Probability Notes I: Dice

Matt Ollis

February 17, 2017

1 Introduction

We want to understand games. Games often use dice. This document considers the mathematics of dice behaviour. What do we expect to happen? Can we judge how likely a particular outcome is? What strategies arise from dice-based play? How do we evaluate these strategies? We'll consider these questions in two particular contexts: Pig and Heroscape. Of course, the lessons learnt have far wider application. In particular, I hope the tools we develop will help when designing your own games (and, of course, allow you to indulge in the noblest of mathematical pursuits: scamming money out of those less numerate than you).

I assume that you have read "Games as Systems of Uncertainty" [4, p. 172] and will occasionally refer to it. However, I have tried to make theses notes more-or-less self-contained. I'll soft-pedal the formal proofs in favour of a practical understanding: while a more formal approach is valuable for many reasons the first step for game designers is to play with the concepts.

Before getting stuck into Pig and Heroscape, let's look at five fundamental rules of probability from which the rest of our discussion will flow.

2 Five basic rules

An *event* is a precisely defined potential outcome in a given situation. Examples include getting a 6 when rolling a single standard die, rolling a double on a pair of dice and getting at least three heads when tossing ten coins at once. We talk of the "probabilities" of events; the five rules capture what it is we mean by this. We denote the probability of an event X by $P(X)$.

Rule 1 *For any event X we have $0 \leq P(X) \leq 1$. If $P(X) = 0$ then the event is* impossible*; if $P(X) = 1$ then the event is* certain.

This rule simply says that probability operates on a scale from zero to one, with zero probability meaning that the event cannot happen and a probability of one meaning that the event is certain.

Example 1 *Roll a standard die and let X be the outcome. Then we have $P(X = 7) = 0$ and $P(X \leq 6) = 1$. Nothing profound here: it's impossible to roll a 7 on regular die; it's certain that you'll roll at most a 6. You knew this already, but it's good to see the notation in action in familiar situations.*

Our next rule gives meaning to probabilities that lie between zero and one.

Rule 2 *If there are n equally likely outcomes then the probability of any one outcome is $\frac{1}{n}$.*

Example 2 *Roll a standard die and let X be the outcome. Also, abbreviate $P(X = 1)$ by $P(1)$ and so on. We have*

$$P(1) = P(2) = P(3) = P(4) = P(5) = P(6) = \frac{1}{6}.$$

Similarly,

$$P(\text{X is even}) = P(\text{X is odd}) = \frac{1}{2}.$$

The "equally likely" condition is important here. We saw in *Rules of Play* [4, p. 178] that when rolling a pair of dice the probability of total of 5 is different to the probability of rolling a total of 7. Rules 4 and 5 will lead to an understanding of this situation.

Rule 3 says "something happens":

Rule 3 *The sum of the probabilities all of the different possible basic outcomes is 1.*

Example 3 *Roll that standard die again. We have:*

$$P(1) + P(2) + P(3) + P(4) + P(5) + P(6) = 1.$$

Rule 3 suggests that we can add up probabilities. For example, let E be the event that we roll at least one 6 when we roll two dice in turn. We might hope that:

$$P(E) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}.$$

Taking this hope further, let F be the event that we roll at least one 6 when we roll seven dice in turn.

$$P(F) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{7}{6} > 1.$$

Oops. What went wrong? The problem is that we gave undue weight to the possibility that more than one die would come up with a 6. The probability $P(E)$ that we get a six at least once on two dice needs to be smaller: consulting Tables 1 and 2 in *Rules of Play* [4, p. 178] we see that $P(E) = \frac{11}{36}$. Only 11 of the "basic outcomes" include a 6, not the 12 we'd expect from our naïve hope.

Salen and Zimmerman [4] leave "basic outcome" undefined in general—presumably it means an event that is somwhow indivisible into smaller events. This calculation shows us that we need to be wary. However, all is not lost. From this unsuccessful addition-of-probabilities experiment we can salvage Rule 4 that tells when we can add probabilities. Call two events *mutually* *exclusive* if they cannot happen simultaneously. (Whatever the strict definition of "basic outcome", it will be the case that basic outcomes are mutually exclusive but mutually exclusive events might not be basic outcomes.) So, rolling a 3 and rolling a 4 are mutually exclusive events when rolling a standard die, whereas when rolling two dice in turn rolling a 6 on the first and rolling a 6 on the second are not mutually exclusive as both can happen.

