Dice-Rolls in Role-Playing Games

Torben Mogensen
email: torbenm@di.ku.dk

15. marts 2016

Resumé

Most RPGs (role-playing games) use some sort of randomizer when resolving actions. Most often dice are used for this, but a few games use cards, rock-paper-scissors or other means of randomization.

There are dozens of different ways dice have been used in RPGs, and we are likely to see many more in the future. This is not an evolution from bad methods to better methods – there is no such thing as a perfect dice-roll system suitable for all games (though there are methods that are suitable for none). But how will a designer be able to decide which of the existing dice-roll method is best suited for his game, or when to invent his own?

There is no recipe for doing this – it is in many ways an art. But like any art, there is an element of craft involved. This paper will attempt to provide some tools and observations that, hopefully, will give the reader some tools for the craftsmanship involved in the art of choosing or designing dice-roll mechanisms for RPGs.

Parts of this document have been published as a column called "Roll the Bones" on `https://www.rpg.net`.

1 Introduction

Ever since Dungeons & Dragons was published in 1974, randomization has, with a few exceptions, been a part of role-playing games. Randomization has been used for creating characters, determining if actions are successful, determining the amount of damage dealt by a weapon, determining encounters ("wandering monsters", etc.) and so on. We will look mainly at randomizers for action resolution – the act of determining how successful an attempted action is. The reason for this is that this is in many ways the most critical part of an RPG and the part that is hardest to get right.

Dice of various types are the most common randomizers in role-playing games, and D&D was indeed known for introducing non-cubic dice into modern games, but a few games (such as Castle Falkenstein and the Saga system) use cards as randomizers, and some "diceless" games, like Amber, use no randomizers at all, apart from the inherent unpredictability of human behaviour. We will focus on dice in this article, but briefly touch on other randomizers.

We start by discussing some aspects of action resolution that it might be helpful to analyse when choosing a dice-roll mechanism, then a short introduction to probability theory followed by an analysis of some existing and new dice-roll mechanisms using the above.

2 Action resolution

When a character attempts to perform a certain action during a game, there are several factors that can affect the outcome. We will classify these as *ability*, *difficulty*, *circumstance* and *unpredictability*.

Ability is a measure of how good the character is at performing the type of action he or she attempts. This can be a matter of natural talent, training and tools. The quality of the ability will typically be given by one or more numbers, such as attribute, skill, level, weapon/tool bonus, feats or whatnot. This can be modified by temporary disabilities such as injury, fatigue or magic.

Difficulty is a measure of how hard the action is. This can be in the form of active opposition, inherent difficulty or a combination thereof. This is usually also given as one or more numbers.

Circumstance is a measure of external factors that may affect the outcome, making it harder, easier or less predictable. This can be terrain, time of day, lunar phase, weather and so on. Often, these factors are modeled as modifiers to ability or difficulty, but they can also be modeled separately.

Unpredictability is a measure of how random the outcome is. This often depends on the type of action performed – if a person tries to beat another in a game of Poker (where shuffling the deck introduces uncertainty), the outcome is more random than if the game was Chess (where there is no chance or hidden information), especially if there is a large difference in skill between the two opponents. But the outcome may be unpredictable even in situations with no explicit random element, such as when throwing a ball through a hoop. Chaotic systems such as this will be unpredictable even if the physical system is fully deterministic.

We will now look at some properties that action resolution systems might or might not have. We believe that a designer should think about these, even if only to conclude that they are irrelevant for the particular game design in consideration.

2.1 Detail and complexity

So, is the best action resolution mechanism the one that models these aspects most realistically or in most detail? Not necessarily. First of all, more realism will usually also mean higher complexity, which makes your game more difficult to learn and play, and more detail will typically mean more categories (of skills, tasks, etc.) and larger numbers (to more finely distinguish between degrees of ability, success, etc.), which will require larger character sheets and more calculation. Nor is utmost simplicity necessarily the best way to go – the result may be too inflexible and simplistic for proper use.

There is no single best compromise between simplicity and realism and detail, it depends on the type of game you want to make. For a game that is designed to emulate the silliness of Tex Avery cartoons, the fifty-fifty rule of Toon may be suitable: Regardless of ability, difficulty and circumstance, there is 50% chance that you will succeed in what you do. But for a game about WW2 paratroopers, you would want somewhat finer distinctions and take more factors into account. Nor does detail and realism have to be consistent in a single game – if the game wants to recreate the mood in The Three Musketeers, it had better have detailed rules for duels and seduction, but academic knowledge can be treated simplisticly, if at all. On the other hand, if the game is about finding lost treasure in ruins of ancient civilizations, detailed representation of historic and linguistic knowledge can be relevant, but seduction ability need not even be explicitly represented.

In short, you should not decide on an action resolution mechanism before you have decided what the game is about and which mood you want to impart.

2.2 Interaction of ability and difficulty

In many games, ability and difficulty (including aspects of circumstance and predictability) are combined into a single number that is then modified by a randomizer. In other games, ability and difficulty are separately modified by randomizers, and the results are then compared. You can even have cases where ability and difficulty are randomly modified in different ways. Similar issues are whether proactive or reactive actions (e.g., attack versus defense) are treated the same or differently, whether opposed and unopposed actions are distinguished, and how multiple simultaneous or chained actions are handled.

2.3 Degrees of success and failure

In the simplest case, all a resolution system needs to determine is "did I succeed?", i.e., yes or no. Other systems operate with degrees of success or failure. These can be numerical indications of the quality of the result, or they might be verbal characterisations such as "fumble", "failure", "success" and "critical success". Systems with numerical indications usually use the result of the dice roll more or less directly as degree of success/failure, while systems with verbal characterisations often use a separate mechanism to identify extreme results.

2.4 Nonhuman scales

Some games, in particular superhero or SF games, operate with characters or character-like entities at scales far removed from humans in terms of skill, size or power. These games need a resolution mechanism that can work at vastly different scales and, preferably, also handle interactions across limited differences in scale. Large differences in scale will usually make interactions impossible or trivially one-sided, so they need not be covered by the usual interaction mechanisms.

2.5 Luck versus skill

Let us say we set a master up against a novice. Should the novice have any chance at all, however remote, of beating the master? In other words, shall an unskilled character have a small chance of succeeding at an extremely difficult task and shall a master have a small chance of failing at a routine task?

The luck versus skill ratio may depend on the task in question – some tasks are inherently more random than others, such as Poker versus Chess.

On a related note, the amount of random variability may depend on skill. In the "real world", you would expect experienced practitioners to be more consistent in their performance than unskilled dabblers.

2.6 Hiding difficulty from the players

A GM might not always want to reveal the exact level of ability of an opponent to the players until they have seen him in action several times. Similarly, the difficulty of a task may not be evident to a player before it has been attempted a few times, and the GM may not even want to inform the players of whether they are successful or not at the task they attempt if this is not readily evident. For example, if a PC tries to determine if an NPC is lying, the GM may simply say "you don't think so", regardless of whether the NPC told the truth, or was lying through his teeth, but the PC just failed to realise this.

In all cases, the GM can decide to roll all dice and tell the players only as much as he wants them to know. But players often like to roll for their own characters, so you might want a system where the GM can keep, e.g., the difficulty level secret so the players are unsure if they succeed or fail or by how much they do so, even if they can see the numbers on their own dice rolls.

2.7 Diminishing returns

Many games makes it increasingly more difficult to improve the ability of a character. This is most often done by making the cost of increasing an ability increase with its level.

It can also be done through chance: The cost of improving an ability is constant, but there is an increasing chance that the ability will fail to improve. Such mechanisms are used both in Avalon Hill's RuneQuest and in Columbia Games' HârnMaster, where a dice-roll result has to exceed the current level of the ability in order for the improvement to happen.

