When Does a Random Robin Hood Win?

William Gasarch *
University of Maryland at College Park
Evan Golub †
University of Maryland at College Park
Aravind Srinivasan ‡
University of Maryland at College Park

May 25, 2007

Abstract

A certain two-person infinite game (between "Robin Hood" and the "Sheriff") has been studied in the context of set theory. In certain cases, it is known that for any deterministic strategy of Robin Hood's, if the Sheriff knows Robin Hood's strategy, he can adapt a winning counter-strategy. We show that in these cases, Robin Hood wins with "probability one" if he adopts a natural random strategy. We then characterize when this random strategy has the almost-surely-winning property. We also explore the case of a random Sheriff versus a deterministic Robin Hood.

1 Introduction

The Robin Hood game has been studied by logicians [2, 3, 4] in the context of set theory[1]. The key theorems about it state "if the Sheriff

*Dept. of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, U.S.A. gasarch@cs.umd.edu

†Dept. of Computer Science, University of Maryland, College Park, MD 20742, U.S.A. egolub@cs.umd.edu

‡Dept. of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, U.S.A. Supported in part by NSF Award CCR-0208005. srin@cs.umd.edu

[1]The game is not called "the Robin Hood game" in these papers; however, the authors of these papers have called it that in private correspondence.

knows Robin's Strategy then the Sheriff can play a counter-strategy that wins." What if the Sheriff does not know Robin's strategy? How can we formalize this notion? We will have Robin play a simple randomized strategy so that the Sheriff cannot know Robin's next move.

For 2-player matrix games [5] randomized strategies are provably better than deterministic strategies. This paper shows that for a class of infinite games (which are not matrix games) randomized strategies are provably better than deterministic ones. In Section 2 we show exactly when a Random Robin Hood beats any deterministic Sheriff strategy. In Section 3 we show some cases where a Random Sheriff beats a particular deterministic Robin Hood strategies and speculate on other possibilities.

Let ω denote the set of positive integers, and ω_1 denote the first uncountable ordinal.

Def 1.1 Let $r : \omega \rightarrow \omega$ and $s : \omega \rightarrow \omega$ be two functions. Let A be a set. The *Robin Hood Game* with parameters (r, s, A) (henceforth $RH(r, s, A)$) goes as follows.

1. On day i, the Sheriff of Nottingham (henceforth 'the Sheriff') puts a set of $\leq s(i)$ bags of gold into a cave. He labels the bags with elements of A. No label can ever be used twice over the course of the game.
2. On night i, Robin Hood (henceforth "Robin") removes $\leq r(i)$ bags of gold from the cave.

The game goes on forever. If every single bag that ever enters the cave is eventually removed then Robin wins; otherwise the Sheriff wins.

Notation 1.2 If r or s is a constant function we can denote it by that constant. So the phrase "$RH(2, 3, \omega_1)$" makes sense.

Here are some easy facts about the game.

1. If there are no limits on Robin's strategy then he has an easy winning strategy: always remove the bag that has been in the cave the longest.
2. If A is countable then Robin has an easy winning strategy. Before the game begins, Robin fixes an ordering of A which we denote $A = \{a_1, a_2, \ldots\}$. During the game, if Robin is looking at a cave with $\{a_{i_1}, a_{i_2}, \ldots, a_{i_k}\}$ (with $i_1 < i_2 < \cdots < i_k$) then Robin removes a_{i_1}. More succinctly, Robin removes the least indexed bag. Any bag that is placed into the cave will eventually be removed since there are only a finite number of possible bags with lower index.
3. If $(\forall i)$ $[r(i) \geq s(i)]$ then Robin wins easily.

Because of these three facts the game has been studied when Robin's strategy is restricted, A is uncountable, and $(\forall i)[r(i) < s(i)]$.

Def 1.3 A strategy of Robin's is *memoryless* if it depends only on the set of bags of gold (and their labels) that he sees in the cave, and not on when they came in, or how many times Robin has visited the cave.

Here are some theorems about the game from [4]. See [2, 3] for more theorems.

1. Assume Robin has a memoryless strategy. If (r, s) lies in the set $\{(1, 2), (2, 3), (3, 4)\}$ and if the Sheriff knows Robin's strategy then he can devise a counter-strategy that wins $RH(r, s, \omega_1)$. The $(4, 5, \omega_1)$ case is open.
2. Assume Robin has a memoryless strategy. There are models of set theory where, for all constants c, if the Sheriff knows Robin's strategy then he can devise a counter-strategy that wins $RH(c, c + 1, \omega_1)$. (In these models of set theory the axiom of choice is false.)