Rule 4 *If X and Y are mutually exclusive events then $P(X \text{ or } Y) = P(X) + P(Y)$.*

Example 4 *Let's roll two dice again with Rule 4 to hand. Again, let E be the event that we roll at least one six. We can divide E into three mutually exclusive possibilities:*

- *A: we roll a 6 on the first die but not the second,*
- *B: we roll a 6 on the second die but not the first,*
- *C: we roll a 6 on both dice.*

Table 1 in Rules of Play *[4, p. 178] tells us the probability of each of these events. $P(A) = \frac{5}{36}$, $P(B) = \frac{5}{36}$ and $P(C) = \frac{1}{36}$ and Rule 4 then gives*

$$P(E) = P(A) + P(B) + P(C) = \frac{5}{36} + \frac{5}{36} + \frac{1}{36} = \frac{11}{36}$$

which we know to be the correct answer.

Fired up with enthusiasm, one may now consider $P(F)$ as defined above: what is the probability that we roll at least one 6 when we roll seven dice in turn. More general questions are now also answerable: what is the probability that we roll at least four 6s when we roll ten dice? At least fourteen 6s when we roll twenty-three dice? And so on. While these are theoretically within reach with our current toolkit, some more results will make the job much easier.

Before Rule 5, here is an immediate consequence of the rules we've seen so far, where "not X" is the event that X does not happen:

Corollary 1 *For any event X, we have $P(X) + P(\text{not } X) = 1$.*

This corollary will frequently be useful.

Example 5 *Corollary 1 gives us a more efficient way to solve the problem of Example 4. Let D be "not E": the event that we do not roll at least one 6. Put differently, D is the event that we roll no sixes. Now $P(D)$ can be read from the tables as $\frac{25}{36}$ and we have*

$$P(E) = 1 - P(D) = 1 - \frac{25}{36} = \frac{11}{36}$$

as required.

Maybe there was not such a saving in effort with this example, but in conjunction with Rule 5 this technique will bring many more problems within range.

One more definition before Rule 5. Two events X and Y are *independent* if whether or not X occurs has no bearing on $P(Y)$. Whether a 6 is rolled on each of two dice are independent events. The probability of a 6 on the second die is $\frac{1}{6}$ regardless of the outcome on the first die.

Rule 5 *If X and Y are independent then $P(X \text{ and } Y) = P(X) \times P(Y)$.*

For the final time, let's look at the probability of at least one 6 during successive die rolls.

Example 6 *Rule 5 allows us to perform these calculations without recourse to the tabulation of all of the outcomes when two dice are rolled. This holds out (justified) hope that we can move towards more complex situations such as the seven dice example. Consider D as defined above. Knowing that $P(\text{not } 6) = 1 - \frac{1}{6} = \frac{5}{6}$ when rolling a single die, we calculate that*

$$P(D) = \frac{5}{6} \times \frac{5}{6} = \frac{25}{36}$$

and may proceed as in the previous example.

Example 7 *Now let's return to $P(F)$. What is the probability that at least one 6 is rolled when seven dice are rolled in turn? While it is possible to enumerate the possibilities and calculate their various probabilities, we more efficiently use the technique of the last example. Let G be the probability that we roll no 6s on the seven dice. Rule 5 gives*

$$P(G) = \frac{5}{6} \times \frac{5}{6} \times \frac{5}{6} \times \frac{5}{6} \times \frac{5}{6} \times \frac{5}{6} \times \frac{5}{6} = 0.279$$

and now $P(F) = 1 - P(G) = 0.721$ by Corollary 1. In other words, there is about a 72% chance of rolling at least one 6 when rolling seven dice.

That completes our grounding in probability. Not all questions are yet answerable (or, at least, easily answerable): the probability that we roll at least four 6s when we roll ten dice, for example, needs a little more theory to be tractable. That theory, if we choose to pursue this aspect of the tutorial further, will appear in the "Counting" section of Part II of these notes.