A third alternative is to have linear cost of increasing ability, but reduce the effectiveness of higher skills through the way abilities are used in the randomization process, i.e, by letting the dice-roll mechanism give ever decreasing benefits for added ability.

3 Elementary probability theory

In order to fully analyse a dice-roll mechanism, we need to have a handle on the probability of the possible outcomes, at least to the extent that we can say which of two outcomes is most likely, and if a potential outcome is extremely unlikely. This section will introduce the basic rules of probability theory as these relate to dice-rolling, and describe how you can calculate probabilities for simple systems. The more complex systems can be difficult to analyse by hand, so we might have to rely on computers for calculations, so we will briefly talk about this too.

Calculating probabilities of dice is both easy and hard: You need only use a few very simple rules to figure out what the probabilities of the possible outcomes are, but for rolls involving many dice or possible rerolls, the calculations may be quite lengthy. In such cases, we can use a computer to do the calculation.

3.1 Events and probabilities

Probabilities usually relate to events: What is the chance that a particular event will happen in a particular situation?

The probability of an event e happening is modeled as a number between 0 and 1, with 0 meaning that the event can never happen and 1 meaning it is certain to happen. Numbers between these mean that it is possible, but not certain for the event to happen, and larger numbers mean greater likelihood of it happening. For example, a probability of $\frac{1}{2}$ means that the likelihood of an event happening is the same as the likelihood of it not happening. We use letters p and q to denote probabilities, and if we want the event e that a probability p concerns to be explicit, we write $p(e)$.

This brings us to the basic rules of probabilities:

The rule of negation: If an event has probability p of happening, it has probability $1-p$ of not happening.

The rule of coincidence: If two events are *independent* and have probabilities p and q, respectively, of happening, then the chance that both happen is $p \times q$ (p times q).

For example, if a die has probability $\frac{1}{6}$ of showing a 1, the probability that two rolls of this die will both show a 1 is $\frac{1}{6} \times \frac{1}{6} = \frac{1}{36}$. The probability that they will *not* both be 1 is $1 - \frac{1}{36} = \frac{35}{36}$.

Two events are *independent* if the outcome of one event does not influence the outcome of the other. For example, when you roll a die twice, the two outcomes are independent (the die doesn't remember the previous roll). On the other hand, the events "the die lands with an odd number facing up" and "the die show 4 or more" are not independent, as knowing one of these will help you predict the other more accurately. Taking any one of these events alone (on a normal d6), will give you a probability of $\frac{1}{2}$ of either one happening, but if you know that the result is odd, there is only $\frac{1}{3}$ chance of it being at least 4, as only one of 4, 5 and 6 is odd.

In the above, we have used an as yet unstated rule of dice: Dice are *fair*. A die is fair, all sides have the same probability of ending on top. In games, we usually deal with fair dice (as unfair dice are considered cheating), so we will, unless otherwise stated, assume this to be the case. So, if there are n sides to the die, each has a probability of $\frac{1}{n}$ of being on top after the roll. The number on the top face or vertex is usually taken as the result of the roll (though some d4s read their result at the bottom edges). Most dice have results from 1 to n, where n is the number of sides of the die, but most ten-sided dice go from 0 to 9 and some have numbers $00, 10, 20, \ldots, 90$. We will use the term dn about an n-sided die with numbers 1 to n with equal probability, and zn about an $n+1$-sided die with numbers 0 to n with equal probability. Hence, a typical ten-sided die will be a z9.

If we have an event e, we use $p(e)$ to denote the probability of this event. So, the rules of negation and coincidence can be restated as

$$\begin{aligned} p(\text{not } e) &= 1 - p(e) \\ p(e_1 \text{ and } e_2) &= p(e_1) \times p(e_2) \end{aligned}$$

3.2 Calculating with probabilities

We can use the rules of negation and coincidence to find probabilities of rolls that combine several dice, like the example above of two rolls of a die both showing 1. But what about the probability of rolling two dice such that at least one of them is a one? It turns out that we can use the rules of negation and coincidence for this too: The chance of having at least one die land on one is 1 minus the chance that neither land on ones. The chance of neither landing on 1 is the chance that the first is not a 1 times the chance that the other is not a 1. So we get that the probability of getting at least one one is $1 - \frac{5}{6} \times \frac{5}{6} = \frac{11}{36}$. We can calculate a general rule as

$$\begin{aligned} p(E_1 \text{ or } E_2) &= 1 - p(\text{not } E_1) \times p(\text{not } E_2) \\ &= 1 - (1 - p(E_1)) \times (1 - p(E_2)) \\ &= p(E_1) + p(E_2) - p(E_1) \times p(E_2) \end{aligned}$$

For another example, what is the chance of rolling a total of 6 on two d6? We can see that we can get 6 as $1+5$, $2+4$, $3+3$, $4+2$ and $5+1$, so a total of 5 of the possible 36 outcomes yield a sum of 6, so the probability is $\frac{5}{36}$. Note that we need to count $1+5$ and $5+1$ separately, as there are two ways of rolling a 1 and a 5 on two d6, but there is only one way of getting two 3s. To see this, think of the dice having two different colours, say, blue and red. Getting a 1 and a 5 can happen either if the blue die is 1 and the red die is 5 or if the blue die is 5 and the red die is 1, but to get two 3s, both the blue die and the red die have to show 3.

In general, when you combine several dice, you count the number of ways you can get a particular outcome and divide by the total number of rolls to find the probability of that outcome. When you have two d6, this isn't difficult to do, but if you have, say, five d10, it is unrealistic to by hand enumerate all outcomes and count those you want. In these cases, you either use a computer program to enumerate all possible rolls and count those you want, or you find a way of counting that doesn't require explicit enumeration of all possibilities, usually by exploiting the structure of the roll.

For simple cases, such as the chance of rollinga sum of S or more on m dn, some people have derived formulae that don't require enumeration. These formulas are, however, often cumbersome (and error-prone) to calculate by hand, albeit not as badly as explicitly counting all combinations, so you might as well use a computer anyway. For finding the chance of rolling a sum of S or more on 3d6, we can write the following program (in sort-of BASIC, though it will be similar in other languages):

```
count = 0
for i1 = 1 to n
 for i2 = 1 to n
   for i3 = 1 to n
    if i1+i2+i3 >= S then count = count + 1
   next i3
 next i2
next i1
print count/(6*6*6)
```

Each loop runs through all values of one die, so in the body of the innermost loop, you get all combinations of all dice. You then count those combinations that fulfill the criterion you are looking for. In the end, you divide this count by the total number of combinations (which in this case is $6 \times 6 \times 6$).

Such programs are not difficult to write, though it gets a bit tedious if the number of dice can change, as you need to modify the program every time (or use more complex programming techniques, such as recursive procedure calls or stacks). To simplify this task, I have developed a programming language called **Troll** specifically for calculating dice probabilities. In **Troll**, you can write the above as

```
sum 3d6
```

and you will get the probabilities of the result being equal to each possible value, as well as the probability of the result being greater than or equal to each possible value. Alternatively, you can write

```
count S <= (sum x d n)
```

which counts only the results that are at least S. You can find **Troll**, including instructions an examples, at `http://topps.diku.dk/~torbenm/troll.msp`.