The key to prior results is that the Sheriff *knew* Robin's strategy. What if he did not? One way to accomplish this is to have Robin adopt a random strategy; it is well-known that random (or *mixed*) strategies can be crucial in games [5]. We show in particular that for any positive constants r and s, and for any infinite set A, Robin wins $RH(r, s, A)$ almost surely if he adopts a natural random strategy. (Note that this is in contrast with the above-seen results on deterministic strategies for Robin.) We use the simplest possible random strategy for Robin:

Def 1.4 The *Random Strategy* for Robin in stage i is to choose $r(i)$ elements at random out of the cave. (If there are less than $r(i)$ elements then he takes them all out of the cave. This cannot happen in our discussion, since, as mentioned above, we will only deal with the case where $r(i) \leq s(i)$ for all i.)

We also characterize when this random strategy has the almost-surely-winning property for Robin.

Clearly, if Robin adopts such a strategy, the choice of A does not matter, as long as it is an infinite set. Also, since Robin plays randomly, there is no advantage for the Sheriff to put in less than $s(i)$ bags, nor for Robin to take out less than $r(i)$ bags. Hence we will assume that in round i the Sheriff puts in exactly $s(i)$ bags and Robin takes out exactly $r(i)$ bags.

We cannot speak of Robin "winning" or "losing" with this strategy; however, we can speak of the probability of winning or losing. This probability is taken over Robin's sequence of coin-tosses. Since we are dealing with an infinite probability space, we need some care in defining Robin's probability of winning. There does not appear to be a unique definition but the following one appears natural. Basically, Robin would like to have removed all bags that arrived in the first n days within some finite time, with high probability; this should hold for arbitrarily large n. Guided by this notion, we define Robin's probability of losing (which is 1 minus his probability of winning) as follows. Let $f : \omega \to \omega$ be an arbitrary function. Define $E_f(n)$ to be the event "there exists some bag that had arrived on or before day n, which had not been removed by night $n + f(n) - 1$". Then, we define Robin's failure probability to be

$$p(r, s) = \inf_{f:\omega\to\omega} \lim_{n\to\infty} \Pr[E_f(n)]. \tag{1}$$

It is not clear that this limit should exist, but we shall show that it exists and is always either 0 or 1 (respectively denoting Robin succeeding and failing almost surely). Let us now understand these two limiting values better. Suppose $p(r, s) = 0$. This means that for any $\epsilon > 0$, there exist $f : \omega \to \omega$ and n_0 such that for all $n \geq n_0$, $\Pr[E_f(n)] \leq \epsilon$; in other words, Robin succeeds almost surely (since the waiting time of bags that arrived by day n is bounded by the finite function $f(n) + n$ almost surely, for any n large enough). Next, suppose $p(r, s) = 1$. This means that for any $\epsilon > 0$, there exists n_0 such that for all $n \geq n_0$ and *for all* $f : \omega \to \omega$, $\Pr[E_f(n)] \geq 1 - \epsilon$; i.e., Robin fails with arbitrarily large probability.

Notation 1.5 Let $L(i) = \sum_{j=1}^{i}(s(i) - r(i))$. This is the number of bags left in the cave after the ith round.

2 When Does a Random Robin Hood Win Beat a Determinisitc Sheriff

?

Our main result, Theorem 2.1, shows when Robin wins or loses, almost surely: basically this is shown to depend on whether the series $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)}$ converges or not. Theorem 2.1 in fact says something stronger. If the series $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)}$ diverges, it is shown that with arbitrarily high probability, there is some f such that $E_f(n)$ fails simultaneously for *all* n. On the other hand, if $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)} < \infty$, then any bag that arrived after a certain day, will stay forever with arbitrarily high probability. (We assume without loss of generality that $L(i) + r(i) > 0$ for all i, so that this series is well-defined. Indeed, $L(i) + r(i) = 0$ is possible, under the condition "$r(j) \leq s(j)$ for all j", only if $s(i) = r(i) = 0$. Such indices i clearly contribute nothing to the game, and can be ignored.)

