ANALYSIS OF THE CHOW-ROBBINS GAME

JON LU

May 10, 2016

Abstract

Flip a coin repeatedly and stop whenever you want. Your payoff is the proportion of heads and you wish to maximize this payoff in expectation. In this paper, we will derive upper and lower bounds for the expected value of this game. We will also examine stopping conditions for several different strategies and expected payoffs for these strategies.

1 Statement of the Game

Consider the following game:

We flip a fair coin repeatedly until we wish to stop. Our payoff for the game is the proportion of heads we've observed up to that point.

This game is a variation of the Robbins's problem of optimal stopping, or Secretary Problem, which was first introduced by Yuan-Shih Chow and Herbert Robbins [2] in 1964.
For the rest of this paper, we will analyze strategies by their expected payoff. Occasionally we will also use the term *expected payoff* when referring to a state of the game, in which case we are referring to the expected payoff under optimal play given we are at that state. In the following sections, we will establish upper and lower bounds on the optimal payoff of this game.

2 Recursive Relation and Lower Bound

Let (h, n) represent our position, where n is the total number of times we have flipped the coin and h is the number of heads we have observed.

From (h, n), we have two choices: stop or continue.
If we stop, we receive h/n in payoff.
If we continue, with $1/2$ probability we end up at $(h, n+1)$ and with $1/2$ probability we end up at $(h+1, n+1)$.
Thus, if we let $V(h, n)$ denote our expected payoff from (h, n), we have the following recursive relation:

$$V(h, n) = \max\left(\frac{h}{n}, \frac{V(h, n+1) + V(h+1, n+1)}{2}\right). \tag{1}$$

At first glance this may seem like a standard recursion but upon closer inspection we notice something strange: each $V(h, n)$ depends on terms with successively larger n, so the "base case" for this recursion is $V(h, \infty)$, which we can't compute!
To get around this problem, we will choose a "horizon" N such that we find lower bounds for all $V(h, N)$ and then recursively apply (1) to get lower bounds for all $V(h, n), n < N$. We can think of N as the cutoff point where we cease to use the recursive formula (1) and as a result don't compute any of the terms $V(h, n), n > N$.
To obtain a lower bound, notice that at position (h, n), we can guarantee h/n by stopping. Moreover, if $h/n < 1/2$, we can wait until the proportion of heads is at least $1/2$ to stop which occurs with probability 1.[1] Thus, we attain the following simple lower bound:

$$V(h, n) \geq \max\left(\frac{h}{n}, \frac{1}{2}\right). \tag{2}$$

While this lower bound holds in general, for smaller values of n it is not very tight and so we will avoid using it when possible. Choosing a larger horizon N to apply the lower bound from (2) on gives a tighter bound on each of the $V(h, n)$s after recursing. As N approaches infinity, the lower bound we obtain for each $V(h, n)$ approaches the true value. Using $N = 10^7$, Häggström and Wästlund [1] compute the following lower bound:

$$V(0, 0) > 0.79295301 \tag{3}$$

3 Upper Bound on $V(h, n)$

The following section is motivated by Häggström and Wästlund [1].

Let $\tilde{V}(h, n)$ be the expected payoff from position (h, n) under *infinite clairvoyance*,

[1]For a rigorous proof, see Theorem 4.2.

i.e., we have complete knowledge of the results of all future coin flips and $\tilde{V}(h,n)$ is the largest proportion of heads ever attained.

Since we know that $V(h,n) \leq \tilde{V}(h,n)$, any upper bound on $\tilde{V}(h,n)$ is necessarily an upper bond on $V(h,n)$. In order to establish an upper bound on $\tilde{V}(h,n)$, we first prove the following:

Let $P(h,n,p)$ be the probability that $\exists\, m, m \geq n$ such that the proportion of heads in the first m flips exceeds p. $P(h,n,p)$ can also be thought of as the probability that starting from (h,n), at some point now or in the future the overall proportion of heads exceeds p.

