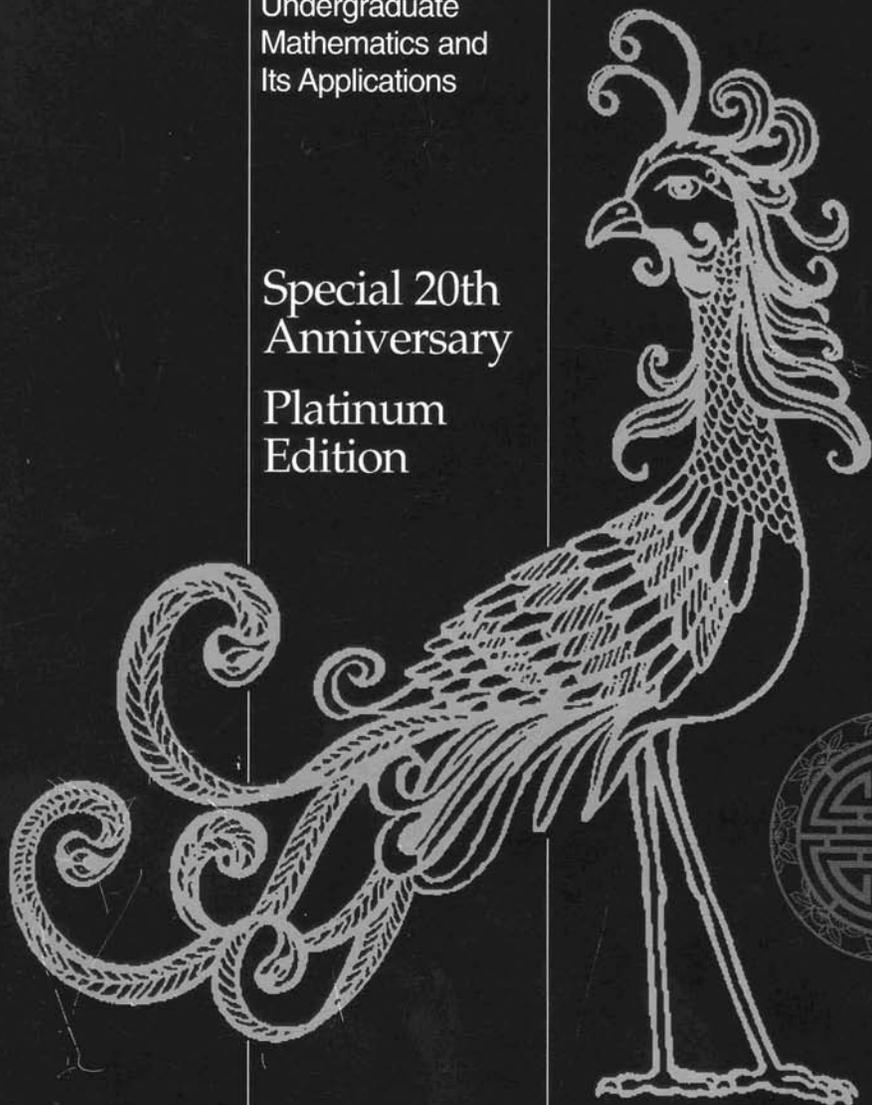The

UMAP
Journal

The Journal of
Undergraduate
Mathematics and
Its Applications

Special 20th
Anniversary

Platinum
Edition

COMAP

Spring 1999

1

The Best Way to Knock 'm Down

Arthur T. Benjamin
Matthew T. Fluet
Dept. of Mathematics
Harvey Mudd College
Claremont, CA 91711
{benjamin, matthew_fluet} @hmc edu

Introduction

"Knock 'm Down" is a game of dice that is so easy to learn that it is being played in classrooms around the world. Although this game has been effective at developing students' intuition about probability [Fendel et al. 1997; Hunt 1998], we will show that lurking underneath this deceptively simple game are many surprising and highly *un*intuitive results.
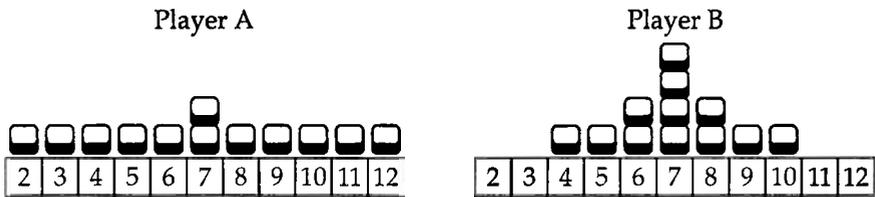

Figure 1. Allocations for Player A and Player B.