Theorem 2.1 *Suppose $r(i) \le s(i)$ for all i. Then:*

(a) If $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)} = \infty$, then $p(r,s) = 0$ (i.e., Robin wins almost surely). In fact, in this case, for any $\epsilon > 0$, there exists an f such that $\Pr[\exists n:\ E_f(n)] \le \epsilon$.

(b) if $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)} < \infty$, then $p(r,s) = 1$ (i.e., Robin loses almost surely). In fact, in this case, for any $\epsilon > 0$, there exists n_0 such that for any given bag b that arrived on day n_0 or later, the probability that b is ever removed is at most ϵ.

Proof: Suppose a bag b is in the cave on night i, before Robin makes his random choices. Note that the probability that b is removed on night i is

$$\frac{r(i)}{L(i-1)+s(i)} = \frac{r(i)}{L(i)+r(i)}. \tag{2}$$

(a) Fix some $n \in \omega$, some $\epsilon \in (0,1]$, and some bag b that arrived on some day $d \le n$. By (2), the probability that b is *not* removed by the end of night $n+m-1$ (where m is any positive integer) is

$$\begin{aligned} \prod_{i=d}^{n+m-1} \left(1 - \frac{r(i)}{L(i)+r(i)}\right) &\le \prod_{i=n}^{n+m-1} \left(1 - \frac{r(i)}{L(i)+r(i)}\right) \\ &\le e^{-\sum_{i=n}^{n+m-1} \frac{r(i)}{L(i)+r(i)}}. \end{aligned}$$

Thus, by a union bound, the probability that at least one bag that had arrived on day n or earlier is not removed by the end of night $n+m-1$ is at most

$$\left(\sum_{j=1}^{n} s(j)\right) \cdot e^{-\sum_{i=n}^{n+m-1} \frac{r(i)}{L(i)+r(i)}}. \tag{3}$$

Now since $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)} = \infty$, we also have $\sum_{i=n}^{\infty} \frac{r(i)}{L(i)+r(i)} = \infty$. Therefore, we can choose $m = m(n,\epsilon)$ large enough so that the expression in (3) is at most, say, $6\epsilon/(\pi^2 n^2)$. Thus we see that for all ϵ, there exists f such that

$$\Pr[\exists n:\ E_f(n)] \le \sum_{n=1}^{\infty} \Pr[E_f(n)] \le \sum_{n=1}^{\infty} 6\epsilon/(\pi^2 n^2) = \epsilon.$$

This implies that $p(r,s) = 0$.

(b) In this case, we proceed as follows. Suppose we are given some positive ϵ. Since $\sum_{i=1}^{\infty} \frac{r(i)}{L(i)+r(i)} < \infty$, there exists some integer n_0

such that
$$\sum_{i=n_0}^{\infty} \frac{r(i)}{L(i)+r(i)} \leq \epsilon.$$

Now consider an arbitrary bag that arrived on some day $d \geq n_0$. Then, the probability that it is never removed is

$$\prod_{i=d}^{\infty} \left(1 - \frac{r(i)}{L(i)+r(i)}\right) \geq 1 - \sum_{i=d}^{\infty} \frac{r(i)}{L(i)+r(i)} \geq 1 - \epsilon.$$

This easily implies that $p(r,s) = 1$. ∎

We obtain some easy corollaries of interest.

Corollary 2.2 *Let A be any infinite set.*

1. *If $s \in \omega$ and $s > 1$ then Robin almost surely wins $RH(1, s, A)$. (This follows from the divergence of $\sum_{i=1}^{\infty} 1/i$.)*
2. *If $s(i) = \lceil \log(i+1) \rceil$ then Robin almost surely wins $RH(1, s, A)$. (This follows from the divergence of $\sum_{i=1}^{\infty} \frac{1}{i \log(i+1)}$.)*
3. *Let $\epsilon > 0$. If $s(i) = \lceil (\log(i+1))^{1+\epsilon} \rceil$ then Robin almost surely loses $RH(1, s, A)$. (This follows from the convergence of the series $\sum_{i=1}^{\infty} \frac{1}{i(\log(i+1))^{1+\epsilon}}$.)*
4. *Let $\epsilon > 0$. If $s(i) = i$ then Robin almost surely loses $RH(1, s, A)$. (This follows from the convergence of $\sum_{i=1}^{\infty} \frac{1}{i^2}$.)*

3 When Does a Random Sheriff Beat a Deterministic Robin

We now study a question suggested by one of the referees: what if Robin is deterministic and the Sheriff plays randomly? Concretely, consider $RH(1, s, [0,1])$ for any integer $s \geq 2$. Suppose the Sheriff chooses s *random* labels from $[0,1]$ on each day; with probability 1, no two labels will ever be the same. What are some good *deterministic* strategies for Robin? We are not able to answer this question. We show in Section 3.1 that certain "natural" strategies for Robin do not work, and present a strategy in Section 3.2 that we conjecture to work.