Let $k = k(h,n,p)$ be the minimum number of coin flips from (h,n) for which it is possible to achieve an overall proportion of heads exceeding p in k more flips. For example, $k(1,3,0.5) = 2$ and $k(2,3,0.5) = 0$.

Lemma 3.1.

$$P(h,n,p) \leq \frac{1}{(2p)^k}$$

Proof. One can check that if $p \leq \max(a/n, 1/2)$, the statement holds. If $p < a/n$ or $p \leq 1/2$, the RHS $1/(2p)^k \geq 1$. The case of $p = a/n$ is left as an exercise for the reader.

Consider Pr(next k flips are heads | proportion p will eventually be exceeded). On one hand,

$$\begin{aligned}
&Pr(\text{next } k \text{ flips are heads} \mid \text{proportion } p \text{ eventually exceeded}) \qquad (4)\\
=&\frac{Pr(\text{next } k \text{ flips are heads } \cap \text{ proportion } p \text{ eventually exceeded})}{\text{proportion } p \text{ eventually exceeded}}\\
=&\frac{Pr(\text{next } k \text{ flips are heads })}{\text{proportion } p \text{ eventually exceeded}}\\
=&\frac{(1/2)^k}{P(h,n,p)}, \qquad (5)
\end{aligned}$$

where the second equality holds because the next k flips being heads guarantees that a proportion p of heads will eventually be exceeded.

However, we can also establish a lower bound for this probability. In (4), we condition on the existence of an $m \geq n$ for which the proportion of heads in the first m

flips exceeds p. Since we assumed $p > 1/2$, we know by the Weak Law of Large Numbers that with probability 1 there exists a maximal $m = M$ for which the first m coin flips have a proportion greater than p of heads.

Because we are operating under infinite clairvoyance, we know the value of M and so we also know the exact number of heads H in flips $n+1, n+2, ..., M$. Furthermore, we know that all permutations of those $M - n$ flips containing H heads are equally likely by symmetry.

Consider flip $n + 1$. If $k > 1$, the proportion of heads up to this point is less than or equal to p but at flip M exceeds p so the proportion of heads q_1 among flips $n + 1$, ..., M necessarily exceeds p. Since flip $n + 1$ is just as likely as any of flips $n + 2, ..., M$ of being among the $q_1 > p$ that are heads, the probability that flip $n+1$ is heads is $q_1 > p$.

Similarly, consider flip $n + 2$. Conditioned on flip $n + 1$ being heads, the proportion of heads up to this point is still at most p but exceeds p at flip M. If we let q_2 be the proportion of heads among flips $n + 2, ..., M$, we know $q_2 > p$ and consequently the probability of flip $n + 2$ being heads given flip $n + 1$ is heads is greater than p.

Continuing this argument for $i = 3, ..., k$, we have that flip $x+i$ has a probability greater than p of being heads given flips $x + 1, ..., x + i - 1$ were heads. Thus, the conditional probability of flips $x + 1, ..., x + k$ being heads

$$Pr(\text{next } k \text{ flips are heads} \mid \text{proportion } p \text{ eventually exceeded}) \geq p^k. \tag{6}$$

Combining (5) and (6) and rearranging gives the desired result.

Theorem 3.2.

$$\tilde{V}(h, n) \leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \min\left(\frac{1}{4}\sqrt{\frac{\pi}{n}}, \frac{1}{2 \cdot |2h - n|}\right). \tag{7}$$

Proof. The following derivation is reproduced from Häggström and Wästlund [1].

