A STOCHASTIC SUBTRACTION GAME

ALEXANDER SIETSEMA

Abstract. The basic subtraction game is well-known example in combinatorial game theory. Many variants which introduce significantly more complexity have been introduced and studied in recent years. We examine a variant of the basic subtraction game which introduces randomness into the outcomes of the moves chosen by each player, and prove some general properties about it as well as a few theorems about more specific games.

1. Introduction

The basic subtraction game is simple: two players begin with a pile of k counters. Alternating turns, each player may choose to remove a number of counters from the pile from the subtraction set $S = \{1, 2, 3\}$. The goal of the game is to be the player to remove the last counter.

Winning strategies for this game, and all games in which each player chooses from a fixed subtraction set, are well-known [1]. For the above game, the moving player always has a winning strategy if the number of counters remaining is not a multiple of four, otherwise the moving player is guaranteed to lose. Similar results for other subtraction sets may also be solved for.

A number of variants of this game exist, including dynamic variants, in which the subtraction set changes based on the number of counters remaining [4]; blocking variants, in which the opposing player may block certain moves for the moving player [2, 3], and multi-player variants, in which there are more than two players [3, 5]. However, all of these versions of the game do not involve any randomness.

In our variant, we replace the subtraction set with a set of (different) dice. The moving player chooses a die from this set, and rolls it to determine how many counters to take from the pile. This is a substantial change, as the players no longer have direct control over their moves, instead only indirectly through their choice of dice. By introducing randomness, we also move away from the combinatorial end of game theory, which is only equipped to study games with perfect information.

This paper is organized as follows: we will begin by formally describing the stochastic subtraction game, introduce a few general properties, and finally prove two theorems about a particular simple example. Afterwards, we will discuss remaining open questions as well as avenues for future work.

Date: April 2022.

This work completed as part of MTH 492H (Honors Senior Thesis), advised by Dr. Albert Cohen.

2. Notation and general properties

We begin by formally restating the rules of the stochastic game. As before, we will have two players and a pile of k counters. However, instead of the subtraction set we will define the *dice set* $\mathbb{D} = \{D_1, D_2, \ldots, D_m\}$, where each die $D_i = (d_{i1}, d_{i2}, \ldots, d_{in})$: is identified by its faces. On each turn, the moving player selects a die D_i from the dice set $\mathbb{D}$, and a face is randomly selected from that die with uniform distribution (i.e., the die is rolled). The moving player then takes the number of counters on the selected face from the pile. As before, the goal of the game is to be the player to remove the last counter.

For the sake of simplicity, we will assume that all faces of a given die are distinct and take positive integer values, and will additionally assume that faces are picked uniformly. We do not, however, assume that all dice have the same number of faces.

Because randomness is involved, we cannot construct a perfect winning strategy, so we instead consider win probabilities. We define $p(k)$ to be the moving player's probability of winning with k counters remaining, assuming that the moving player will play to maximize this probability for every subsequent move, and the other player will play to minimize this probability (this is known as the *expectiminimax* algorithm). As well, we let $q(k)$ be the winning probability of a player when the *other* player is moving with k counters remaining.

Further, let $p(D, k)$ represent the moving player's winning probability when choosing die D. Thus, the moving player will maximize $p(D, k)$ over all available dice. We define $d(k)$ to be the best choice of die for the moving player with k counters remaining.

The proposition below, which we will refer to frequently, follows immediately from the rules definitions above and formalizes them further. We leave the verification as an exercise to the reader.

Proposition 1. *The following properties hold for all games:*

$$d(k) = \arg\max_{D \in \mathfrak{D}} p(D, k) \tag{1}$$

$$p(k) = \max_{D \in \mathfrak{D}} p(D, k) = p(d(k)) \tag{2}$$

$$q(k) = \min_{D \in \mathfrak{D}} q(D, k) \tag{3}$$

$$p(D_i, k) = \frac{1}{n_i} \sum_{d \in D} q(k - d) \tag{4}$$

$$q(D_i, k) = \frac{1}{n_i} \sum_{d \in D} p(k - d) \tag{5}$$

$$q(D_i, k) = 1 - p(D_i, k) \tag{6}$$

$$p(D_i, k) \leq p(k) \tag{7}$$

One point of particular interest is (6) above, which allows us to convert statements about the opposing player's using particular dice to statements about the moving player. In fact, we have symmetry in the winning probabilities for each player.

Proposition 2. *In every game, $q(k) = 1 - p(k)$ for all k.*

Proof.

$$\begin{aligned}
1 - p(k) &= 1 - \max_{D \in \mathfrak{D}} p(D, k) \\
&= \min_{D \in \mathfrak{D}} 1 - p(D, k) && \text{since } 0 \le p(\cdot) \le 1.0 \\
&= \min_{D \in \mathfrak{D}} q(D, k) = q(k) && \text{Prop. 1 (6)}
\end{aligned}$$

□

One useful fact will come from a particular class of dice: dice whose indices are an additively shifted version of another die in the set. This proposition will be critical for the proofs in the next section.