3.3 Average, variance and spread

If you can assign a numeric value (such as numbers between 1 and 6) to each outcome, you can calculate an *average* (or *mean*) value of the possible rolls. This is the sum of the probability of each outcome multiplied by its value. More precisely, if the possible outcomes are $e_1, \ldots, e_n$ and the value of outcome e_i is $V(e_i)$, then the average of the outcomes is $p(e_1) \times V(e_1) + \cdots + p(e_n) \times V(e_n)$. For a single d6, the average is, hence, $1 \times \frac{1}{6} + 2 \times \frac{1}{6} + 3 \times \frac{1}{6} + 4 \times \frac{1}{6} + 5 \times \frac{1}{6} + 6 \times \frac{1}{6} = \frac{21}{6} = 3.5$. In general, a d$n$ has average $\frac{(n+1)}{2}$ and a zn has average $\frac{n}{2}$. So a ten-sided die with numbers from 0 to 9 has an average value of $\frac{9}{2} = 4.5$.

If you add several dice, you also add their averages, so, for example, the average sum of 3d6 is $3 \times \frac{7}{2} = 10.5$.

The *variance* of a number or outcomes with values is the sum of the squares of the distances of the values from the mean, i.e.,

$$p(e_1) \times (V(e_1) - M)^2 + \cdots + p(e_n) \times (V(e_n) - M)^2$$

where M is the mean value, as calculated above. This can be rewritten to

$$(p(e_1) \times V(e_1)^2 + \cdots + p(e_n) \times V(e_n)^2) - M^2$$

I.e., the average of the squares minus the square of the average. For a single dn, this adds up to $\frac{(n^2-1)}{12}$. For example, the variance of a d6 is $\frac{35}{12}$. A zn has the same variance as a d$(n+1)$, which is $\frac{((n-1)^2-1)}{12} = \frac{(n^2+2n)}{12}$.

When you add two dice, you also add their variances (like you do with averages), so the variance of the sum of five d6 is $5 \times \frac{35}{12} = \frac{175}{12}$, and so on.

It is, however, more common to talk about the *spread* (or *standard deviation*) of the outcomes. The spread is simply the square root of the variance. Examples: The spread of a d6 is $\sqrt{\frac{35}{12}} = 1.7078$ and the spread of 5d6 is $\sqrt{\frac{175}{12}} = 3.8188$.

The spread is a measure of how far away from the average value you can expect a random value to be. So if the spread is small, most values cluster closely around the average, but if the spread is large, you will often see values far away from the average. If two rolls have spreads s_1 and s_2, then their sum has spread $\sqrt{s_1^2 + s_2^2}$ (as the spread is the square root of the variance).

Note that the spread is *not* the average distance from the mean value. The latter is called the *mean deviation*, and (while intuitively more natural) isn't used as much as the standard deviation, mostly because it isn't as easy to work with. The mean deviation is defined as

$$p(e_1) \times |V(e_1) - M| + \cdots + p(e_n) \times |V(e_n) - M|$$

where $|x|$ is the absolute value of x. For a single dn, the mean deviation is $\frac{n}{4}$ if n is even and $\frac{n^2-1}{4n}$ if n is odd. A zn has the same mean deviation as a d$(n+1)$.

It gets more complicated when you add several dice, as (unlike for standard deviation), you can't compute the mean deviation of the combined roll from the mean deviations of the individual rolls. For example, d4 and d2+d4 both have mean deviation 1, but d4+d4 has mean deviation $\frac{5}{4}$ while (d2+d4)+(d2+d4) has mean deviation $\frac{11}{8}$.

For both even and odd n, 2dn has mean deviation $\frac{n^2-1}{3n}$, but it quickly gets a lot more complicated.

Troll can calculate the average, spread and mean deviation of a roll.

3.4 Open-ended rolls

The rules above can be used to calculate probabilities, mean and spread of any finite combination of dice (though some require complex enumeration of combinations). But what about open-ended rolls, i.e., rolls that allow unlimited rerolls of certain results? For example, an open-ended (or "exploding") d6 has the rolled value if this is between 1 and 5, but if you roll a 6, you roll again and add 6 to the result of that roll. If the second roll is also a 6, you roll and add again, so the final result is 6 times the number of sixes rolled plus a value between 1 and 5 for the final (non-six) roll. We can not finitely enumerate all the possible rolls (as there are infinitely many), so what do we do?

A simple solution it to limit the rerolls to some finite limit and, hence, get approximate answers (**Troll**, for example, does this). But it is, actually, fairly simple to calculate the average of a roll with unbounded rerolls.

Let us say that we have a roll that without rerolls has average M_0, that you get a reroll with probability p, and that when you reroll, the new roll is identical to the original (including the chance of further rerolls) and added on top of the original roll.

This gives us a recurrence relation for the average M of the open-ended roll: $M = M_0 + p * M$, which solves to $M = \frac{M_0}{(1-p)}$.

For an n-sided die with values $x_1, \dots, x_n$ and rerolls on x_n, this yields $M = \frac{x_1 + \cdots + x_n}{n-1}$ compared to the normal average $M_0 = \frac{x_1 + \cdots + x_n}{n}$. So the "exploding" d6 has average $\frac{1+\cdots+6}{6-1} = \frac{21}{5} = 4.2$

The variance is more complicated. If an n-sided die has values $x_1, \dots, x_n$ and rerolls on x_n, the variance V is

$$V = \frac{x_1^2 + \cdots + x_n^2}{n-1} - \frac{(x_1 + \cdots + x_{n-1})^2 - x_n^2}{(n-1)^2}$$

Compared to the variance V_0 of the same die without reroll:

$$V_0 = \frac{x_1^2 + \cdots + x_n^2}{n} - \frac{(x_1 + \cdots + x_n)^2}{n^2}$$

As an example, let us take the dice-pool system from White Wolf's "World of Darkness" game. In this system, you roll a number of d10s and count those that are 8 or more. Additionally, any 10 you roll adds another d10, which is also rerolled on a 10 and so on.

Without rerolls, the values are $x_1, \dots, x_n = 0, 0, 0, 0, 0, 0, 0, 1, 1, 1$. So the average of one open-ended die is $M = \frac{3}{9} = \frac{1}{3}$. If you roll N dice, the average is $\frac{N}{3}$. The variance of one WoD die is

$$V = \frac{3}{9} - \frac{2^2 - 1^2}{9^2} = \frac{8}{27}$$

As with normal dice, the variance of several open-ended dice add up, so the variance of N WoD dice is $\frac{8N}{27}$.

Another example is an open-ended "normal" dn with reroll on n. The average is

$$M = \frac{1 + \cdots + n}{n-1} = \frac{n(n+1)}{2(n-1)}$$

The variance is

$$V = \frac{1^2 + \cdots + n^2}{n-1} - \frac{(1 + \cdots + (n-1))^2 - n^2}{(n-1)^2} = \frac{n(n+1)(n^2+7n-2)}{12(n-1)^2}$$

In calculating the above, I have used the formulas

$$\begin{aligned} 1 + \cdots + n &= \frac{n(n+1)}{2} \\ 1^2 + \cdots + n^2 &= \frac{n(n+1)(2n+1)}{6} \end{aligned}$$

If a die is re-rolled when it has its maximal value, it is not very hard to find the probability of getting any particular result: Given a dN that is rerolled on a roll of N, a result that can be written as $k \times N + v$, where $0 < v < N$, has probability $\frac{1}{N^{k+1}}$, since you need to roll N k times and then v once, each of which has probability $\frac{1}{N}$. If the reroll is not on the maximal value, it gets more complicated, as there are multiple ways to roll the same value. Let us, for example, say that we reroll on a roll of 1. Then, the result 4 can be obtained by rolling a 4 in the first roll, rolling a 1 and then a 3, or by rolling a 1, a 1 and then a 2. It also gets more complicated if you roll multiple exploding dice, so in both these cases it is better to use a program (such as **Troll**) to figure out the probabilities.