From the inverse-CDF definition of expected value, we have:

$$\begin{aligned}\tilde{V}(h, n) &= \int_0^1 P(h, n, p) dp \\ &= \max\left(\frac{h}{n}, \frac{1}{2}\right) + \int_{\max\left(\frac{h}{n}, \frac{1}{2}\right)}^1 P(h, n, p) dp\end{aligned}$$

By definition, k satisfies

$$\frac{h + k - 1}{n + k - 1} \leq p$$

which reduces to

$$k \leq 1 + \frac{np - h}{1 - p}. \tag{8}$$

Furthermore, k is the largest integer that satisfies (8) so we have:

$$k \geq \frac{np - h}{1 - p}. \tag{9}$$

Substituting (9) into Lemma 3.1. yields:

$$P(h, n, p) \leq \frac{1}{(2p)^{\frac{np-h}{1-p}}}$$

It follows that

$$\tilde{V}(h, n) \leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \int_{\max\left(\frac{h}{n}, \frac{1}{2}\right)}^{1} \frac{dp}{(2p)^{\frac{np-h}{1-p}}}. \tag{10}$$

Substituting $p = (1+t)/2$ and using the well-known inequality

$$\frac{\log(1+t)}{1-t} \geq t,$$

(10) simplifies to

$$\begin{aligned}
\tilde{V}(h, n) &\leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2}\int_{\max\left(\frac{2h-n}{n}, 0\right)}^{1} \frac{dt}{(1+t)^{\frac{(1+t)n-2h}{1-t}}} \\
&= \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2}\int_{\max\left(\frac{2h-n}{n}, 0\right)}^{1} \exp\left(-\frac{(1+t)n - 2h}{1-t} \cdot \log(1+t)\right) dt \\
&\leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2}\int_{\max\left(\frac{2h-n}{n}, 0\right)}^{1} \exp(-(1+t)tn + 2ht) dt
\end{aligned}$$

Using the substitution $u = t\sqrt{n}$, we obtain:

$$\begin{aligned}
\tilde{V}(h, n) &\leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2\sqrt{n}}\int_{\max\left(\frac{2h-n}{\sqrt{n}}, 0\right)}^{\sqrt{n}} \exp\left(-u^2 + \frac{2h-n}{\sqrt{n}} \cdot u\right) du \\
&\leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2\sqrt{n}}\int_{\max\left(\frac{2h-n}{\sqrt{n}}, 0\right)}^{\infty} \exp\left(-u^2 + \frac{2h-n}{\sqrt{n}} \cdot u\right) du
\end{aligned} \tag{11}$$

Using the substitution $w = u - (2h - n)/\sqrt{n}$,

$$\int_{\frac{2h-n}{\sqrt{n}}}^{\infty} \exp\left(-u^2 + \frac{2h-n}{\sqrt{n}} \cdot u\right) du = \int_{0}^{\infty} \exp\left(-w^2 - \frac{2h-n}{\sqrt{n}} \cdot w\right) dw \tag{12}$$

From (11) and (12), we can check that

$$\tilde{V}(h, n) \leq \max\left(\frac{h}{n}, \frac{1}{2}\right) + \frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty} \exp\left(-u^2 - \frac{|2h-n|}{\sqrt{n}} \cdot u\right) du. \tag{13}$$

Furthermore, by discarding terms from the integral in (13), we have

$$\frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty}\exp\left(-u^2-\frac{|2h-n|}{\sqrt{n}}\cdot u\right)du \leq \frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty}\exp\left(-u^2\right)du = \frac{1}{4}\sqrt{\frac{\pi}{n}} \tag{14}$$

and

$$\frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty}\exp\left(-u^2-\frac{|2h-n|}{\sqrt{n}}\cdot u\right)du \leq \frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty}\exp\left(-\frac{|2h-n|}{\sqrt{n}}\cdot u\right)du = \frac{1}{2\cdot|2h-n|}. \tag{15}$$

Combining (14) and (15), we have

$$\frac{1}{2\sqrt{n}}\int_{\frac{2h-n}{\sqrt{n}}}^{\infty}\exp\left(-u^2-\frac{|2h-n|}{\sqrt{n}}\cdot u\right)du \leq \min\left(\frac{1}{4}\sqrt{\frac{\pi}{n}}, \frac{1}{2\cdot|2h-n|}\right) \tag{16}$$

Substituting (16) into (13) gives the desired result.