The game is played by two players, each of whom is given a 6-sided die, 12 tokens, and a card with the numbers 2 through 12. Each player allocates tokens among the eleven numbers on that player's card. Let's suppose that players A and B allocate their tokens as shown in **Figure 1**. Next, the players roll their dice together and each removes a token from their board on the value equal to the sum of the dice. For instance, in **Figure 1**, if a 6 and a 2 are rolled, then both players remove a token from the 8 spot; but if a 6 and a 5 are rolled, player A removes a token from the 11 spot but—since player B has no tokens on the 11 spot—player B's board is unchanged. The first player to remove all tokens is the winner. (If both players remove their last token on the same roll, then the game is a draw.)

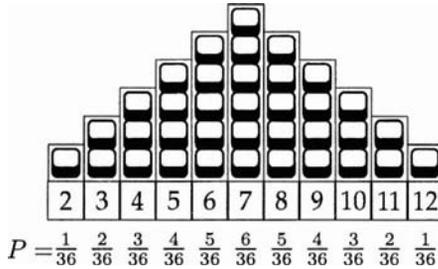In this paper, we investigate the questions:

Who is the favorite to win the game?

Can you find a "better" allocation than the ones in **Figure 1**?

Instinctively, position B seems superior, since it more closely resembles the shape of the histogram of probabilities in **Figure 2**. In fact, B wins against A 75% of the games, draws 9%, and loses only 16%.

$$P = \begin{array}{ccccccccccc} 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 \\ \frac{1}{36} & \frac{2}{36} & \frac{3}{36} & \frac{4}{36} & \frac{5}{36} & \frac{6}{36} & \frac{5}{36} & \frac{4}{36} & \frac{3}{36} & \frac{2}{36} & \frac{1}{36} \end{array}$$

Figure 2. Histogram of probabilities and allocation of 36 tokens.

Given 36 tokens, we could allocate them exactly proportionally to the probabilities (see **Figure 2**). By all that is sensible, we felt that this should be the optimal allocation for 36 tokens. But as we soon learned, in this innocent little dice game, all is not sensible!

Expectations and Results

Before revealing our solution to the original 12-token game, let's find the best allocation for some simpler games. (Here's a hint: The best 12-token allocation can be obtained by moving just one token in player B's allocation in **Figure 1**.)

Consider a 4-valued game consisting of outcomes 1, 2, 3, and 4 with respective probabilities .1, .2, .3, and .4. How should you allocate 10 tokens among the four outcomes? For notational convenience, we call this the *10-token game with* $P = (.1, .2, .3, .4)$. Can you predict which of the two allocations in **Figure 3** is better? Notice that the first allocation has exactly the same triangular shape as the histogram of probabilities.

Player A   Player B

$$P = .1 \quad .2 \quad .3 \quad .4 \qquad P = .1 \quad .2 \quad .3 \quad .4$$
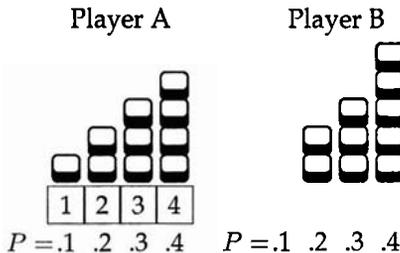
Figure 3. Which allocation is better?

Surprisingly, the answer depends on what you mean by "better." It seems reasonable that we should want the allocation that requires, on average, the fewest number of turns to remove all tokens. Let $A = (1, 2, 3, 4)$, $B = (0, 2, 3, 5)$, and let $E[X]$ denote the average number (i.e., the *expected value*) of the number of rolls needed to clear all the tokens with allocation X. We find $E[A] = 17.7$, whereas $E[B] = 16.3$. In fact, using calculations discussed in a later section, we can show that B has the smallest expectation among *all* allocations of 10 tokens. Armed with this information, it appears that B is the superior position. Or is it?