3.5 Bell curves

When talking about distribution of results (such as dice rolls), people often use the term *bell curve* to mean that the distribution looks somewhat like this:

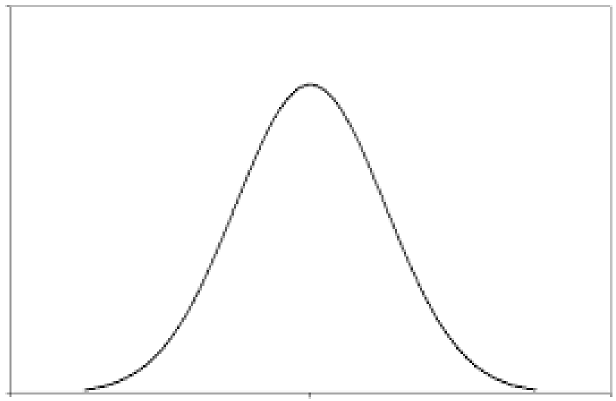

I.e., reminiscent of a normal distribution. Strictly speaking, dice rolls have discrete probability distributions, i.e., the distributions are not continuous curves but disconnected points, such as the classic 3d6 distribution, here shown as a bar graph:

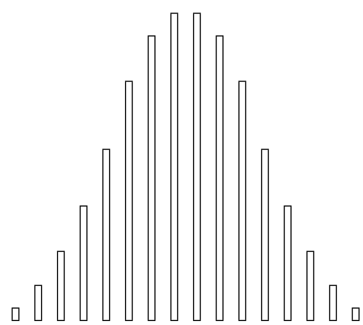

So you can't, strictly speaking, talk about *curves* when considering discrete distributions. Additionally, mathematicians usually reserve the word "bell curve" for normal (or Gauss) distributions, which is one particular kind of bell-shaped curve. A normal distribution is usually given by the formula $p(x) = e^{-x^2}$, but you can center and scale this by different values.

When bar diagrams are approximately bell shaped like the above, we will call these "bell curves", though they are strictly speaking not curves, nor Guassian distributions. We will also use the term about non-symmetric distributions, such as the one you get for the sum of the three highest of four d6 (as used in d20 character generation).

It is not difficult to get bell-shaped distributions: The so-called Central Limit Theorem (see `https://en.wikipedia.org/wiki/Central_limit_theorem`) states that if you add the values of independent events (such as dice rolls), the sum approximates a normal distribution more and more closely when the number of events increase. Note that this is true even if each event has a highly uneven distribution of values. All that is required is that the events are independent, or nearly so, and that the number of events is high enough. When adding normal dice, you get fairly good approximations with as little as three dice.

The original use of bell curves in RPGs was in generating attributes – since "real world" attributes supposedly follow a normal distribution, people want this to be true in the game also. However, this is relevant only insofar as the in-game attributes translate linearly into real-world values. While this may be true for height and weight, etc., there is no indication that, for example, intelligence in D&D translates linearly to IQ (which is defined to follow a normal distribution centered on 100).

You can also argue that the same person performs the same task repeatedly, such as throwing a spear as far as she can, the results (e.g, the distances thrown) are normally distributed. This indicates that a bell-curved dice roll is suitable for action resolution. But again, this requires that the quality of results in the game translates linearly to some real-world scale, which is not always the case.

Some games use a logarithmic scale on attributes or results, so a +1 in the attribute corresponds to, say, a doubling of a real-world value (such as size). A normal distribution of the real-world values would map to a highly skewed and not very bell-like distribution in the logarithmic scale. Look up "Log-normal distribution" to see what it looks like.

4 Analysis of dice-roll mechanisms

In this section, we will look at some existing and new systems and discuss them in terms of the properties discussed in section 2. We will sometimes in this discussion calculate probabilities with the methods from section 3, but in other cases we will just relate observations about the probability distribution, which in most cases are obtained by using the **Troll**.

4.1 One die to rule them all

The simplest dice-roll mechanism is to use a single die. The result can be modified with ability, difficulty, and circumstance in various ways. Probability calculation is easy, as each result on a dN has probability $\frac{1}{N}$.

There are many single-dice system, but the best known is the d20 system that originated in D&D. Here, a d20 is rolled, ability is added, and the sum must exceed a threshold determined by difficulty. Some opposed actions are handled by a contest of who gets the highest modified roll, while other opposed actions are treated by using properties of the opponent (such as armour class) to determine a fixed (i.e., non-random) difficulty rating. If the unmodified die shows 18-20, there is a chance of critical success: If a second roll indicates a successful action, the action is critically successful. As far as I recall, there is no mechanism for critical failures in the standard mechanics. Diminishing returns are handled by increasing costs of level increases.

Another single-die system is HârnMaster, where you roll a d100 and must roll under your skill (rounded to nearest 5) to succeed. If the die-roll divides evenly by 5 (i.e., if the last digit is 0 or 5), the success or failure is critical, so there is a total of four degrees of success/failure. The effective skill may be reduced by circumstance such as wounds and fatigue, but difficulty does not directly modify the roll. In opposed actions, both parties roll to determine their degree of success or failure, and the highest degree of succes (or lowest degree of failure) usually wins, but in some situations (such as combat), a table is cross-indexed by the degree of success or failure on bith sides and the effect of the action is read from this table. Ties in degree of success/failure usually indicate a stalemate that can be resolved in later rounds. Diminishing returns are handled by letting increase of skills be determined by dice rolls that are increasingly difficult to succeed (you must roll over your current ability on a d100 to increase it).

Talislanta (4th edition) uses a d20 to which you add ability and subtract difficulty. An "Action Table" is used to convert the value of the modified roll to one of five degrees of success/failure. Opposed rolls use the opponent's ability as a negative modifier to the active players roll. Diminishing returns in Talislanta is (like in d20) handled by increasing cost of skill increases.

So, even with the same basic dice-roll mechanism (roll a single die), these systems are quite different due to differences in how difficulty and circumstance modify the rolls, how opposed actions are handled, how degree of success is determined, and how diminishing returns are achieved.

If we cut each of the systems down to the bone, we have one system where you roll a die, add your ability and compare to a threshold. In the other, you roll a die and compare directly to your skill. Though these look different, they behave the same way if the threshold in the first method is fixed: Increased ability will linearly increase the probability of success (until success is certain). Modifiers applied to the roll or threshold will also linearly increase or decrease the success probability. Such linear modification of probability is the basic property shared by nearly all single-dice systems.

I will pass no judgment about which of the above systems is best (and, indeed, this will depend on what you want to achieve), just note a few observations:

- If opposed actions have both players roll dice and only one roll is made for unopposed action, there is a larger spread of results on opposed actions than on unopposed actions.

- In the HârnMaster system, 20% of all failures and successes will be critical, regardless of ability. The d20 system and Talislanta both give higher proportion of critical successes to higher abilities, though d20 can give no higher chance than 16.7% for critical success (assuming 18-20 give criticals).
- Even though HârnMaster uses a d100, the fact that skills are rounded to nearest multiple of 5 before adding them to the roll means that there are only 20 essentially different die-roll results (if we ignore criticals). Hence, HârnMaster is not really any more fine-grained than systems that use a d20. The steps between each multiple of 5 are, however, used to keep track of progress in experience towards the next "real" skill level. Essentially, the skill track and experience track are merged.
- If both attribute and skill modify the roll, the typical ranges of these will determine if the game favours training over raw talent or vice-versa. d20 translates attributes to modifiers at the rate of two to one, which makes a skill difference of 2 more significant than a difference of 2 in an attribute. Whether this makes skills overall more significant than attributes depends on the typical range of skills and attributes: If they have similar ranges of values, skills will dominate, but if attributes have a much larger typical range than skills, they may still dominate. In the d20 SRD rules, maximum skill levels increase with the character level, so for low-level characters, attribute differences dominate, but for high-level characters, differences in skill will matter more.

4.2 Adding a few dice

A variant of the above is adding up a few dice instead of a single die, but otherwise use the result as above (i.e., adding it to the ability, require it to be less than the ability, etc.).