Before continuing to the particular setting of Pig, here are a few exercises for you to practise the techniques of this section.

Exercises

1. What is the probability of rolling a number less than 5 on a single die?

2. Roll two dice. What is the probability of each of the following events?

a. Rolling a total of exactly 5?
b. Rolling a total of at most 5?
c. Rolling a double?
d. Rolling exactly one 4?
e. Rolling at least one 4?
f. Rolling at least one 4 or at least one 5?

3. What is the probability of rolling at least one 6 when you roll four dice? When rolling ten dice? How many dice do you need to roll to have at least a 90% chance of rolling a 6? A 99% chance?

3 Pig

Before we consider Pig we consider a much simpler game. Honest Bob tosses a coin. If it's a head you win \$10, if it's a tail you lose. Honest Bob is charging \$8 to play, do you take him up on the offer? How about if he was charging \$2? \$5? If you answered "no", "yes", "maybe" (in that order) then you already have an intuitive grasp of "expected values".

Let $X_1, X_2, \ldots, X_n$ be the n possible (mutually exclusive) outcomes in some situation. Let $v_1, v_2, \ldots, v_n$ be the values that each of the outcomes is worth to us respectively. We then define the *expected value* of the game, $E(G)$, by

$$E(G) = v_1 P(X_1) + v_2 P(X_2) + \cdots + v_n P(X_n).$$

In words, we've multiplied the probability of each event by its value and added it all up.

Let's go back to Honest Bob to get a sense of what this expected value thingy does. Let X_1 be heads and X_2 be tails. Then, in the first example, $v_1 = 2$ because if heads is the outcome we win \$2 (\$10 winnings minus the \$8 to play). Similarly, $v_2 = -8$. Now,

$$E(G) = (2 \times 0.5) + (-8 \times 0.5) = 1 - 4 = -3$$

(the 0.5s are the probabilities of tossing a head or a tail). An expected value of $-\$3$ means that, over the long term, if you keep playing the game, you'll lose about \$3 per game on average. Of course, there is no one game in which you lose exactly \$3: sometimes you'll come \$2 to the good; other times you'll lose \$8. However, if you play a bajillion times you can expect to be about three bajillion dollars down when you're done.

Running the same analysis on the other two versions of the game, we find that the expected value is \$3 for the \$2-to-play game and \$0 for the \$5-to-play game. So, at \$2-to-play we make a profit in the long-term. At \$5-to-play we break even. Expected values give us a sense of the best choice to make if we are in the situation of making the same choice again and again. Often, the same choice is the best in the instance when you just make that choice once. Given one go at Honest Bob's game, the "no", "yes", "maybe" approach to the three options dictated by the expected value calculations is probably the way to go for individual games too.

The tension between expected values that tell you the best choice with respect to return in the long-run and more immediate concerns is—in a newly-formed theory of mine from the reading we've been doing—a crucial aspect of generating meaningful play in dice games.

So, our rule-of-thumb will be that a positive expected value is good and a negative expected value is bad. In situations with more options, the higher the expected value the better. With that in mind, let's turn to Pig. Quoting Salen and Zimmerman [4, p. 182] quoting Knizia [2], here are the rules of Pig:

Object: The aim of the game is to avoid rolling 1s and to be the first player who reaches 100 points or more.

Play: One player begins, then play progresses clockwise. On your turn, throw the die:

· if you roll a 1, you lose a turn and do not score.

· if you roll any other number, you receive the corresponding points.

As long as you receive points you can throw again, and again. Announce your accumulated points so that everyone can easily follow your turn. You may throw as often as you wish. Your turn ends in one of two ways:

· If you decide to finish your turn before you roll a 1, score your accumulated points on the notepad. These points are now safe for the rest of the game.

· If you roll a 1, you lose your turn and your accumulated points.

Record all scores on the notepad and keep running totals for each player. The first player to reach 100 points or more is the winner.