When we play the two positions against each other, we find that B loses to A more than twice as often as B beats A! Why does this happen? Essentially because if a 1 is rolled any time before B is finished, then A's configuration becomes a *proper subconfiguration* of B's, so B cannot possibly win (it must lose or draw). However, *five* 4s must be rolled before B can achieve that status against A.

In fact, allocation A beats out *all* other allocations with 10 tokens, in terms of (on average) winning more often than it loses. Using terminology from Maurer [1980], we call such an allocation an *emperor*.

Does this same phenomenon occur when we increase the number of tokens? Alas, no. When we play the same 4-valued game with 20 tokens, allocation $(1, 3, 6, 10)$ has the lowest expected value *and* it beats the triangular allocation $(2, 4, 6, 8)$ in head-to-head competition. (We note that 1, 3, 6, and 10 are *triangular numbers*, but that's just a coincidence!) But here's the strange part: Allocation $(2, 4, 6, 8)$ beats $(1, 4, 6, 9)$ which in turn beats $(1, 3, 6, 10)$! In other words, we have a situation with *nontransitive probabilities*, as illustrated in **Figure 4**. The arc from $(1, 3, 6, 10)$ to $(2, 4, 6, 8)$ indicates that $(1, 3, 6, 10)$ beats $(2, 4, 6, 8)$ with probability .433 and loses to $(2, 4, 6, 8)$ with probability .388 (and therefore draws with probability .179). Although this game has no emperor, the three allocations above are part of an *emperor cycle* in that each of them defeats all other allocations of 10 tokens.
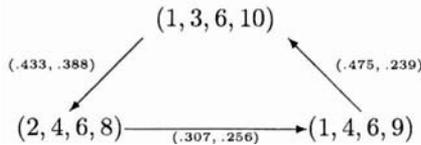
Figure 4. Subset of the emperor cycle in the 20-token game with $P = (.1, .2, .3, .4)$.

An emperor cycle also exists in the 5-token game with probability vector $P = (1/6, 2/6, 3/6)$. Here, allocation $(0, 2, 3)$ defeats $(0, 1, 4)$, who defeats $(1, 2, 2)$, who defeats $(1, 1, 3)$, who defeats $(0, 2, 3)$, as illustrated in **Figure 5**.

When the same game is played with 3 tokens, then allocation $(0, 1, 2)$ is an emperor. However, several non-emperor cycles exist in this game, as illustrated in **Figure 6**. One can easily imagine lucrative scams based on these nontransitive properties, easily played with a single six-sided die.
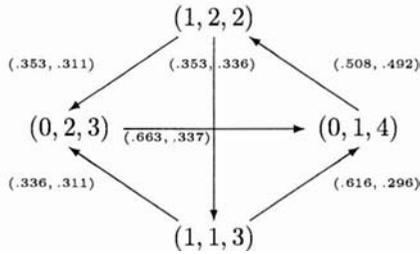
$(1, 2, 2)$

$(.353, .311)$ $(.353, .336)$ $(.508, .492)$

$(0, 2, 3)$ $\overline{(.663, .337)}$ $\longrightarrow$ $(0, 1, 4)$

$(.336, .311)$ $(.616, .296)$

$(1, 1, 3)$

Figure 5. Emperor cycle in the 5-token game with $P = (\frac{1}{6}, \frac{2}{6}, \frac{3}{6})$.

$(1, 0, 2)$

$(.574, .476)$ $(.578, .422)$ $(.328, .193)$

$(0, 2, 1)$ $\longleftarrow$ $\overline{(.507, .417)}$ $(1, 1, 1)$

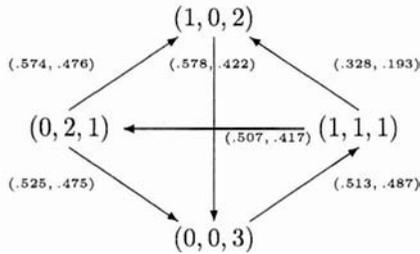$(.525, .475)$ $(.513, .487)$

$(0, 0, 3)$

Figure 6. Cycles in the 3-token game with $P = (\frac{1}{6}, \frac{2}{6}, \frac{3}{6})$.

