Optimal Token Allocations in Solitaire Knock 'm Down

Arthur T. Benjamin
Department of Mathematics
Harvey Mudd College
Claremont, CA 91711
benjamin@hmc.edu

Matthew T. Fluet
Department of Computer Science
Cornell University
Ithaca, NY 14853
fluet@cs.cornell.edu

Mark L. Huber
Department of Statistics
Stanford University
Stanford, CA 94305
mhuber@orie.cornell.edu

Submitted: February 28, 2000; Accepted: August 14, 2000

Abstract

In the game Knock 'm Down, tokens are placed in N bins. At each step of the game, a bin is chosen at random according to a fixed probability distribution. If a token remains in that bin, it is removed. When all the tokens have been removed, the player is done. In the solitaire version of this game, the goal is to minimize the expected number of moves needed to remove all the tokens. Here we present necessary conditions on the number of tokens needed for each bin in an optimal solution, leading to an asymptotic solution.

MR Subject Classifications: primary: 91A60

1 Introduction

Knock 'm Down is a simple game to play that proves surprisingly difficult to analyze. At the beginning of the game, the player allocates t tokens to N bins. An allocation can be described by a vector $X = (x_1, x_2, \ldots, x_N)$ of non-negative integers, with $\sum x_i = t$. At each step of the game, a bin is chosen at random, with the probability of choosing bin i fixed at $p_i > 0$. If at least one token remains in bin i when it is chosen, then one token is removed from the bin. The game continues until all of the tokens have been removed.

In Solitaire Knock 'm Down, the goal is to remove all the tokens in the shortest time possible. Given t tokens, N bins, and a fixed probability vector $P = (p_1, p_2, \ldots, p_N)$, we will say that an initial token allocation $X^* = (x_1^*, x_2^*, \ldots, x_N^*)$ is optimal if it minimizes the expected time needed to remove all of the tokens.

Intuitively, one would expect the optimal token placement to resemble the shape of the probability histogram, so that roughly tp_i tokens are placed in bin i. However, recursive calculations in [3] reveal that even when one can allocate tokens exactly proportional to the probabilities, this allocation is seldom optimal. For example, when $P = (1/3, 2/3)$ and $t = 9$, then $X^* = (2, 7)$. When $t = 1200$ with the same P, $X^* = (393, 807)$. When $P = (.1, .2, .3, .4)$ and $t = 10$, then $X^* = (0, 2, 3, 5)$. In the original version of Knock 'm Down, described in [1], $t = 12$ and P is the "dice total" vector $\frac{1}{36}(1, 2, 3, 4, 5, 6, 5, 4, 3, 2, 1)$. Here $X^* = (0, 0, 1, 2, 2, 3, 2, 1, 1, 0, 0)$.

When $N = 2$ and $P = (p, 1-p)$, it is shown in [2] that with t tokens, the optimal allocation is $X^* = (m, t-m)$, where m is the pth percentile of the Binomial(t, p) distribution. As a consequence, $\lim_{t\to\infty} X^*/t = P$.

In this paper, we generalize these results to obtain necessary conditions for X^* when $N > 2$, and achieve the same asymptotic conclusion.

2 The Token Adding Theorem and Consequences

Given a probability distribution $P = (p_1, \ldots, p_N)$ and token allocation $X = (x_1, \ldots, x_N)$, we define the random variable T_X to be the number of rolls needed to clear allocation X. Consider arbitrary bins a and b, with $p_a < p_b$. Let X_a and X_b be the allocations obtained from X by adding one token to bins a and b respectively. Our first theorem maintains that if x_a/x_b is at least p_a/p_b then X_a has a greater expected clearing time than X_b. More precisely,

Theorem 1. The Token Adding Theorem. *Let X be a token allocation with t tokens such that $p_b x_a \geq p_a x_b$, where $p_a < p_b$. Then $E[T_{X_a}] > E[T_{X_b}]$.*

Proof. To clear allocation X_a, we must first clear suballocation X and then clear the remaining token in bin a, if it has not already been removed. Thus $T_{X_a} = T_X + R_a$, where R_a is the number of extra turns needed to clear the extra token in bin a. R_a is either 0 or a geometric random variable with mean $1/p_a$. Defining R_b analogously, we obtain

$$\begin{aligned} E[T_{X_a}] - E[T_{X_b}] &= \left(E[T_X] + \frac{1}{p_a}Pr(R_a > 0)\right) - \left(E[T_X] + \frac{1}{p_b}Pr(R_b > 0)\right) \\ &= \frac{1}{p_a}Pr(R_a > 0) - \frac{1}{p_b}Pr(R_b > 0). \end{aligned}$$

It remains to show that $\frac{1}{p_a}Pr(R_a > 0) > \frac{1}{p_b}Pr(R_b > 0)$.