Will expected value considerations let us recreate the Pig strategy that Salen and Zimmerman [4, p. 183] report that Knizia calculates as best for Pig? That is, stop rolling once you've amassed 20 or more points in a turn.

At any point within a turn there are two options: roll or do not roll ("stick"). We'll denote them R and S respectively. Suppose you have k points at this point. We'll write $E(R|k)$ for the expected value of strategy R given that we have k points and $E(S|k)$ for the non-rolling analogue.

Example 8 *Suppose we have* 12 *points so far this turn and we must choose whether to roll or stick. What are the expected values for each strategy? Sticking is easy to calculate. If we stick then we bank our points:* $E(S|12) = 12$. *Now suppose that we roll. The probability of a* 1 *is* $\frac{1}{6}$ *and in this case we lose our* 12 *points, a value of* 0. *The probability of a* 2 *is* $\frac{1}{6}$ *and this gives us 14 points; the probability of a* 3 *is* $\frac{1}{6}$ *and this moves us to 15; and so on.*

$$E(R|12) = (0\times\frac{1}{6})+(14\times\frac{1}{6})+(15\times\frac{1}{6})+(16\times\frac{1}{6})+(17\times\frac{1}{6})+(18\times\frac{1}{6}) = 13.333.$$

As $13.333 > 12$ *our expected value is maximised if we roll again.*

Example 9 *Suppose now that we have* 30 *points on the turn.* $E(S|30) = 30$ *and running the numbers as before we find that* $E(R|30) = 27.5$. *In this case we are better off sticking.*

We could continue in the vein of these two examples and zero in on the precise number at which the cut-off between rolling and sticking occurs. However, we can be a little more efficient. Let c be that cut-off number. When we have c points during a turn the expected value will be the same whether we roll or stick, that is $E(R|c) = E(S|c)$. Therefore

$$(0\times\frac{1}{6})+((c+2)\times\frac{1}{6})+((c+3)\times\frac{1}{6})+((c+4)\times\frac{1}{6})+((c+5)\times\frac{1}{6})+((c+6)\times\frac{1}{6}) = c.$$

Simplifying the left hand side, this reduces to

$$\frac{5c+20}{6} = c$$

and hence $c = 20$. Exactly what Knizia prescribes!

Here is where we meet that tension between expected values and best move again. Early on when there are many moves to go and the goal is to get points on the board as quickly as possible, following the stick-on-20-and-higher rule will be the best strategy. However, if your opponent has 99 maybe it's worth taking more risks; if you're comfortably ahead, maybe taking fewer risks is a better strategy.

Exercises

1. Consider the alternative pig dice suggested my last batch of notes:
a. A six-sided die with sides $1, 1, 5, 5, 6, 6$,
b. A six-sided die with sides $1, 1, 1, 1, 10, 25$,
c. A twelve-sided die with sides $1, 2, 2, 2, 2, 2, 3, 3, 3, 4, 4, 5$.
For each of these dice, analyse the expected values as we did for a regular die. Given the choice against a player using a standard die, which one would you choose?

2. Wild Boar is played as follows: Each player selects a die from the choice of varied dice (one die is a standard six-sider; some have different numbers of sides; they tend to have different values on each side; at least one has a negative number on at least one side). Once each player has a die, they follow the rules of Pig with each player using his/her own die. Design some fair dice to go in the bag. Why are they fair? Is it possible to have a fair die in which one side has the number 100? Which range of dice as the choices do you think will lead to the most meaningful play? [Dice are regularly made in the shape of the Platonic solids, making 4, 6, 8, 12 and 20 siders all common. However, by using "spinners" (or computers) rather than dice, it is easy to have something that acts like a 51-sided die (or any other number). Feel free to experiment with strange numbers of sides in this exercise.]

3. Building on the previous two exercises, design a Pig-like game in which players use different dice at different points of the game.

4 Heroscape

Heroscape's dice-based combat mechanic lends itself to the types of analysis that we've been developing. It is a much richer game than Pig and so has the ability to sidestep the near-degenerate strategy that we saw exists in Pig. In this document we simply introduce the mechanic and perform a couple of basic calculations. This sets the stage for a more thorough investigation as interest dictates.