Theoretical Results

Having described some of the surprising results that arise in particular variations of Knock 'm Down, we now explore some of the general results that apply to all instances of the game. Ideally, we would like to have a theorem that could construct directly the emperor or the emperor cycle from the probability distribution and the number of tokens. Unfortunately, the unintuitive nature of the game suggests that such a theorem would be very complicated. However, we have made significant progress towards characterizing the allocation of tokens that minimizes the expected number of rolls required to clear the allocation. While we have already demonstrated that this allocation is not always an emperor, it generally does very well in a tournament and may serve as a starting point for searching for the emperor or an emperor cycle. In the interest of space, the proofs of these results are omitted; the interested reader should see Fluet [1999].

In the following discussion, we are interested in the game played with t tokens on a board with values having probabilities $P = (p_1, p_2, \ldots, p_n)$. Then a *minimal* allocation $X^* = (x_1^*, x_2^*, \ldots, x_n^*)$ is one that minimizes the average number of rolls needed to clear the board, that is, $E[X^*] \leq E[X]$ for all t-token allocations X. Our first result is the following:

$$\text{If } p_a = p_b, \text{ then } |x_a^* - x_b^*| \leq 1. \tag{1}$$

This result largely confirms our intuition that values with equal probability

should have equal numbers of tokens. Alternatively, we can assert that any allocation that does not satisfy this result can be improved by "evening out" the distribution of tokens. Another result suggested by our intuition is that values with greater probability should have more tokens:

$$\text{If } p_a < p_b, \text{ then } x_a^* \leq x_b^*. \tag{2}$$

When $p_a < p_b$, it is possible that $x_a^* = x_b^*$. For example, in the 2-token game with $P = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$, the minimal allocation is $X^* = (0, 1, 1)$.

Our intuition suggests that the minimal allocation of tokens should resemble the histogram of probabilities. The 10-token example in **Figure 3** illustrates that an allocation may have the exact shape as the probability histogram, yet may not minimize the expected number of rolls. The next result, however, implies that the minimal allocation must at least "respect" proportions.

$$\text{If } p_a < p_b, \text{ then } \frac{x_a^* - 1}{x_b^*} < \frac{p_a}{p_b}. \tag{3}$$

We note that the 2-token example above illustrates that the stronger conclusion $x_a^*/x_b^* < p_a/p_b$ is not attainable.

Notice that condition **(3)** implies condition **(2)**, and that together with **(1)** it allows us to reduce drastically the search space for minimal allocations. For example, in the 6-token game with $P = (1/6, 1/3, 1/2)$, there are 28 different allocations. However, only 5 of these satisfy **(3)**. The savings become more extreme as the number of values and tokens are increased. In the 10-token game with $P = (.1, .2, .3, .4)$, there are $\binom{10+4-1}{10} = 286$ different allocations, but only 16 are potentially minimal.

The 2-valued game

For games with two values (i.e., $P = (p, 1 - p)$), we give an exact solution to the t-token game. Even here, we have some nice surprises.

From **(3)**, we know that if $p < 1/2$, then the minimal allocation $X^* = (x_1^*, x_2^*)$ must satisfy

$$\frac{x_1^* - 1}{x_2^*} < \frac{p}{1 - p},$$

which leads to $x_1^* < pt + 1 - p$. Perhaps for this game, we might expect to see a normal looking answer like $x_1^* \approx pt$, $x_2^* \approx (1-p)t$. However, the $P = (1/3, 2/3)$ game with $t = 9$ tokens yields $X^* = (2, 7)$ instead of $(3, 6)$. When $t = 1200$, we have $X^* = (393, 807)$, not $(400, 800)$. What's going on here?