Proposition 3. *Let* $D_1 = (d_1, d_2, \ldots, d_n)$ *and* $D_2 = (d_1 + i, d_2 + i, \ldots, d_n + i)$ *for* $i \in \mathbb{Z}$ *be in an arbitrary dice set* $\mathbb{D}$. *Then in the* $(\mathbb{D}, k)$ *game,*

$$p(D_1, k) = p(D_2, k + i)$$

.

Proof.

$$\begin{aligned}
p(D_2, k+i) &= \frac{1}{n} \sum_{d \in D_2} q((k+i) - d) \\
&= \frac{1}{n} \left(q\left((k+i) - (d_1 + i)\right) + q\left((k+i) - (d_2 + i)\right) + \ldots + q\left((k+i) - (d_n + i)\right) \right) \\
&= \frac{1}{n} \left(q(k - d_1) + q(k - d_2) + \ldots + q(k - d_n) \right) \\
&= p(D_1, k)
\end{aligned}$$

□

We finish this section by outlining some of the relevant questions pertaining to this game. Aside from the obvious problem of finding an efficient method to compute $p(\cdot)$, we also observe two interesting phenomena. The first is that, it appears, that $p(k) \to \frac{1}{2}$ as $k \to \infty$ for all choices of dice set so long as at least one die has more than two faces (i.e., some amount of chance is involved). While intuitive, the proof of this is non-obvious.

The second interesting phenomenon is that it is possible for a particular die never to be optimal for all numbers of counters remaining. The next section will focus on proving this for two simple cases, but a list of necessary and sufficient conditions for this to occur in general would be ideal.

3. Theorems for a particular game

The classical subtraction game uses a subtraction set $S = \{1, 2, 3\}$. To extend this to the stochastic case, a natural choice of die set is $\mathbb{D} = \{(1, 2), (2, 3), (1, 3)\}$, where we have three dice with two faces each. We will refer to each of these dice as D_1, D_2, and D_3, respectively.

From the definitions above we may inductively (albeit inefficiently) compute the winning probabilities for each die with every number of counters remaining, as seen in Table 1.

k	$p((1,2),k)$	$p((2,3),k)$	$p((1,3),k)$	$d(k)$	$p(k)$
1	1.0000	1.0000	1.0000	1	1.0000
2	0.5000	1.0000	0.5000	2	1.0000
3	0.0000	0.5000	0.5000	2	0.5000
4	0.2500	0.0000	0.2500	1	0.2500
5	0.6250	0.2500	0.3750	1	0.6250
6	0.5625	0.6250	0.4375	2	0.6250
7	0.3750	0.5625	0.5625	2	0.5625
8	0.4062	0.3750	0.4062	1	0.4062
9	0.5156	0.4062	0.4844	1	0.5156
10	0.5391	0.5156	0.4609	1	0.5391
11	0.4727	0.5391	0.5273	2	0.5391
12	0.4609	0.4727	0.4727	2	0.4727
13	0.4941	0.4609	0.4941	1	0.4941
14	0.5166	0.4941	0.4834	1	0.5166
15	0.4946	0.5166	0.5054	2	0.5166

TABLE 1. Optimal play for the $\{(1,2),(2,3),(1,3)\},k\}$ game for $k \leq 15$. The highlighted cells illustrate the equality of $p((1,2),2)$ and $p((2,3),3)$.

In the table, we can see the reflection of Proposition 3 by the diagonal equality of the first two columns. For example, we see that $p((1,2),2) = p((2,3),3)$. The following lemma states this fact for this particular game, and may be verified by direct application of Proposition 3.

Lemma 1. *In the* $(\{(1,2),(2,3),(1,3)\},k)$ *game, we have*

$$p(D_1,k) = p(D_2,k+1)$$

Again looking at the table, the more striking observation is that $d(k) \neq D_3$ for all $k \leq 15$, that is, $D_3 = (1,3)$ is never the best choice of die. It turns out that this is true for all k. To prove this, we will use induction combined with Lemma 1, as well as some computation that is mainly dealt with in the following lemma.

Lemma 2. *In the* $(\{(1,2),(2,3),(1,3)\},k)$ *game, we have*

$$p(D_3,k) \leq p(D_1,k) \iff p(k-3) \geq p(k-2)$$

$$p(D_3,k) \leq p(D_2,k) \iff p(k-1) \geq p(k-2)$$

Proof.

$$p(D_1,k) = \frac{1}{2}(q(k-1)+q(k-2))$$

$$p(D_2,k) = \frac{1}{2}(q(k-2)+q(k-3))$$

$$p(D_3,k) = \frac{1}{2}(q(k-1)+q(k-3))$$

Thus,

$$p(D_3,k) \leq p(D_1,k) \iff q(k-3) \leq q(k-2)$$

$$p(D_3,k) \leq p(D_1,k) \iff q(k-1) \leq q(k-2)$$

The result follows by applying Proposition 2. ☐

Theorem 1. *In the $(\mathbb{D}_1, k)$ game, D_3 is never the best choice, that is, $d(k) \neq D_3$ for all k.*

Proof. By Lemma 2, it suffices to show that $p(k-2) \leq p(k-3)$ or $p(k-2) \leq p(k-1)$ for all k. We proceed by induction on k, the initial number of counters remaining.