Notice that $Pr(R_a > 0)$ is the probability that suballocation X is cleared with bin a chosen exactly x_a times. Thus, we define $A_{n,m}$ to be the set of length n sequences of bin choices such that X is cleared on the n^{th} turn with bin a selected exactly x_a times and bin b selected $x_b + m$ times, where $m \geq 0$. $Pr(A_{n,m}) = \sum_{\alpha \in A_{n,m}} Pr(\alpha)$, where $Pr(\alpha)$ is the product of the probabilities of the values in the sequence. Note that $A_{n,m}$ may be empty for some values of t, n, and m. In particular, $A_{n,m} = \emptyset$ for $n < t$ and for $m > n - t$. Therefore,

$$Pr(R_a > 0) = \sum_{n=t}^{\infty} \sum_{m=0}^{n-t} Pr(A_{n,m}).$$

Similarly, let $B_{n,m}$ be the set of length n sequences of bin choices such that X is cleared on the n^{th} turn with bin b selected exactly x_b times and bin a selected $x_a + m$ times. Thus,

$$Pr(R_b > 0) = \sum_{n=t}^{\infty} \sum_{m=0}^{n-t} Pr(B_{n,m}).$$

Hence to prove our theorem, it suffices to prove the following

Lemma 1. *Let X be a token allocation with t tokens and N values such that $p_b x_a \geq p_a x_b$, where $p_a < p_b$. Then for $n \geq t > 0$ and $0 \leq m \leq n - t$, $\frac{1}{p_a} Pr(A_{n,m}) \geq \frac{1}{p_b} Pr(B_{n,m})$. Further, the inequality is strict in the case $n = t$ and $m = 0$.*

We begin with the case $n = t$ and $m = 0$. Recalling the definitions of $A_{n,m}$ and $B_{n,m}$, we note that both $A_{t,0}$ and $B_{t,0}$ are the set of sequences of values of length n such that X is cleared on the t^{th} turn with exactly x_a a's and x_b b's. Hence $A_{t,0} = B_{t,0}$. Further, since $n = t$, neither $A_{t,0}$ nor $B_{t,0}$ is empty. Since, $p_a < p_b$, then $\frac{1}{p_a} Pr(A_{t,0}) > \frac{1}{p_b} Pr(B_{t,0})$.

Next, consider the case $n > t$ and $m = 0$. Again, we note that $A_{n,0} = B_{n,0}$. Thus, $\frac{1}{p_a} Pr(A_{t,0}) \geq \frac{1}{p_b} Pr(B_{t,0})$. The inequality is strict except when $N = 2$ and $m > t$, where $Pr(A_{n,0}) = Pr(B_{n,0}) = 0$.

Finally, consider the case $n > t$ and $0 < m \leq n - t$.

Let X' be the allocation of $t - x_a - x_b$ tokens with all $x_a + x_b$ tokens removed from bins a and b in allocation X. Let T' be the number of turns needed to clear X' using probability vector $P' = (p'_1, \ldots, p'_n)$, where $p'_a = p'_b = 0$, and $p'_i = p_i/(1 - p_a - p_b)$ for $i \neq a, b$. In other words, T' is the number of non-a and non-b turns needed to clear X'.

We may write $Pr(A_{n,m})$ as the sum of two probabilities:

$$Pr(A_{n,m}) = Pr(A_1) + Pr(A_2),$$

where A_1 is the event that X is cleared on the n^{th} turn with exactly x_a a's and $x_b + m$ b's and the n^{th} token removed is from bin a. A_2 is the event that X is cleared on the n^{th} turn with exactly x_a a's and $x_b + m$ b's and the n^{th} token removed is not in bin a nor bin b. Note that since $m > 0$, the n^{th} token removed can not be from bin b.

Likewise, we may write $Pr(B_{n,m})$ as the sum of two probabilities:

$$Pr(B_{n,m}) = Pr(B_1) + Pr(B_2),$$

where B_1 is the event that X is cleared on the n^{th} turn with exactly $x_a + m$ a's and x_b b's and the n^{th} token removed is from bin b. B_2 is the event that X is cleared on the n^{th}

turn with exactly $x_a + m$ a's and x_b b's and the n^{th} token removed is neither in bin a nor bin b. Since $m > 0$, the n^{th} token removed can not be from bin a.