Heroscape has two basic types of six-sided dice, the attack die and defence die. The attack die has three skulls and three blanks. The defence die has two shields and four blanks. Immediately we see that the probability of rolling a skull with the attack die is $\frac{1}{2}$ and the probability of rolling a shield with the defence die is $\frac{1}{3}$. Some heroscape dice combine the two functions into a single die with three skulls, two shields and a single blank. As it is always clear whether it is being used as an attack or defence die, this does not cause confusion.

The combat mechanic is this: the attacking figure has a set number of attack dice and the defending figure has a set number of defence dice. Both sides roll their dice. If the attacking figure has more skulls than the defending figure has shields then wounds are inflicted. The number of wounds is equal to the number of excess skulls. There are two types of figure, squad figures and hero figures. Typically squad figures have a single life and heroes have multiple lives. Different characters can attack from different ranges; that need not concern us here. There is one extra rule that is of interest. If an attacking figure is higher than the defending figure, then the attacker gets an extra attack die. Likewise, if the defender is higher than the attacker, the defender gets an extra defence die. On level ground neither side gets an extra die. Unless otherwise stated, we will assume level ground.

Example 10 *Eldgrim has an attack of 2 and chooses to attack an Arrow Grut squad figure (single life) that has defence 1. What is the probability $P(kill)$ that Eldgrim kills the Arrow Grut? There are three options for Eldgrim's dice: 0, 1 or 2 skulls with probabilities of $\frac{1}{4}$, $\frac{1}{2}$ and $\frac{1}{4}$ respectively (why?). There are two options for the Arrow Grut: 0 or 1 shield, with probabilities of $\frac{2}{3}$ and $\frac{1}{3}$ respectively (again, why?). Letting $P(X, Y)$ denote the probability of Eldgrim getting exactly X skulls and the Arrow Grut getting exactly Y shields we have*

$$\begin{aligned} P(kill) &= P(1,0) + P(2,0) + P(2,1) \\ &= (\frac{1}{2} \times \frac{2}{3}) + (\frac{1}{4} \times \frac{2}{3}) + (\frac{1}{4} \times \frac{1}{3}) \\ &= \frac{7}{12} \end{aligned}$$

That's a slightly better than 50:50 chance.

Table 1 gives the results of similar calculations for numbers of attack and defence dice up to 10. Our value of $\frac{7}{12}$ is 0.583 as a decimal—this matches the table's value.

Table 1: Wound Chance [5]

The entry in column i and row j gives the probability of inflcting at least one wound when attacking with i dice against a defence of j dice.

	1	2	3	4	5	6	7	8	9	10
0	0.500	0.750	0.875	0.938	0.969	0.984	0.992	0.996	0.998	0.999
1	0.333	0.583	0.750	0.854	0.917	0.953	0.974	0.986	0.992	0.996
2	0.222	0.444	0.625	0.757	0.847	0.906	0.944	0.967	0.980	0.989
3	0.148	0.333	0.509	0.655	0.766	0.846	0.901	0.938	0.961	0.976
4	0.099	0.247	0.407	0.556	0.679	0.776	0.848	0.899	0.934	0.958
5	0.066	0.181	0.321	0.463	0.592	0.700	0.786	0.851	0.898	0.932
6	0.044	0.132	0.250	0.380	0.507	0.622	0.718	0.795	0.855	0.899
7	0.029	0.095	0.192	0.308	0.429	0.544	0.647	0.734	0.804	0.859
8	0.020	0.068	0.146	0.246	0.358	0.470	0.576	0.670	0.749	0.813
9	0.013	0.049	0.111	0.195	0.295	0.401	0.506	0.603	0.689	0.762
10	0.009	0.035	0.083	0.153	0.241	0.339	0.440	0.538	0.628	0.707

As you might expect, we can also count the expected number of wounds from an attack.