For the t-token game with $P = (p_1, p_2)$, Fluet [1999] has shown that the expected time to clear allocation $X = (x_1, x_2)$ is

$$E[X] = t + \frac{p_2^t}{p_1} \sum_{k=0}^{x_1} (x_1 - k) \binom{t}{k} \left(\frac{p_1}{p_2}\right)^k + \frac{p_1^t}{p_2} \sum_{k=0}^{x_2} (x_2 - k) \binom{t}{k} \left(\frac{p_2}{p_1}\right)^k.$$

Also, it can be shown that as x_1 increases from 0 to t (and thus x_2 decreases from t to 0), $E[X]$ decreases and then increases, achieving its minimum at x_1^*, the smallest number x satisfying

$$\sum_{k=0}^{x} \binom{t}{k} p_1^k (1-p_1)^{t-k} > p_1.$$

Put another way, x_1^* is the p_1th percentile of the binomial distribution with parameters t and p_1. For example, if a weighted coin with heads probability $1/3$ is flipped 9 times, then H, the number of heads obtained, has a binomial distribution with

$$Pr(H = k) = \binom{9}{k} \left(\frac{1}{3}\right)^k \left(\frac{2}{3}\right)^{9-k}.$$

Thus, $Pr(H = 0) = 0.026$, $Pr(H = 1) = 0.143$, and $Pr(H = 2) = 0.234$; hence, $x_1^* = 2$, since $P(H \leq 1) < 1/3 < P(H \leq 2)$. With 1200 tokens, we have $P(H \leq 400) \approx 1/2$, whereas $P(H \leq 393)$ is just over $1/3$, whence $x_1^* = 393$.

In general, when t is large, x_1^* can be estimated quite accurately using a normal approximation. Specifically,

$$x_1^* \approx tp_1 + z_{p_1} \sqrt{tp_1(1 - p_1)},$$

where z_{p_1} is the p_1th percentile of the standard normal distribution. For example, with $P = (\frac{1}{3}, \frac{2}{3})$ and $t = 1200$, the normal approximation yields

$$x_1^* \approx 400 - 0.43 \sqrt{\frac{800}{3}} \approx 392.98,$$

and similarly, $x_2^* \approx 807.02$. As t gets larger, x_1^* gets farther away from tp_1. But it cannot get *too* far away, since

$$\lim_{t \to \infty} \frac{x_1^*}{t} = p_1,$$

agreeing with whatever is left of our intuition!

Although allocation X^* minimizes the expected number of rolls to bear off, it may still not be an emperor. For example, when $P = (1/3, 2/3)$, the minimal allocation with 2 tokens is $(0, 2)$, but it loses to $(1, 1)$ more than half the time. (Specifically, one can easily show that $E[0, 2] = 3$ and $E[1, 1] = 3.5$, but $(1, 1)$ beats $(0, 2)$ with probability $5/9$.) On the other hand, we have not discovered any examples of nontransitive behavior in a two-valued game.

Computational Details

We mention some details about the computations used to generate the results of this paper. We compute $E[X]$ by *conditioning* on the outcome of the first

roll and taking the appropriate weighted average. For instance, for the game with $P = (.1, .2, .3, .4)$, we have

$$E[1, 2, 3, 4] = 1 + .1E[0, 2, 3, 4] + .2E[1, 1, 3, 4] + .3E[1, 2, 2, 4] + .4E[1, 2, 3, 3]$$

and

$$
\begin{aligned}
E[0, 2, 3, 5] &= 1 + .1E[0, 2, 3, 5] + .2E[0, 1, 3, 5] + .3E[0, 2, 2, 5] + .4E[0, 2, 3, 4] \\
&= \frac{10}{9}(1 + .2E[0, 1, 3, 5] + .3E[0, 2, 2, 5] + .4E[0, 2, 3, 4]).
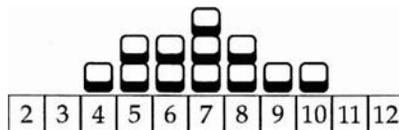\end{aligned}
$$

Notice that if X is an allocation of t tokens, then $E[X]$ depends only on allocations of $(t-1)$ tokens (except of course $E[0, 0, \ldots, 0] = 0$), and also allocation $X = (x_1, x_2, \ldots, x_n)$ has $(1 + x_1)(1 + x_2) \cdots (1 + x_n)$ subconfigurations. How many t-token allocations must we check to find one that minimizes $E[X]$? Theoretically, there are

$$\binom{t + n - 1}{t}$$

t-token allocations $(x_1, x_2, \ldots, x_n)$. In practice, we usually do not need to inspect this many, as we now illustrate.