Base cases $\mathbf{k \leq 3}$: Verified by computation (see table).

Inductive step: Assume that $d(j) \neq D_3$ for all $j \leq k$. Then we know that $d(k-2) \neq D_3$, so $d(k-2)$ is either D_1 or D_2.

If $d(k-2) = D_1$, then we have

$$p(k-2) = p(D_1, k-2) = p(D_2, k-1) \leq p(k-1).$$

Likewise if $d(k-2) = D_2$, we have

$$p(k-2) = p(D_2, k-2) = p(D_1, k-3) \leq p(k-3).$$

In both cases one of the inequalities is satisfied, so we are done. □

Before discussing this result, we will prove another similar result where we replace D_3 with a die with faces $(1, 2, 3)$. For the sake of brevity we omit the probability table, but we observe that the new die is also never the best choice. The method of proof will be similar; the proof of the lemma is omitted (verify by computation), and the inductive proof is similar.

Lemma 3. *In the $(\{(1,2), (2,3), (1,2,3)\}, k)$ game, we have*

$$p(D_3, k) \leq p(D_1, k) \iff p(k-1) + p(k-2) \leq 2p(k-3)$$
$$p(D_3, k) \leq p(D_2, k) \iff p(k-2) + p(k-3) \leq 2p(k-1)$$

Theorem 2. *In the $(\{(1,2), (2,3), (1,2,3)\}, k)$ game, D_3 is never the best choice, that is, $d(k) \neq D_3$ for all k.*

Proof. Assuming by induction that $d(k-2) \neq D_3$, we again have two cases to show. Rather than verifying individual inequalities, in each case we will leverage both at once.

If $D(k-2) = D_1$, then we may find separately that

$$\begin{aligned} p(k-1) + p(k-2) &= p(k-1) + p(D_1, k-2) \\ &= p(k-1) + p(D_2, k-1) \\ &\leq p(k-1) + p(k-1) \\ &= 2p(k-1), \end{aligned}$$

so if $p(k-1) \leq p(k-3)$, then $p(k-1)+p(k-2) \leq 2p(k-3)$ as desired. Likewise, we have

$$\begin{aligned} p(k-2) + p(k-3) &= p(D_1, k-2) + p(k-3) \\ &= p(D_2, k-1) + p(k-3) \\ &\leq p(k-1) + p(k-3) \end{aligned}$$

so if $p(k-3) \leq p(k-1)$, by substitution we get $p(k-2) + p(k-3) \leq 2p(k-3)$. Since $p(k-1)$ must be either greater than or less than $p(k-3)$, we have proven the statement for $d(k-2) = D_1$. The same holds reasoning symmetrically for $d(k-2) = D_2$, so we are done. $\square$

These theorems are useful as small examples that inform general heuristics about which kinds of dice are preferable. In general, we see that dice with smaller variance in the faces tend to be optimal more frequently than other kinds of dice. The theorems show an example of each: compared to $(1,2)$ and $(2,3)$ the $(1,3)$ die has the same number of faces but larger variance, while the (1,2,3) die has more variance through more faces. Neither die is ever best compared to the other two. We also see another general pattern reflected in these theorems.

Conjecture 3. *If all dice in a set $\mathbb{D}$ have faces less than or equal to n, and $\mathbb{D}' = \{(1,2),(2,3),\ldots.(n-1,n)\} \in \mathbb{D}$, then only dice from $\mathbb{D}'$ will be chosen.*

We expect that verifying this is not difficult, but we have not yet completed a proof.

4. Conclusion and future work

As stated earlier, the three main questions for this game are identifying conditions for a die never to be optimal, proving the convergence of $p(k)$ to $\frac{1}{2}$, and finding a closed-form solution for $p(k)$ in a general game. We still have a ways to go before finding general results, and proving results for specific examples may be helpful.

There are many avenues of continued work on this game. One natural generalization is to define a die as a discrete (or ultimately continuous) random variable and consider games resulting from those dice sets. There are alterations to the game rules, for example, allowing consecutive turns, that are trivial in the deterministic game but nontrivial in the stochastic game. Finally, many of the other deterministic variants mentioned in the introduction directly apply to the stochastic game.

References

[1] Elwyn R Berlekamp, John H Conway, and Richard K Guy. *Winning ways for your mathematical plays*. AK Peters/CRC Press, 1982.

[2] Achim Flammenkamp, Arthur Holshouser, and Harold Reiter. Dynamic one-pile blocking nim. *the electronic journal of combinatorics*, pages N4–N4, 2003.

[3] Arthur Holshouser and Harold Reiter. Three pile nim with move blocking.

[4] Arthur Holshouser, Harold Reiter, and James Rudzinski. Dynamic one-pile nim. *Fibonacci Quarterly*, 41(3):253–262, 2003.

[5] Wen An Liu and Jia Wei Duan. Misère nim with multi-player. *Discrete Applied Mathematics*, 219:40–50, 2017.