An example is Stefan O'Sullivan's Fudge system, that to the ability number adds four "Fudge dice" that each have the values -1, 0 and 1 (so a single Fudge die is equivalent to $\mathrm{d}3-2$). This gives sums from -4 to 4 that are added to the ability, which is then compared to the difficulty. The roll has a bell-like distribution centered on 0. Centering rolls on 0 has the advantage that ability numbers and difficulty numbers can have the same range, so you can use the opponent's ability directly as difficulty without adding or subtracting a base value.

Another way of getting zero-centered rolls is the $\mathrm{d}n - \mathrm{d}n$ method: Two dice of different colours (or otherwise distinguishable) are rolled, and the die with the "bad" colour is subtracted from the die with the "good" colour. The distribution is triangular and equivalent to $\mathrm{d}n + \mathrm{d}n$ shifted down $n+1$ places (i.e., to $2\mathrm{d}n - n - 1$). Yet another way of getting the same distribution is, again, to roll a good die and a bad die, but instead of subtracting the bad from the good, you select the die with the smallest number showing and let this count as negative if it is on the bad die and as positive if it on the good die. Ties count as 0. For example, if the good die shows 6 and the bad die shows 4, the result is -4. This way, you replace a subtraction by a comparison, which many find faster. It takes a bit more effort yo explain, though.

Also equivalent to $\mathrm{d}n - \mathrm{d}n$ is to have both sides in a conflict add a $\mathrm{d}n$ to their abilities and then compare the results. For unopposed actions, the GM acts opponent, so he adds the $\mathrm{d}n$ to a predetermined difficulty number. This allows the GM to hide the exact difficulty of an action (or ability of an NPC) from the players by rolling the opposing die secretly. It also means that players get to roll whenever they are involved in an action (even if they are on the receiving end), which keeps them active.

In general, having one side roll $\mathrm{d}n$ and the other roll $\mathrm{d}m$ is equivalent to letting the first side roll $\mathrm{d}n - \mathrm{d}m$ or $\mathrm{d}n + \mathrm{d}m - (m+1)$. So there is no basic difference (apart from constant offsets) between having both sides roll and only one side roll, as long as the the total number of dice rolled is the same.

The advantage of $\mathrm{d}n - \mathrm{d}n$ (or equivalent) over Fudge dice is that you don't need special dice, but you do need players and GMs to be in agreement of which dice are good and bad *before* the dice are rolled. If you use $\mathrm{d}n + \mathrm{d}n - (n+1)$, you don't need this agreement, but you need one more arithmetic operation.

Since zero-centered, symmeteric dice-rolls by construction always have average 0, you can fairly easily take different degrees of randomness into account. For example, with $\mathrm{d}n - \mathrm{d}n$, you can use different n for different tasks: If the task has a low degree of variability, use d4s, if it has average variability, use d8s and if it has high variability, use d12s or even d20s. With Fudge dice, you can use three, four or five Fudge dice in a roll depending on how variable you want the result to be.

All of the above have non-flat distributions (and if more than two dice are involved, the distribution will be a bell curve), so adding a constant modifier will not increase the probability of success by a fixed percentage (as it does in single-dice systems). Some people dislike this by saying that the same modifier benefits some people more than others, but you can argue that this is the case for single-die systems too (if you, for example, look at the *relative* increase in success chance). So, again, it boils down to what the designer wants to achieve.

Instead of adding a number of dice and comparing the sum to an ability rating, you can compare each die value individually to a (lower) ability rating and count how many are below the threshold. You can then directly translate this number into a degree of success. If three or more dice are rolled, the degree of success will have an asymmetric bell-like distribution which is skewed towards low or high results depending on whether the ability is lower or higher than the mean value of a die (though the skew becomes less if more dice are rolled). This mechanism limits the effective range of abilities to the range of a single die, but for games that operate with low granularity of abilities, this won't be a problem. And a d20 should accommodate enough ability levels to satisfy most.

4.3 Linear dice pools

Many games use a system where the ability of the character is translated into a number of dice that are rolled to determine success. We call these "dice-pool systems". Note that we will use this name only when the number of dice used is determined solely by the ability of the roller. Systems that use a variable number of dice that is determined in other ways are treated in Section 4.5.

In some of these systems, the dice are added to a single value. In others, each die is independently compared to a threshold and the number of dice that meet or exceed (or, in some cases, do not exceed) this threshold are counted. Modifiers can modify either the number of dice rolled, the required sum or success count, the threshold towards which the dice are compared, or combinations of these.

Some examples:

- West End Games' d6 system basically adds a number of d6s equal to the ability and compares the sum to a difficulty level.
- White Wolf's "World of Darkness" system rolls a number of d10s equal to the ability and counts the number of results that are 8 or more. Rerolls on 10s increase the average somewhat (from 0.3 per dice to 0.333... per dice) and allow results in excess of the character's ability. See also section 3.4.
- Earlier White Wolf systems were similar but some had a variable threshold for the dice (determined by the complexity of the task) and there were no rerolls (though, in some variants, a 10 count as 2 successes and a 1 cancels a success).

If each increase in ability adds a die to the pool, you will quickly have to roll a very large number of dice unless the range of abilities is limited. White Wolf's systems limit attributes to a range of 1-5 and skills to a range of 0-5, so no more than 10 dice need to be rolled, and that only rarely. West End Games' d6 system (in some versions) has levels between adding a full die, which is another way to limit the number of dice rolled: You go from d6 to d6+1 to d6+2 to 2d6, and so on. Nevertheless, dice pools are usually used in games where there is a relatively narrow numerical range of abilities.

Since the results of each die (which may be the straight value of the die or reduced to a smaller range of values, e.g., 0 or 1) are added, the average result increases linearly with ability, as does the variance, and the distribution approaches a normal distribution when the number of dice is high. This has the effect that characters with higher ability have a larger spread in performance than do novice characters (although only in an absolute sense – the spread divided by the average result decreases). If the results of action rolls translate to real-world figures, this may seem counter-intuitive, but since such translations rarely exist, it is largely a matter of taste whether this is good or bad.

4.4 Nonlinear dice pools

The above-mentioned dice-pools are *linear* in the sense that adding more dice gives a linear increase in the average result. There are also games that use nonlinear dice-pools of various kinds.

One of the more complicated examples is the so-called "One Roll Engine" used in "Godlike" by Hobgoblynn Press. Here, you roll a number of d10 equal to attribute + skill, like in many other dice pool systems, but how you determine your result is quite different: You search for sets (pairs, triples, etc.) of identical dice-values and select one such set. The value on the dice determines how well you succeed, and the number of dice in the set (called the width of the set) how quickly you accomplish your goal. There are optional "hard" and "wiggle" dice, which complicates things s bit, but we will ignore these and just look at the basic idea of looking for sets of dice and taking the set with the highest showing numbers.

We observe that the higher the value, the more likely it is. The reason is that the probability if getting a pair (or more) of a value doesn't depend on the value. All pairs are equally likely, but since you will choose the highest-valued pair if you get more than one, the final result is skewed towards higher values. If the number of dice is low, the skew is fairly small (as the chance of getting two or more sets is small), but at eight d10, a result of 10 is nearly seven times as likely as a result of 1. The chance of getting no sets (i.e., all different values) is initially quite high, but it drops to under 50% at five dice and is less than 2% at eight dice. The average result actually increases more than linearly with the number of dice (as doubling the number of dice more than doubles the average), at least up to the maximum of 10 dice used in Godlike.