Consider the original 12-token game with $P = (p_2, p_3, \ldots, p_{12})$ equal to $\frac{1}{36}(1, 2, 3, 4, 5, 6, 5, 4, 3, 2, 1)$. There are $\binom{22}{12} = 646,646$ possible allocations, but most need not be considered for minimality. By exploiting symmetry and conditions **(1)** and **(2)**, we need consider only allocations that are *monotonic* in the following sense. Since $p_{12} \leq p_2 \leq p_{11} \leq p_3 \leq p_{10} \leq p_4 \leq p_9 \leq p_5 \leq p_8 \leq p_6 \leq p_7$, there must exist a minimal allocation X^* satisfying $x_{12}^* \leq x_2^* \leq x_{11}^* \leq \cdots \leq x_7^*$. The number of such allocations is $\pi(12, 11) = 76$, where $\pi(t, n)$ denotes the number of partitions of the integer t into at most n parts.

But this number can be reduced to 49 once we factor in condition **(3)**. The solution to the 12-token game is the allocation $X^* = (0, 0, 1, 2, 2, 3, 2, 1, 1, 0, 0)$ shown in **Figure 7**. **Table 1** provides minimal allocations of up to 12 tokens in the original game.

| 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |

Figure 7. Minimal allocation of 12 tokens.

Based on **Table 1** and other tables like it, we make the following:

Conjecture 1 *If X^* is a minimal solution to the t-token game defined by a vector P, then there exists a minimal allocation to the $(t + 1)$-token game that properly contains X^*.*

Table 1.

Minimal allocations for the original game with t tokens, $1 \le t \le 12$.

t	\multicolumn{11}{c}{X^*}											$E[X^*]$
	2	3	4	5	6	7	8	9	10	11	12	
01	0	0	0	0	0	1	0	0	0	0	0	6.0
02	0	0	0	0	1	1	0	0	0	0	0	9.927
03	0	0	0	0	1	1	1	0	0	0	0	12.505
04	0	0	0	1	1	1	1	0	0	0	0	15.476
05	0	0	0	1	1	1	1	1	0	0	0	17.768
06	0	0	0	1	1	2	1	1	0	0	0	19.762
07	0	0	0	1	2	2	1	1	0	0	0	22.279
08	0	0	0	1	2	2	2	1	0	0	0	24.306
09	0	0	1	1	2	2	2	1	0	0	0	26.430
10	0	0	1	1	2	3	2	1	0	0	0	28.267
11	0	0	1	1	2	3	2	1	1	0	0	29.865
12	0	0	1	2	2	3	2	1	1	0	0	31.922

If **Conjecture 1** is true, then we need consider at most n allocations. For instance, consider finding the minimal 13-token allocation. By monotonicity, we would have to consider only adding a token to one of the values 3, 6, 7, and 9. In fact, (0,0,1,2,2,3,2,2,1,0,0) is the minimal 13-token allocation. If we continue this process of searching for minimal allocations by building on previous ones, we find for the 36-token game the allocation $X^* = (0, 1, 3, 4, 6, 8, 6, 4, 3, 1, 0)$, with $E[X^*] = 69.569$ (see **Figure 8**). Further calculations verify that X^* is indeed minimal.
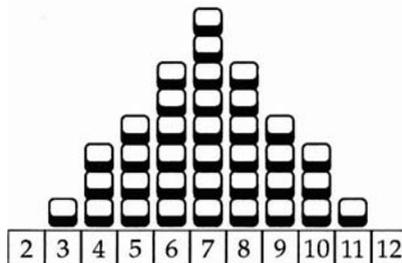
Figure 8. Minimal allocation of 36 tokens (but not an emperor).