Example 11 *Again, Eldgrim is the attacker. This time he is attacking Marcu, a hero with 1 defence and 6 lives. The expected number of wounds,* $E(wounds)$ *is calculated as follows (using the notation and numbers of the last example):*

$$\begin{aligned} E(\text{wounds}) &= (1 \times P(1 \text{ wound})) + (2 \times P(2 \text{ wounds})) \\ &= (P(1,0) + P(2,1)) + (2 \times P(2,0)) \\ &= \frac{3}{4} \end{aligned}$$

This means that Eldgrim does about 0.75 wounds of damage per turn. As in the previous discussion of expected values, we are not claiming that Eldgrim will do exactly 0.75 wounds on any turn—that is impossible. But if Eldgrim keeps attacking he will score 0, 1, or 2 wounds each time with an average of 0.75 per turn.

Table 2 gives the expected damage to a hero with many lives for given attack and defence values. Our 0.75 number is where it should be.

Table 2: Average Damage [5]

The entry in column i and row j gives the expected damage when rolling i attack dice against a multi-life hero with j defence dice.

	1	2	3	4	5	6	7	8	9	10
0	0.500	1.000	1.500	2.000	2.500	3.000	3.500	4.000	4.500	5.000
1	0.333	0.750	1.208	1.688	2.177	2.672	3.169	3.668	4.167	4.667
2	0.222	0.556	0.958	1.403	1.872	2.354	2.845	3.339	3.837	4.335
3	0.148	0.407	0.750	1.150	1.589	2.052	2.530	3.017	3.510	4.005
4	0.099	0.296	0.580	0.932	1.334	1.770	2.230	2.705	3.189	3.680
5	0.066	0.214	0.444	0.747	1.107	1.511	1.947	2.405	2.878	3.361
6	0.044	0.154	0.337	0.593	0.910	1.278	1.685	2.121	2.579	3.050
7	0.029	0.110	0.254	0.466	0.741	1.071	1.446	1.856	2.294	2.750
8	0.020	0.078	0.190	0.363	0.598	0.889	1.230	1.612	2.026	2.464
9	0.013	0.055	0.141	0.281	0.479	0.733	1.038	1.389	1.776	2.193
10	0.009	0.039	0.105	0.216	0.380	0.599	0.869	1.187	1.546	1.939

Now suppose that Eldgrim has a choice of attacking the Arrow Grut or Marcu. Can expected values shed any light on the options? We need to know a bit more information first. Each unit in Heroscape has a cost: a squad of 4 Arrow Gruts is 40pts and Marcu is 20pts. (Each player brings an army of set total point value, usually in the range 400-600pts.) Taking the value of wounding Marcu or killing the Arrow Grut in proportion to their points we find that killing an Arrow Grut is worth 10pts and each wound we put on Marcu is worth 3.333pts. We can now work out the expected point value of attacking the Arrow Grut to be 5.83pts ($= 10 \times 0.583 + 0 \times 0.417$) and of attacking Marcu to be 2.5pts ($= 0.75 \times 3.333$). Eldgrim will get more value by attacking the Arrow Grut.

The tables allow us to perform more complicated calculations without going back to first principles each time:

Example 12 *The three Krav Maga Agents have a Minion of Utgar (defence 6 squad figure) in range. What is the probability that they can kill it in one turn (each agent attacks with three dice in a turn)?*

$$\begin{aligned} P(\text{kill}) &= 1 - P(\text{fail to kill}) \\ &= 1 - P(\text{one agent fails to kill})^3 \\ &= 1 - 0.750^3 \qquad \text{(from Table 1)} \\ &= 0.679 \end{aligned}$$

Example 13 *Major Q10 has potentially four different choices with which to attack. A normal attack of 4 dice "Norm"; a normal attack enhanced by height "NwH"; the Machine Pistol Special Attack comprising four attacks with 2 dice "Mach"; and the Wrist Rocket Special Attack comprising two attacks of 4 dice "Wrist". Q10 can use one of these each turn. They come with various range restrictions that might exclude some options (Wrist Rocket has range 4, Machine Pistol has range 7 and the normal attack has range 8) and it might not be possible to get height. Using the tables and calculations in the style of those we've been doing, we obtain the following expected-damage table against many-life heroes of varying defences:*

Def	E(Norm)	E(NwH)	E(Mach)	E(Wrist)
2	1.403	1.872	2.224	2.806
3	1.15	1.589	1.628	2.3
4	0.932	1.334	1.184	1.864
5	0.747	1.107	0.856	1.494
6	0.593	0.91	0.616	1.186
7	0.466	0.741	0.44	0.932

From this we conclude that we should use the Wrist Rocket special if it is available. Otherwise, go for the Machine Pistol unless the opposing hero has defence 4 or more and you have height, or has defence 7 if you don't have height. In these cases use the normal attack.