3.1 Negative results

Our negative result is that strategies for Robin such as "always remove the bag with the smallest label (or largest label) currently seen" do not work. More generally, if $s \geq 3$, any strategy that on any given day either removes the minimum-labeled or the maximum-labeled bag, also does not work. We start by recalling the Chernoff lower-tail bound [1]. Suppose X is a sum of a finite number of independent random variables each of which takes values in $[0,1]$. Letting $\mu = \mathbf{E}[X]$ (the expected value of X), the bound shows that for any $\delta \in (0,1)$,

$$\Pr[X \leq \mu(1-\delta)] \leq e^{-\mu\delta^2/2}.$$

Fix any $s \geq 2$ and any $\epsilon \in (0,1)$. We first show that if Robin always removes the minimum-labeled bag, then there is a bag that will never be removed with probability at least $1-\epsilon$. Let $0 < \theta < 1/2$ be any constant such that $2(1-\theta)^2 > 1$; thus,

$$s(1-\theta)^2 > 1. \tag{4}$$

Also define

$$T = \lceil \frac{c_0}{\theta^2} \ln(1/\epsilon) \rceil \tag{5}$$

for a suitably large constant c_0. Our proof approach will be as follows. Let the random variable t be the smallest integer that is at least as large as T, such that the Sheriff placed a bag b with label greater than $1-\theta$ on day t; such a t exists with probability 1. We aim to show that the bag b will never be removed with high probability.

In more detail, define the following random variables:

$$n_0 = t;\ n_1 = \lceil t \cdot (s(1-\theta)^2 - 1) \rceil;\ \text{for } i \geq 1,\ n_{i+1} = \lceil n_i s(1-\theta)^2 \rceil.$$

Call a bag "good" iff its label is at most $1-\theta$. We propose to show that with probability at least $1-\epsilon$, all of the following events happen:

- Event $\mathcal{E}_0$: at least $st(1-\theta)^2$ good bags are placed in days 1 through t;
- Event $\mathcal{E}_i$ for $i \geq 1$: at least n_{i+1} good bags are placed in the n_i days that belong to the interval

$$[1 + n_0 + n_1 + \cdots + n_{i-1},\ 1 + n_0 + n_1 + \cdots + n_i].$$

Since Robin can only remove one bag per night, event $\mathcal{E}_0$ implies that at the end of night t (i.e., the night of the day when the bag b arrives), at least n_1 good bags will be left; these will clearly be removed by Robin before he can remove bag b. However, event $\mathcal{E}_1$ implies that in the n_1 days that follow, at least n_2 good bags arrive. Thus, at the end of night $t + n_1$, at least n_2 good bags will have to be removed before bag b can be removed. $\mathcal{E}_2$ then shows that n_3 good bags arrive in the next n_2 days, etc; thus, if all these events hold, b can never be removed. We now show using the Chernoff bound that the probability of even one of the $\mathcal{E}_i$ not holding is at most ϵ. From now on, let us condition on an arbitrary but fixed value of t, and recall that such a value is at least T; all probabilities from now on will be conditional on this value of t. In particular, all values n_i are deterministic from now on.