So, how well do these minimal allocations do in competition with other allocations? To determine how allocation X performs against allocation Y, we define the WDL *(pronounced "widdle") function* as follows. We say $\mathrm{WDL}(X, Y) = (w, d, l)$, where $w = Pr(X$ wins against $Y)$, $d = Pr(X$ draws against $Y)$, and $l = Pr(X$ loses against $Y)$. Naturally, it must always be true that $w + d + l = 1$. As we did for $E[X]$, we can compute $\mathrm{WDL}(X, Y)$ recursively by conditioning on the outcome of the first roll. For example, in the game with $P = (.1, .2, .3, .4)$,

$\mathrm{WDL}\big((1,2,3,4),(0,2,3,5)\big)$ is equal to

$$\begin{aligned} & .1\,\mathrm{WDL}\big((0,2,3,4),(0,2,3,5)\big) + .2\,\mathrm{WDL}\big((1,1,3,4),(0,1,3,5)\big) \\ & + .3\,\mathrm{WDL}\big((1,2,2,4),(0,2,2,5)\big) + .4\,\mathrm{WDL}\big((1,2,3,3),(0,2,3,4)\big). \end{aligned}$$

Generally, $\mathrm{WDL}(X,Y)$ depends only on "smaller" allocations, until we hit a base case. Specifically, we have

$$\mathrm{WDL}(X,Y) = \begin{cases} (0,1,0), & \text{if } X = Y; \\ (1,0,0), & \text{if all nonzero entries of } X \text{ are strictly less than} \\ & \quad \text{all corresponding nonzero entries of } Y; \\ (0,0,1), & \text{if all nonzero entries of } Y \text{ are strictly less than} \\ & \quad \text{all corresponding nonzero entries of } X. \end{cases}$$

For a more detailed discussion of implementation issues, see Fluet [1999].

Returning to our original game, we find that the 12-token allocation in **Figure 7** that minimizes the average number of rolls also does well in competition. It is at least a *local* emperor in that it defeats all *neighboring* allocations, that is, all allocations reachable by moving a single token. We speculate that it is a global emperor as well. On the other hand, the minimal 36-token solution to the same game fails to be even a local emperor, losing to 4 of its monotonic neighbors. For this game, we speculate that $(0,2,3,4,6,7,5,4,3,2,0)$ is the emperor, since it defeats all of its neighbors as well as the triangular allocation $(1,2,3,4,5,6,5,4,3,2,1)$.

Variations and Open Questions

In the original game, the same dice rolls are used for both players. Suppose instead that each player has his/her own two dice, and the players take turns. The same techniques discussed in this paper can be used to analyze this game as well. Although much rarer, some nontransitive behavior is exhibited here too, as described in Fluet [1999].

Many interesting questions remain unanswered that we hope the reader will explore. Although we have a simple description of the minimal solution to the 2-valued games, a solution to the n-valued games remains elusive for $n \geq 3$. Failing that, we would like to find more properties of solutions. Are local emperors necessarily global emperors? No, as $(1,2,2)$ in **Figure 5** demonstrates; $(0,1,4)$ beats it. Are local minimal allocations necessarily global minimal allocations? We don't believe so, but we have yet to find a counterexample. How about **Conjecture 1**? If it's true, is there a simple rule to determine which value deserves the next token? Are there variations of condition **(3)** that provide interesting *lower* bounds on x_a/x_b, when $p_a < p_b$? All of these questions are still standing, but we hope that some of our readers will be able to knock 'em down!

Acknowledgment

References

Fendel, Dan, and Diane Resek, with Lynne Alper and Sherry Fraser. 1997. *Interactive Mathematics Program*. Year 1. Key Curriculum Press.

Fluet, Matthew T. 1999. Searching for optimal strategies in Knock 'm Down. Senior Thesis. Harvey Mudd College.

Hunt, Gordon. 1998. Knock 'm Down. *Teaching Statistics* 20: 59–62.

Maurer, Stephen B. 1980. The King Chicken theorems. *Mathematics Magazine* 53: 67–80.

About the Authors

Arthur Benjamin, famous as a "mathemagician" and lightning calculator, is the all-time point leader of the American backgammon Tour. He is an associate professor of mathematics at Harvey Mudd College, where he has taught since earning his Ph.D. in Mathematical Sciences from Johns Hopkins in 1989. He has co-authored with students several previous articles in *The UMAP Journal* about games.

Matthew Fluet is a senior mathematics major at Harvey Mudd College, pursuing the Mathematics with Computer Science option. In addition to mathematics, he is completing a humanities concentration in Medieval Studies. After graduation, he plans to begin graduate studies in Computer Science.