In order to show that $\frac{1}{p_a}Pr(A_{n,m}) \geq \frac{1}{p_b}Pr(B_{n,m})$, we show that $\frac{1}{p_a}Pr(A_1) \geq \frac{1}{p_b}Pr(B_1)$ and $\frac{1}{p_a}Pr(A_2) \geq \frac{1}{p_b}Pr(B_2)$.

First, we show $\frac{1}{p_a}Pr(A_1) \geq \frac{1}{p_b}Pr(B_1)$. The inequality is obviously true when $Pr(B_1) = 0$, (e.g., when $x_b = 0$). Now when $Pr(B_1) \neq 0$,

$$Pr(A_1) = \binom{n-1}{x_a-1}\binom{n-x_a}{x_b+m} p_a^{x_a} p_b^{x_b+m} Pr(T' \leq n - x_a - x_b - m.)$$
$$Pr(B_1) = \binom{n-1}{x_b-1}\binom{n-x_b}{x_a+m} p_b^{x_b} p_a^{x_a+m} Pr(T' \leq n - x_a - x_b - m.)$$

Then,

$$\begin{aligned}
\frac{\frac{1}{p_a}Pr(A_1)}{\frac{1}{p_b}Pr(B_1)} &= \frac{\frac{1}{p_a}\binom{n-1}{x_a-1}\binom{n-x_a}{x_b+m} p_a^{x_a} p_b^{x_b+m} Pr(T' \leq n - x_a - x_b - m)}{\frac{1}{p_b}\binom{n-1}{x_b-1}\binom{n-x_b}{x_a+m} p_b^{x_b} p_a^{x_a+m} Pr(T' \leq n - x_a - x_b - m)} \\
&= \frac{\frac{(n-1)!}{(x_a-1)!(x_b+m)!(n-x_a-x_b-m)!} p_b^{m+1}}{\frac{(n-1)!}{(x_b-1)!(x_a+m)!(n-x_a-x_b-m)!} p_a^{m+1}} \\
&= \left(\frac{p_b}{p_a}\right)^{m+1} \frac{(x_b-1)!(x_a+m)!}{(x_a-1)!(x_b+m)!} \\
&= \left(\frac{p_b}{p_a}\right)^{m+1} \frac{x_a(x_a+1)\cdots(x_a+m)}{x_b(x_b+1)\cdots(x_b+m)} \\
&= \prod_{k=0}^{m} \frac{p_b(x_a+k)}{p_a(x_b+k)} \\
&\geq 1.
\end{aligned}$$

The last step follows from our theorem's initial assumptions, since

$$p_b(x_a + k) - p_a(x_b + k) = p_b x_a - p_a x_b + k(p_b - p_a) \geq 0.$$

By a similar calculation, when $Pr(B_2) \neq 0$,

$$\begin{aligned}
\frac{\frac{1}{p_a}Pr(A_2)}{\frac{1}{p_b}Pr(B_2)} &= \frac{\frac{1}{p_a}\binom{n-1}{x_a}\binom{n-x_a-1}{x_b+m}p_a^{x_a}p_b^{x_b+m}Pr(T'=n-x_a-x_b-m)}{\frac{1}{p_b}\binom{n-1}{x_b}\binom{n-x_b-1}{x_a+m}p_b^{x_b}p_a^{x_a+m}Pr(T'=n-x_a-x_b-m)} \\
&= \frac{\frac{(n-1)!}{x_a!(x_b+m)!(n-1-x_a-x_b-m)!}p_b^{m+1}}{\frac{(n-1)!}{x_b!(x_a+m)!(n-1-x_a-x_b-m)!}p_a^{m+1}} \\
&= \left(\frac{p_b}{p_a}\right)^{m+1}\frac{x_b!(x_a+m)!}{x_a!(x_b+m)!} \\
&= \left(\frac{p_b}{p_a}\right)^{m+1}\frac{(x_a+1)(x_a+2)\cdots(x_a+m)}{(x_b+1)(x_b+2)\cdots(x_b+m)} \\
&= \frac{p_b}{p_a}\prod_{k=1}^{m}\frac{p_b(x_a+k)}{p_a(x_b+k)} \\
&\geq 1,
\end{aligned}$$

as desired. The lemma and theorem are now established. □

The Token Adding Theorem leads to practical necessary conditions for optimal token allocation. The first corollary establishes a relationship between the ratio of tokens in any two bins and their ratio of probabilities. Specifically,

Corollary 1. *Let $X^* = (x_1^*, \ldots, x_N^*)$ be an optimal allocation of t tokens with probability vector $P = (p_1, \ldots, p_n)$. If $p_a < p_b$, then $p_b(x_a^* - 1) < p_a x_b^*$.*

Proof. Suppose $p_b(x_a^* - 1) \geq p_a x_b^*$. Then by applying the Token Adding Theorem to the $t-1$-token allocation $X = (x_1^*, \ldots, x_a^*-1, \ldots, x_N^*)$, we see that X^* can not be optimal since it has a higher average clearing time than allocation $(x_1^*, \ldots, x_a^*-1, \ldots, x_b^*+1, \ldots, x_N^*)$. □

As an immediate corollary, we have the following intuitive result.