To compute an upper bound on $\tilde{V}(0,0)$, we turn to the a similar strategy to the one we used to calculate a lower bound. In particular, we will use Theorem 3.2 to compute upper bounds on terms of the form $\tilde{V}(h,N)$ for some large N, from which we also have upper bounds on the $V(h,N)$s. We will then apply our fundamental recursive relation from (1) to compute upper bounds on all $V(h,n), n < N$. Using $N = 10^7$ gives the upper bound [1]:

$$V(0,0) < 0.79295560 \tag{17}$$

Lastly, combining the lower bound (3) we derived earlier with this upper bound (17) gives us fairly tight bounds on the expected value of the game under optimal play:

$$0.79295301 < V(0,0) < 0.79295560 \tag{18}$$

4 Strategies for Playing the Game

Now that we've derived bounds on the expected payoff of this game, it's worth examining some different strategies and to see how they fare in comparison.

We will find the following results helpful in our analysis of these strategies:

Lemma 4.1. *If we flip a fair coin forever, the probability of at some point obtaining k more heads than tails or k more tails than heads is 1 for all values of k.*

Proof. Let h_n be the number of heads and t_n be the number of tails we observe in the first n flips respectively. Let $d_n = h_n - t_n$ be the heads-tails differential after n flips.

We are trying to show that with probability 1, $|d_n| = k$ for some n. The key realization we make is that if $|d_N| < k$, flipping $2k$ heads or $2k$ tails in a row causes $|d_{N+2k}| > k$, which means for some value of $n \in \{N+1, ..., N+2k-1\}$, $|d_n| = k$. Thus, in order for the condition of $|d_n| < k$ to hold for all n, we can never flip $2k$ heads or $2k$ tails consecutively.
We now consider blocks of $2k$ flips at a time. The probability that we avoid flipping $2k$ consecutive heads or $2k$ consecutive tails within the first $2k$ flips is simply $\frac{2^k-2}{2^k}$. Likewise, this is also true for the next $2k$ flips, the $2k$ flips after that, etc. Thus, the probability that in the first b blocks of $2k$ flips each, no individual block contains all heads or all tails is

$$\left(\frac{2^{2k}-2}{2^{2k}}\right)^b.$$

As b goes to infinity, this expression goes to 0, so the probability that some block contains $2k$ heads or $2k$ tails is 1, completing the proof.

Theorem 4.2. *[Simple Random Walks] If we flip a fair coin forever, the probability of at some point obtaining k more heads than tails is 1 for all values of k.*

Proof. Reusing our notation from Lemma 4.1, we seek to show that with probability 1, $d_n = k$ for some n. From Lemma 4.1, we know that with probability 1, $|d_n| = k$ for some n. By symmetry, d_n is equally likely to reach k or $-k$ first. Thus, with probability 1/2, d_n reaches k first and we are done. With probability 1/2, d_n reaches $-k$ first. However, reapplying this argument, the heads-tails differential starting from the $n+1$-th flip is equally likely to reach $2k$ or $-2k$. With probability 1/2, it reaches $2k$ first, resulting in $d_n = -k+2k = k$, in which case we are done. With probability 1/2, it reaches $-2k$ first, causing $d_n = -k-2k = -3k$, in which case we repeat the argument again. Thus, the probability of $d_n = k$ at some point is:

$$\frac{1}{2} + \frac{1}{2}\cdot\frac{1}{2} + \frac{1}{2}\cdot\frac{1}{2}\cdot\frac{1}{2} + \cdots = 1.$$

This completes the proof.

