of Large Numbers", *Annals of Mathematical Statistics* 16, 1945.

[L] Lucia, "Mean if i.i.d. Random Variables with No Expected Value", *MathOverflow* 159222, 2014.