Corollary 2. *Let $X^* = (x_1^*, \ldots, x_N^*)$ be an optimal allocation of t tokens with probability vector $P = (p_1, \ldots, p_n)$. If $p_a < p_b$, then $x_a^* \leq x_b^*$.*

By modifying the proof of the Token Adding Theorem, one can also obtain, as done in [3], the following two corollaries.

Corollary 3. *Let $X^* = (x_1^*, \ldots, x_N^*)$ be an optimal allocation of t tokens with probability vector $P = (p_1, \ldots, p_n)$. If $p_a = p_b$, then $|x_a^* - x_b^*| \leq 1$.*

Corollary 4. *If the bins are arranged so that P satisfies $p_1 \geq p_2 \geq \cdots \geq p_N$, then there exists an optimal allocation X^* such that $x_1^* \geq x_2^* \geq \cdots \geq x_N^*$.*

The four previous corollaries can be used to drastically cut down the number of candidates for optimal allocation. For instance, in the original game of *Knock 'm Down* with $t = 12$ tokens, the candidates for optimal allocation are reduced from $646,646$ to 49.

However, some allocations, such as putting all tokens in the most probable bin, are not eliminated by the previous corollaries. In the next section, we obtain lower bounds for the optimal number of tokens in each bin, resulting in a natural allocation of tokens in the long run.

3 Lower Bounds

As in the previous section, given a configuration $X = (x_1, \ldots, x_N)$ where $\sum x_i = t$, we let T_X denote the first time that all of the tokens have been removed. Then an optimal solution X^* will satisfy $E[T_{X^*}] \leq E[T_X]$ for all configurations X.

Theorem 2. *Let $X^* = (x_1^*, \ldots, x_N^*)$ be an optimal allocation of $t > 0$ tokens, with probability vector $P = (p_1, \ldots, p_N)$. Then for $N \geq 2$, in each bin i, x_i^* is at least as big as the $p_i/(N-1)$ percentile of the Binomial distribution with parameters t and p_i.*

Proof. When $N = 2$, this is proved in [2]. For $N > 2$, we shall prove something slightly stronger. The idea behind our proof is that in an optimal configuration, the probability of choosing from bin i more than x_i^* times before the end of the game should be small. For each bin i, let τ_i be the first time that bin i is chosen $x_i^* + 1$ times. If bin i has been chosen more than x_i^* times by the end of the game, then $\tau_i < T_{X^*}$. Otherwise, when bin i is chosen exactly x_i^* times at the end of the game, we set $\tau_i = \infty$. Note that τ_i cannot equal T_{X^*} since the game cannot end on a "wasted" turn. Summarizing, $\tau_i > T_{X^*}$ if and only if bin i is chosen exactly x_i^* times at the end of the game. Our theorem is based on the following lemma.

Lemma 2. *For $N > 2$, let X^* be an optimal allocation of t tokens. Then for each bin i,*

$$Pr(\tau_i > T_{X^*}) > \frac{p_i}{N-1}.$$

Proof. We shall show the contrapositive. Let X be a configuration with $Pr(\tau_i > T_X) \leq p_i/(N-1)$. We shall construct a new configuration X' with $E[T_{X'}] < E[T_X]$.

Let E_j be the event that a token in bin j was the last one removed in the game. Consider the event $\tau_i < T_X$. When this occurs, all of the tokens in bin i are gone before time τ_i, but at least one token remains in some other bin. In particular, there exists at least one bin j with $Pr(E_j|\tau_i < T_X) \geq 1/(N-1)$ and $x_j \geq 1$. Construct the new configuration X' by setting $x_i' = x_i + 1$, $x_j' = x_j' - 1$ and $x_k' = x_k$ for k different from i and j. In other words, X' is created from X by transferring a token from bin j to bin i.

Since $\tau_i \neq T_X$, we have,

$$E[T_X] = E[T_X|\tau_i < T_X]Pr(\tau_i < T_X) + E[T_X|\tau_i > T_X]Pr(\tau_i > T_X),$$

and

$$E[T_{X'}] = E[T_{X'}|\tau_i < T_X]Pr(\tau_i < T_X) + E[T_{X'}|\tau_i > T_X]Pr(\tau_i > T_X).$$

Now we show how these expectations relate to one another.