The first strategy we will consider is a very simple strategy:

Strategy 4.3.1 *If the first flip is heads, stop. If it is tails, flip until the number of heads equals the number of tails.*

Analysis. With probability 1/2, we flip heads and our payoff is 1 and with probability 1/2 we flip tails and eventually receive a payoff of 1/2 (we can guarantee termination because of Theorem 4.2). Thus the expected payoff of this strategy is just

$$E = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}. \tag{19}$$

Notice that this simple strategy actually does reasonably well compared to the bound obtained in (18).

We now consider a more interesting strategy:

Strategy 4.3.2 *Stop when the number of heads first exceeds the number of tails.*

Analysis. We first note that termination is guaranteed since this corresponds to $k = 1$ from Theorem 4.2. We also note that termination will always occur on an odd numbered flip. With this in mind, the expected payoff of this strategy is then

$$E = \sum_{n=0}^{\infty} p_{2n+1} \cdot \frac{n+1}{2n+1},$$

where p_{2n+1} is the probability we terminate on the $2n+1$-th flip and $(n+1)/(2n+1)$ is our payoff from terminating on the $2n+1$-th flip.

We now try to compute p_{2n+1}:

$$p_{2n+1} = \frac{\text{\# of ways to terminate on } 2n+1\text{-th flip}}{\text{\# of possibilities for first } 2n+1 \text{ flips}} \tag{20}$$

Clearly, the denominator of (20) is just 2^{2n+1}. To compute the numerator, notice that the terminating flip will always be a head since if it were tails we would have terminated sooner. Thus, computing the numerator reduces to finding the number of sequences $X_1, ..., X_{2n}$ containing exactly n heads and n tails such that for no $i \in \{1, ..., 2n\}$ does $X_1, ..., X_i$ contain more heads than tails – this last condition prevents early termination.

However, this is equivalent to the number of Dyck paths[2] of length $2n$, which is known to be the n-th Catalan number C_n. Thus, we have:

$$\begin{aligned} E &= \sum_{n=0}^{\infty} p_{2n+1} \cdot \frac{n+1}{2n+1} \\ &= \sum_{n=0}^{\infty} \frac{C_n}{2^{2n+1}} \cdot \frac{n+1}{2n+1} \\ &= \frac{1}{2} \sum_{n=0}^{\infty} \frac{1}{2^{2n}(2n+1)} \binom{2n}{n} \\ &= \frac{1}{2} \arcsin(1) \\ &= \frac{\pi}{4}. \end{aligned}$$

Note that $\pi/4 \approx 0.7854$, so this is remarkably close to the bounds obtained in (18).

Comment: Because of Theorem 4.2, all strategies of the form "stop when the number of heads first exceeds the number of by tails by k" are valid, i.e., terminate with probability 1. It turns out that the strategy above ($k = 1$) has the highest expected value out of all strategies in this family.
Question to the reader: Intuitively why does this make sense?

5 Further Exploration

As of today, no known strategy has an expected payoff within the bounds we derived (18) and so an optimal strategy has not yet been discovered. To my knowledge, the tightest known bounds for the expected value of this game are given by Julian Wiseman [3], who used a larger horizon of $N = 2^{28}$. Possible areas for research would be improving these bounds either through increased computational power or by using a more computationally efficient bounding strategy. Another possible extension of this problem which was suggested by a classmate was to repeat the same analysis we performed but for biased coins with probability $p \neq 1/2$ of landing heads.

[2]Sequences containing equal numbers of up and down steps in which every prefix has at least as many up steps as down steps.

References

[1] Olle Häggtröm and Johan Wästlund, *Rigorous computer analysis of the Chow-Robbins game*, The American Mathematical Monthly 120(10): 893-900 (2013)

[2] Yuan-Shih Chow and Herbert Robbins, *On optimal stopping rules for* s_n/n, Ill. J. Math., 9: 444-454, 1965.

[3] Julian D. A. Wiseman, *The Chow & Robbins Problem: Stop at h=5, t=3*, web page, `http://www.jdawiseman.com/papers/easymath/chow_robbins.html`