There are other nonlinear dice-pool systems. One of the simplest is to roll a number of dice equal to the ability and then pick the highest result, as is done in Dream Pod 9's "Silhouette". Here is definitely a case of diminishing returns: With d10s, the average result starts at 5.5 and gets closer and closer to 10 when the number of dice increase, but will never reach it. The spread of the results decrease with the number of dice, so you can say that you reflect that more able persons are more consistent. However, the effect of the diminishing returns is maybe too strong: Even a rank novice with ability 1 has 10% chance of getting the best possible result (10) and will have an average result that is more than half of what is maximally possible. Additionally, 10 (the maximum) is the most likely result already at ability 2.

To solve this, you can take the *second-highest* result of n dice (where $n \geq 2$). There is still diminishing returns and decreasing spread, but much slower than before. In particular, the chance of getting a result of 10 increases much slower, so it isn't until 14 dice that it becomes the most likely result. The distributions (when $n > 2$) are bell curves skewed towards higher and higher values when n increases. Additionally, a character with skill 2 (the minumim if you pick the second-highest die) has only 1% chance of getting the best possible result, so you don't see novices achieve masterful results as often.

Both the take-highest and take-second-highest method allow a low-skilled person a (low) probability of achieving the best possible result. Some like this possibility, but others want to put an upper limit on the results obtainable by low-skilled persons. A nonlinear dice pool that achieves the latter is that you (as always) roll a number of dice equal to you ability, but then count how many different results you get. This is bounded upwards by both the number of dice and the size of the dice used. There is also diminishing returns and decreasing spread. The main disadvantage is that it takes slightly longer to count the number of different values than to find the highest or second-highest of the values (though not by much). If you use n dM, the average number of different values is $M(1-(\frac{M-1}{M})^n)$. For n d10, this simplifies $10(1-0.9^n)$.

Instead of counting different dice, you can also count identical dice. This is like the "width" of rolls in "Godlike", but not requiring at least two identical dice and not caring about the numbers shown on the dice (only the size of the largest set). This method also bounds the result by the number of dice rolled (but not by their size) and also has diminishing returns, albeit less so than if you count different dice. While counting different dice gives you a saturation curve (when plotting the average against the number of dice), counting identical dice gives you a curve that grows only a bit slower than linearly. Note that when you count different dice, larger dice will give you larger averages (as there are more different values), but when you count identical dice, larger dice give lower averages (as dice are less likely to show the same value). I would prefer using the count-different system with fairly large dice (at least d10s) to get a reasonable range of results, while I would not use the count-identical system with dice larger than d6, as larger dice give too small increases in average when you add dice. For example, if you use d10s, the average increases only by around 0.2 for each extra dice you add.

4.4.1 Summary of dice pools

Dice pools are more cumbersome than systems that use a single die or a small fixed number of dice, and they tend to scale badly (since you don't want to roll more than a handful of dice). But they give a direct translation of ability to roll, so you avoid having to add or compare the ability to the outcome of a roll.

You can modify difficulty by modifying the number of dice rolled in addition to modifying the number you must beat. For example, if you use the pick-second-highest method, reducing the number of dice by one will affect low-skilled persons more than high-skilled persons (due to the diminishing returns) while adding one to the target number affects all more evenly. This gives the GM a bit more flexibility when setting up challenges.

Dice pools allow all sorts of extra effects, such as treating specific numbers in special ways, such as rerolling tens or letting ones cancel successes. Or you can read several numbers out of the same pool, such as the width and height of "Godlike" rolls.

Overall, dice pools allow a wide range of design options that you can consider in your design. But by the same token, you must take care not to design a system that looks cool but is far too cumbersome in practice.

4.5 Other dice-roll systems

Some dice-roll systems defy categorisation in the above classes. I will look at a few of these below.

4.5.1 Letting ability determine dice-size

The original "Sovereign Stone" game from Corsair Publishing (before it was assimilated by the d20 Borgs) used a then fairly novel idea: Attributes and skills are given as dice types. So an attribute can range from d4 to d12 (with non-human attributes of d20 or d30 possible) and skills can range from 0 (nonexistent) through d4 to d12. When attempting a task, you roll one die for your relevant attribute and another for your relevant skill and add the results. If the sum meets or exceeds the difficulty of the task, you succeed. The more recent "Serenity" RPG by Margaret Weis use a similar system, but starts from d2 instead of d4 and when you go past d12, you go to d12+d2, d12+d4, etc., instead of to d20 and d30. In both systems, ratings over d12 are considered exceptional.

An advantage of this system is that you (until you exceed a d12 rating) only need one addition to make a roll that takes attribute and skill into account, where adding skill and attribute to a die roll requires two additions. Additionally, the numbers are likely to be smaller, which makes addition faster.

The disadvantage is that you need to have all types of dice around, preferably at least two of each, and you need time to select those that you need in a particular roll. Additionally, the range of dice gives only five (or six) different values for attributes in the unexceptional range. This is fine for many genres, but not for all.

It gets a bit more interesting if we look at the average and spread of results as abilities increase. If you add a dm and a dn, the average is $\frac{m+n}{2}+1$ and the variance is $\frac{m^2+n^2-2}{12}$. If $m = n$, the average is $n+1$ and the variance $\frac{n^2-1}{6}$. This makes the spread (which is the square root of the variance) increase slightly faster than linearly in the average, which means that more skilled persons have higher spread – even relative to their average – than persons of lower skill. The main visible effect is that even very able persons have a high chance of failing easy tasks.

This observation has made someone suggest that higher abilities should equate smaller dice and low rolls be better than high. Though this makes higher skilled persons more consistent and prevents them from getting the worst possible results, it gives novices a fairly high chance of getting the best achievable result ("snake eyes"), which may be a problem if you want to make sure extremely able characters will almost always beat fumbling amateurs.

A system similar to the Sovereign Stone / Serenity system is used in Sanguine Productions' games, such as "Ironclaw" and "Usagi Yojimbo". Here, three dice are rolled: One for attribute, one for skill and one for career. Instead of adding the dice, each is compared against two dice the GM rolls for difficulty. If one of the player's dice is higher than the GM's highest, the player succeeds, if two are higher, the player gets an overwhelming success. If all are smaller than the GM's *smallest* die, the player gets an overwhelming failure (the remaining cases are normal failures). Like the Sovereign Stone / Serenity system, you get increased spread of results with higher abilities. Also, since difficulties are rolled rather than being constants, high difficulties can sometimes be quite easily overcome (if the GM rolls low). All in all, this makes results quite unpredictable, with experts sometimes failing at simple tasks and novices sometimes succeeding at complex tasks. This will fit some genres, but not all.

Another system that lets ability determine dice size is "Savage Worlds" by Pinnacle Entertainment Group. It differs from the above in several ways: Only one die is determined by the ability, another is (for player characters) always a d10, you take the largest result instead of adding the two dice, and each die can "Ace", i.e., is re-rolled and added to the original value whenever it shows the highest possible value. This is an example of unbounded rolls as described in Section 3.4. You succeed if the result is greater than or equal to a target number. This has the strange behaviour that it is sometimes easier to succeed with a d4 than with a d6, and so on. Consider a target number of 6. A d6, obviously, has one chance in 6 of reaching this. To reach a value of 6, a d4 needs to first roll a 4 (to Ace) and then at least 2. This has a combined probability of $\frac{1}{4} \times \frac{4}{5} = \frac{1}{5}$, which is greater than $\frac{1}{6}$. With similar reasoning, a d6 is preferable to a d8 if the target number is 8, a d8 is better than a d10 if the target number is 10 and a d10 is better than a d12 if the target number is 12. Adding the extra d10 does not change this fact, though it makes the difference smaller, as it is likely that the d10 will give a higher value than the other die.