Suppose that $\tau_i > T_X$. Then when all of the tokens of X have been removed, one token still remains in bin i under configuration X'. The expected number of additional steps needed to remove this final token is $1/p_i$. Therefore,

$$E[T_{X'}|\tau_i > T_X] = E[T_X|\tau_i > T_X] + \frac{1}{p_i}.$$

Suppose that $\tau_i < T_X$. Then we have chosen bin i at least $x_i + 1 = x'_i$ times, and so no tokens remain in this bin. If any bin other than bin j has the last token under X, then this will also be the last token under X' as well, and $T_X = T_{X'}$. On the other hand, if j is the last token remaining under X, at this point there are no tokens left under X', so $T_{X'} < T_X$. The expected time to remove this final token from bin j is $1/p_j$, so altogether

$$\begin{aligned}
E[T_{X'}|\tau_i < T_X] &= E[T_{X'}|\tau_i < T_X, E_j]Pr(E_j|\tau_i < T_X) \\
&\quad + E[T_{X'}|\tau_i < T_X, \bar{E}_j]Pr(\bar{E}_j|\tau_i < T_X) \\
&= \left(E[T_X|\tau_i < T_X, E_j] - \frac{1}{p_j}\right) Pr(E_j|\tau_i < T_X) \\
&\quad + E[T_X|\tau_i < T_X, \bar{E}_j]Pr(\bar{E}_j|\tau_i < T_X) \\
&= E[T_X|\tau_i < T_X] - \frac{1}{p_j} Pr(E_j|\tau_i < T_X) \\
&\leq E[T_X|\tau_i < T_X] - \frac{1}{(N-1)p_j} \\
&< E[T_X|\tau_i < T_X] - \frac{1}{(N-1)(1-p_i)}.
\end{aligned}$$

Putting it all together, we have that

$$\begin{aligned}
E[T_{X'}] &< \left(E[T_X|\tau_i > T_X] + \frac{1}{p_i}\right) Pr(\tau_i > T_X) \\
&\quad + \left(E[T_X|\tau_i < T_X] - \frac{1}{(N-1)(1-p_i)}\right) Pr(\tau_i < T_X) \\
&= E[T_X] + Pr(\tau_i > T_X)\left(\frac{1}{p_i} + \frac{1}{(N-1)(1-p_i)}\right) - \frac{1}{(N-1)(1-p_i)}.
\end{aligned}$$

Therefore, since $Pr(\tau_i > T_X) \leq \frac{p_i}{N-1}$,

$$\begin{aligned}
E[T_{X'}] - E[T_X] &< \frac{1}{N-1} + \frac{p_i}{(N-1)^2(1-p_i)} - \frac{1}{(N-1)(1-p_i)} \\
&= \frac{(N-1)(1-p_i) + p_i - (N-1)}{(N-1)^2(1-p_i)} \\
&= \frac{p_i(2-N)}{(N-1)^2(1-p_i)} \\
&< 0.
\end{aligned}$$

Thus, $E[T_{X'}] < E[T_X]$, and X is not an optimal solution, establishing our lemma. □

To complete the proof of Theorem 2, suppose that x_i^* was less than the $p_i/(N-1)$ percentile of the Binomial distribution with parameters t and p_i. Then since $T_{X^*} \geq t$,

$$\begin{aligned} Pr(\tau_i > T_{X^*}) &\leq Pr(\tau_i > t) \\ &= Pr(\text{Bin } i \text{ is chosen at most } x_i^* \text{ times among the first } t \text{ turns}) \\ &< p_i/(N-1), \end{aligned}$$

contradicting our lemma. □

Because $p_i/(N-1)$ is a constant, and the binomial distribution concentrates near tp_i for large t, we have as an immediate corollary,

Corollary 5. *Let $X^* = (x_1^*, \ldots, x_N^*)$ be an optimal allocation of t tokens with probability vector $P = (p_1, \ldots, p_n)$. Then $\lim_{t\to\infty} X^*/t = P$.*

4 Acknowledgment

We thank Professor Janet Myhre and The Reed Institute for Applied Mathematics for support of this research.

References

[1] A.T. Benjamin and M.T. Fluet, The Best Way to Knock 'm Down, *The UMAP Journal*, vol 20.1, 11–20, 1999.

[2] A.T. Benjamin and M.T. Fluet, What's Best?, *The American Mathematical Monthly*, vol 107, No. 6, 560–562, 2000.

[3] M.T. Fluet, *Searching for Optimal Strategies in Knock 'm Down*, senior thesis, Harvey Mudd College, Claremont, CA, 1999.