The events $\mathcal{E}_i$ for $i \geq 1$ are a bit easier to handle: consider any one such $\mathcal{E}_i$. In any given interval of n_i days, the expected number of good bags placed is $sn_i(1-\theta)$. Thus, the probability of the complement $\overline{\mathcal{E}_i}$ of $\mathcal{E}_i$ can be bounded by the Chernoff bound:

$$\forall i \geq 1, \ \Pr[\overline{\mathcal{E}_i}] \leq e^{-sn_i(1-\theta)\theta^2/2}. \tag{6}$$

We now bound $\Pr[\overline{\mathcal{E}_0}]$. Since (at least) one non-good bag has been known to be placed on day t, we effectively have $st-1$ random label choices in the first t days. Further, we know that in the days from T up to $t-1$, all bags were good; on these days, we thus have s independent random choices of labels, each of which is uniformly distributed in $[0, 1-\theta]$. Thus, even conditional on the value of t, the number X of good bags that arrive in the first t days is a sum of independent binary random variables; also, $\mathbf{E}[X] \geq st(1-\theta) - 1$, which is approximately $st(1-\theta)$ in the relative sense if the constant c_0 is chosen large enough. Thus, if c_0 is large enough, we get

$$\Pr[\overline{\mathcal{E}_0}] \leq e^{-st(1-\theta)\theta^2/3}. \tag{7}$$

Therefore,

$$\Pr[\exists i \geq 0 : \ \overline{\mathcal{E}_i}] \leq e^{-st(1-\theta)\theta^2/3} + \sum_{i \geq 1} e^{-sn_i(1-\theta)\theta^2/2}.$$

Now, we can see from (4) that for any $t \geq T$, the sequence $n_1, n_2, \ldots$ essentially increases geometrically. Thus, the sum on the r.h.s. is dominated by its first term, and the r.h.s. can be made at most ϵ by choosing c_0 large enough. This completes our proof.

For the more general case where on each night, Robin can choose to remove either the minimum-labeled or the maximum-labeled bag, we proceed as follows. Assume $s \geq 3$ now, and consider a bag b with label in the range $(1/2 - \theta/2, 1/2 + \theta/2)$. By essentially the same proof, we can show that bag b will be "swamped" by bags with labels both smaller and larger, and hence will never be removed (with high probability).

3.2 An approach conjectured to work

We now suggest a "maximum likelihood" type scheme that we conjecture will work for Robin. Specifically, suppose we are again dealing with $RH(1, s, [0, 1])$ for some fixed $s \geq 2$. Also suppose that we allow Robin to know the day i. Now suppose D is some known distribution on sets of s distinct labels, each of which lies in $[0, 1]$; furthermore, in any countable number of independent samples from D, we should have with probability 1 that all generated labels are distinct. Consider the case where the Sheriff samples from D independently on each day (for instance, in the discussion above, D is the uniform distribution on $[0,1]^s$), and suppose Robin plays as follows. On each night i Robin examines the labels that he sees, and calculates for each bag b (recursively knowing his strategy for nights 1 through $i-1$, as well as knowing the distribution D) its expected duration of stay so far, Y_b. Then, he chooses to remove the bag with largest Y_b value (breaking ties by a uniformly random choice). We conjecture that this strategy works for Robin in the case where D is the uniform distribution on $[0,1]^s$.

4 Conclusions

In prior papers on the Robin Hood game the results were of the form "for such-and-such settings of the parameters, if Robin plays a deterministic strategy, and the Sheriff knows that strategy, then the Sheriff can win." In this paper we have explored both the assumption that the strategy is deterministic and that the Sheriff knows it. We have shown that even a simple randomized strategy can beat the Sheriff. Hence the key to prior results was that Robin used a deterministic strategy. Hence this paper adds to the literature showing that randomization helps. Also, Section 3 studies a question raised by one of the referees. It will be interesting to resolve the conjecture posed therein; also, if the conjecture is true, can we go further and characterize the distributions D for which it holds? In addition, a study of a random Robin Hood versus a Random Sheriff may be of interest.

Acknowledgements. We thank Marion Scheepers for writing the paper that encouraged this one and for helpful email communication. We thank Srinivasan Parthasarathy for proofreading and commentary. We thank the referees for their helpful comments; we also thank one of the referees for suggesting the problem considered in Section 3.

References

[1] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Math. Stat.*, 23:493–509, 1952.

[2] K. Ciesielski and R. Laver. A game of D. Gale in which one of the players has limited memory. *Periodica Mathematica Hungaria*, 22:153–158, 1990.

[3] M. Scheepers. Variations of a game of Gale (ii): Markov stategies. *Theoretical Comput. Sci.*, 129, 1994.

[4] M. Scheepers and W. Weiss. Variations of a game of Gale (iii): Remainder stategies. *Journal of Symbolic Logic*, 62(4), 1997.

[5] P. Straffin. *Game Theory and Strategy.* MAA, 1993.