4.5.2 Median of three dice

The usual way of getting bell curves is by adding several dice (or counting successes, which is more or less the same), but you can do it also by comparing dice. A simple method is to roll three dice (e.g., d20s) and throw away the largest and smallest result, i.e., pick the *median* (middle) result. The advantage is that it requires no addition, so it is slightly faster than, say, adding 3d6. We will use the abbreviation "median 3dn" for the median of three dn. This notation is recognised by **Troll**.

We can calculate the probability of getting a result of x with median 3dn by the following observation: The median is x either if either two or three dice come up as x, or one dice is less than x, one is equal to x and one is higher than x. The probability of all three coming up as x is $\frac{1}{n^3}$. The chance that exactly two come up as x is $3 \times \frac{1}{n^2} \times \frac{n-1}{n}$ (the 3 comes from the three places the non-x die can be). The chance that there is one less than x, one equal to x and one greater than x is $6 \times \frac{x-1}{n} \times \frac{1}{n} \times \frac{n-x}{n}$ (the 6 comes from the 6 ways of ordering the three dice). We can add this up to

$$
\begin{aligned}
& \frac{1}{n^3} + 3 \times \frac{1}{n^2} \times \frac{n-1}{n} + 6 \times \frac{x-1}{n} \times \frac{1}{n} \times \frac{n-x}{n} \\
= & \frac{1 + 3(n-1) + 6(x-1)(n-x)}{n^3} \\
= & \frac{3n - 2 + 6(x-1)(n-x)}{n^3} \\
= & \frac{-6x^2 + 6(n+1)x - (3n+2)}{n^3}
\end{aligned}
$$

For example, the chance of getting 7 on mid 3d10 is $\frac{3 \times 10 - 2 + 6(7-1)(10-7)}{10^3} = \frac{136}{1000}$.

Since the curve is a parabola, it is arguable whether it can be called a bell-curve (it lacks the flattening at the ends), but it is a better approximation to a bell-curve than 2dN. You can get closer to a "real" bell curve by taking the median of 5 dice instead of 3. The probabilityy of getting a result of x with median 5d$n$has the formula

$$\frac{30(nx - x^2 + x - n)(nx - x^2 + x - 1) + 10n^2 - 15n + 6}{n^5}$$

For a given range of values, the curve obtained by taking the median of three dice is somewhat flatter than what you get by adding three dice, while the one you get by taking the median of five dice is slightly steeper than that of adding three dice. For example, all of the dice rolls below have ranges from 1 to 10 and averages of 5.5:

Roll	Spread
mid 5d10	1.91
3d4-2	1.94
mid 3d10	2.25
d10	2.87

The value of a median roll is typically used in the same way as the value of a single die or the sum of a few dice (as described in Sections 4.1 and 4.2).

If the value of the median roll is compared to ability (i.e., you must roll under your ability), the method is a special case of the method described at the end of section 4.2, except that you don't distinguish degrees of success and failure: If at least half of the individual dice meet the target, it is a success, otherwise a failure.

A variant is to combine median rolls with the idea of letting abilities equal dice types. So you would roll one dice for your attribute and another for your skill, but what about the third? You can add a third trait (like the career in "Ironclaw"), you can let the third die always be the same (e.g., always a d10), or you can duplicate the skill die, so you roll two dice for your skill and one for your attribute. This makes skills more significant than attributes and limits the maximum result to the level of the skill. Regardless, you still have the effect of spread increasing with ability that you get when the dice-sizes increase with ability.

4.5.3 Adding a subset of the rolled dice

In d20, it is common to generate attributes by rolling four d6 and then add only the largest three results. This generates an assymetrical bell curve but retains the same range of results as just adding 3d6. Compare the two figures below:

```
sum 3d6:

   Value   2.16 bars per %
      3 :  |
      4 :  |||
      5 :  ||||||
      6 :  ||||||||||
      7 :  |||||||||||||||
      8 :  |||||||||||||||||||||
      9 :  |||||||||||||||||||||||||
     10 :  |||||||||||||||||||||||||||
     11 :  |||||||||||||||||||||||||||
     12 :  |||||||||||||||||||||||||
     13 :  |||||||||||||||||||||
     14 :  |||||||||||||||
     15 :  ||||||||||
     16 :  ||||||
     17 :  |||
     18 :  |

Average = 10.5
Spread = 2.95803989155
Mean deviation = 2.41666666667
```

```
sum largest 3 4d6:

   Value   2.16 bars per %
      3 :
      4 :  |
      5 :  ||
      6 :  ||||
      7 :  ||||||
      8 :  ||||||||||
      9 :  |||||||||||||||
     10 :  ||||||||||||||||||||
     11 :  |||||||||||||||||||||||||
     12 :  ||||||||||||||||||||||||||||
     13 :  |||||||||||||||||||||||||||||
     14 :  |||||||||||||||||||||||||||
     15 :  ||||||||||||||||||||||
     16 :  ||||||||||||||||
     17 :  |||||||||
     18 :  |||

Average = 12.2445987654
Spread = 2.84684444531
Mean deviation = 2.31853947569
```

Note that the graphs are turned sideways compared to the usual presentation[1]. The main difference is that the very high values are much more common: It is 3.5 times as likely to roll 18 with the add-largest-three-of-four roll than it is with an unmodified 3d6 roll.

You can generalize this method to rolling $n+m$ dice, select n of these and add them. The result will have the same range of values as just adding n dice, but you can modify the distribution. Selecting the lowest dice will slant the distribution towards low values and selecting the highest dice will slant it towards the high end of the range. You don't have to select the largest or smallest dice, you can, for example, add the smallest and the two largest of 4d6 or discard the smallest and highest and add the rest (which shows that this is a generalisation of the median-of-three method described earlier).

You can combine a kind of dice pool with picking one of the dice. Consider this system:

When a character with ability A attacks an opponent with ability D (or attempts a task with difficulty D), roll $A+D+1$ d6, remove the A lowest and the D highest results, leaving one value between 1 and 6. If this is 4 or more, the attacker succeeds, otherwise he fails. You can add more degrees of success and failure by letting 5 and 6 give better successes than 4 and 1 and 2 worse failures than 3.

This system is symmetric in the sense that it doesn't matter whether the attacker rolls to see how successful his attack is or the defender rolls to see how successful his defence is, as long as an outcome of 6 for an attacker is the same as a 1 for the defender and so on. It allows abilities from 0 upwards, as there

[1] The graphs were made with **Troll**

will always be at least one die to roll even if both sides have ability 0. It is always possible to get any result from 1 to 6 regardless of the difference in ability, but it becomes increasingly unlikely to succeed (let alone well) if the difficulty is much higher than the ability, just as it gets increasingly unlikely to fail if the difficulty is much lower than the ability. Two opponents of equal ability have equal chance of winning, but the spread of results decrease when they both increase in ability. There are diminishing returns: The advantage of increasing ability by one step is smaller if it is already high.

The main disadvantage of this method is that the number of dice rolled is high compared to traditional dice pools, where the number of dice depends only on the ability of the attacker. So ability scores should be kept relatively low.

5 Avoiding rolls

If you use dice pools or other systems that involve multiple dice, dice rolls tend to take fairly long to execute: You have to pick up a largish number of dice (and be sure that the number is correct) and then do something afterwards that involves inspecting all the dice. Even more time is needed if there is a possibility of rerolls (as in the "World of Darkness" or "Savage Worlds" systems). Such systems are best suited to games where you don't roll all the time, i.e., where nearly all conflicts are resolved in one or two rolls or, even better, where most conflicts don't even require rolls to resolve.