Similar analysis may be employed in each of the situations with up to four squad figures in range [3].

So do expected value calculations tell us how best to play Heroscape? Are these calculations required to play to a high standard? The answer is no in both cases.

Expected value calculations have a narrow focus. When choosing which of Q10's attacks to use they are very useful. They do not tell us whether it is worth the risk of getting closer to the enemy to use the more powerful attack. They assume that wounds/kills are of value in proportion to the points removed, an assumption that is an approximation at best as the state of play of the game has a large influence on the current value of figures. Many (most?) of the top players have never done a calculation like this. They develop an instinct for the best choice via playing a lot of games.

That said, I think that these calculations could provide a solid basis for default play. Maybe we're programming a computer to play. If it is computationally feasible to analyse all of the possible moves and choose the one with the highest expected value, I suspect that we would have created a reasonably strong Heroscape-playing program. However, I think it would struggle against most experienced players. Expected value calculations are like the stick-on-20-or-higher rule in Pig: a solid basis for reasonable play (assuming they can be calculated, which is not usually the case for humans mid-game) but a good player will override the results at crucial points.

I don't know whether the designers and playtesters use calculations or instinct to develop figures. However, in keeping the game balanced, it is clear they give a lot of consideration to the relative strengths of units.

Exercises

1. Check some numbers in the tables of this section. How far can you get before the calculations become too unwieldy? [Look for more tools to help with this in Part II.]

2. The Krav Maga Agents have a special ability called "Stealth Dodge". This lets them defend successfully against any number of skulls from a non-adjacent attacker as long as a single shield is rolled. The Agents' base defence

is 3 and they have one life each. What is the probability that a single attack of 3 from a non-adjacent figure will destroy an agent? Raelin gives two extra defence dice to all figures within 4 spaces. Repeat this question with the assumption that the agents are within Raelin's aura. [Raelin and the Krav Maga Agents form an "engine" in the terminology of Salen and Zimmerman [4].]

3. Instead of attacking, Braxas may use her Acid Breath: choose three squad figures within four spaces of Braxas and roll the 20-sided die for each. If you roll an 8 or higher, destroy that figure. Assuming that three squad figures are within range, what is the expected number of squad figures killed?

5 Where to now?

We've seen the fundamental types of probability calculation and how expected values can be a useful game tool. We haven't yet seen how to practically find the probability $P(14)$ that we have at least fourteen 6s when we roll twenty-three dice. The probability that we have fourteen 6s in a row followed by nine non-6s is $(\frac{1}{6})^{14} \times (\frac{5}{6})^{9}$ but what about all of the other situations in which we are successful? The theoretical section of Part II of these notes there will be called "Counting" and will enable us to answer this and similar questions. This theory naturally applies to some of the more difficult Heroscape calculations (taking off from wherever you got stuck in Exercise 1 of Section 4) and leads into the consideration of card games. My current intention for much of Part II is to consider Poker. Lobby me to use something different if you like.

Other potential future directions include considering conditional probabilities and approximating discrete distributions with continuous ones.

References

[1] www.heroscapers.com

[2] R. Knizia, *Dice Games Properly Explained*, Tadworth, Surrey: Elliot Right Way Books, 2001.

[3] "Ollie", Major Q10 Analysis, `www.heroscapers.com/community/showpost.php?p=621988postcount=62`, 2008.

[4] K. Salen and E. Zimmerman, *Rules of Play: Game Design Fundamentals*, Cambridge Massachusetts: MIT Press, 2004.

[5] "Sisyphus", Sisyphus' Probability Tables, `www.heroscapers.com/community/showthread.php?t=16263`, 2008.