You can always resolve conflicts without rolling if you can see that even the best (or worst) possible roll won't make a difference to the outcome. In other cases, players may insist on rolling just for the off chance that they (or their opponent) will make an extreme roll, no matter how unlikely this is. To avoid this, you may want a system where results are capped both upwards and downwards by ability, so people with low ability can not hope to beat persons of high ability (unless there are special circumstances that modify the odds). Some of the above systems do that. For example, a simple additive dice pool with $n\mathrm{d}X$ will have a range from n to $n \times X$, so a person can not hope to beat an opponent that has a skill that is X times higher than his own. If X is high, this is rarely going to happen, though, so you might want a narrower range of results, for example by counting successes instead of adding values.

Even when success (or failure) is not certain with a normal roll, you can use a rule similar to the "take 10" of the d20 system: Instead of rolling, a player (or GM) can simply "take" a predetermined value. If one side in a conflict can clearly win the conflict by taking instead of rolling, you don't need to roll. Players in unstressed situations could be allowed to "take" slightly below their average result (the average of a d20 is 10.5, so "take 10" does this), and in stressed situations they can be allowed to "take" a value that they would achieve in most rolls (say, 80% or more of all rolls), but thereby forgoing the possibility of a better result.

You would need an easy way from the ability of a character to determine the result of "taking" a roll in either stressed or unstressed situations. In linear dice pools, the average is fairly easy to calculate: For n dice, it is simply n times the average of one roll. Finding a threshold that you will meet in 80% of all rolls is less easy, however, as the spread does not increase linearly with the number of dice. For nonlinear dice pools, even the average may have a nontrivial relation to the number of dice rolled (see, for example, the formula for averages for the "count different" method described above). You can precalculate the distribution for all possible ability numbers and provide a table of "take" values, but that is a bit cumbersome to use in play unless the table is small enough to fit on the character sheet. So it might be necessary to make fairly coarse approximations. For the method where you roll n d10 and use the second-highest value, it may be reasonable to say that you can take n (though max 8) in stressed situations and $n + 1$ (but max 9) in unstressed situations.

6 Scale

If you want to handle powers of vastly different scales in the same game, some mechanisms break down or become unmanageable. For example, if the number of dice you roll is equal to your ability, what do you do if the ability is a thousand times higher than human average? If you compare a roll against ability, all abilities higher than what the dice can yield are effectively equal and impossible to beat. And if you add a die to an ability, the die becomes increasingly irrelevant as abilities increase.

The first thing to do is consider how in-game values translate to real-world values. If this translation is linear (e.g., each extra point of strength allows you to lift 10kg more), you get very large in-game numbers. If you, instead, use a logarithmic scale (e.g., every point of strength doubles the weight you can lift), you can have very large differences in real-world numbers with modest differences in in-game numbers.

Doubling the real-world values for every step is rather coarse-grained, so you may want a few steps between each doubling. A scale that has been used by several designers is to double for every three steps. This means that every increase multiplies the real-world value by the cube root of two (1.25992). This is sufficiently close to 1.25 that you can say that every step increases the real-world value by 25%. Another thing that makes this scale easy to use is that 10 steps multiply the real-world value by almost exactly 10 (10.0794, to be more precise), so you can say that 10 steps in a factor of 10 with no practical loss of precision. Some adjust the scale slightly, so 10 steps is exactly a factor of 10, which makes three steps slightly less than a doubling, but close enough for practical use. The decibel scale for sound works this way: Every increase of 10dB multiplies the power of the sound by 10.

Even with a logarithmic scale, you may get numbers that are too large to be practically manageable (e.g., in dice pools or when you use dice-types to represent abilities), so you can add a scaling mechanism that say that for every increase of N in ability, you are at one higher scale.

When resolving an action between two entities, you reduce (or increase) both scales such that the weakest of the two entities is at scale 0. For example, if $N = 10$ (i.e., every increase of 10 increases scale by one) and a struggle between one entity of ability 56 and another of ability 64, you reduce it to a battle between abilities of 6 and 14. You can add an additional rule that if the difference in scale is more than, say, 2, then you don't roll: The higher scale automatically wins. This seems reasonable enough: If a step in scale corresponds to ten times the power, two steps correspond to 100 times the power, which shold be enough to give certain victory.

Issaries' "Heroquest" game integrates a scale system directly: A unit of scale (or "Mastery") is a difference of 20, and you denote an ability as a number between 1 and 20 plus a number of Masteries. You must roll under your ability number on a d20 to succeed, but each Mastery increases the degree of success by one (or reduces that of the opponent by one). Effectively, this means that if you have sufficiently more masteries than your opponent, you can never lose.

One thing to note, though, is that not all abilities have quantifiable real-world measures: How do you measure beauty numerically? Or agility, leadership or intelligence? The latter does have a numeric measure (IQ), but this is an artificial construction, defined to average to 100 and be normally distributed around this with a predefined spread, so saying that someone with an IQ of 160 is twice as intelligent as one with IQ 80 is meaningless. When games assign numbers to unquantifiable properties and abilities like these, the numbers are as much a construction as IQ numbers, and can really only be used to determine which of two characters are better (or how well they do against equally abstract difficulty numbers).

7 Other randomizers

Here I will briefly look at other ways of bringing randomness into games.

7.1 Cards

Next to dice, cards seem to be the most common randomizer in RPGs. Some games (like R. Talsorian's "Castle Falkenstein") use standard playing cards, others (like TSR's "Saga" system) use custom card decks.

Drawing cards multiple times from the same deck are, unlike die rolls, not independent events: Drawing an ace in the first draw reduces the chance of drawing an ace in the second draw as well. Some argue that this makes cards superior – "luck" would tend to even out, as you will get all numbers equally often if all cards in a deck are drawn before it is reshuffled. But this assumes that all draws are equally important, which I find questionable.

The main advantages of using cards over dice are:

- Players can keep hands of cards and choose which to play.
- You can use the suit or colour of the card as well as its value to affect the outcome in different ways. Both "Saga" and "Castle Falkenstein" do this.

- With custom-made cards, you can have text on the cards that provide for special effects, such as critical results.

Not all of these will be relevant to all games, though, and some may dislike having the players choose their "luck" from a hand of cards, as it brings meta-game decisions into the game world. Additionally, players may do unimportant tasks simply to get rid of bad cards in a safe way.

Probability calculation of card draws is more complicated than for dice since the draws are not independent events but depend on the history of already-drawn cards. The probability of drawing, say, an ace is fairly simple if you know how many aces are left in the deck and the size of the deck: You simply divide one by the other. It gets more complicated if you want to figure out, for example, the probability of drawing two cards of the same suit end even worse if you do not know all the cards that area drawn (such as the hands of other players).

Drawing coloured beads or some such from a bag is not essentially different from drawing cards from a deck.

7.2 Spinners

Spinners are really just dice of a different shape – you (usually) get equal probabilities of a finite number of outcomes. The main advantage of spinners is that you can make them in sizes (number of outcomes) that you don't find on dice, such as 7 or 11, and you can get unequal probabilities of different results, (depending on the size of the corresponding pie slices). Additionally, spinners can be cheaper to make than specialized dice, such as Fudge dice.

7.3 Rock-paper-scissors

While this strictly speaking isn't random, it is unpredictable enough that it can be used as a randomizer. Its sole advantage is that it doesn't require any equipment or playing surface, which makes it popular in live role-playing.

There are really only three outcomes: Win, lose and draw, each of which have equal probabilities (assuming random or unpredictable choices), but you can add in special cases, such as special types of characters winning draws on certain gestures. For example, warriors could win when both hands are "rock", magicians when both are "paper" and thieves when both are "scissors", but that will make people second-guess and not choose "rock" when going against a warrior, so the value of having such special cases is somewhat dubious.

Rock-paper-scissors only works with two players, as you otherwise can get circular results (A beats B, who beats C, who beats A), which can be hard to interpret. But since most RPG dice-roll systems share this limitation, it is not a major issue.

There are generalisations of rock-paper-scissors to five or seven different values (adding "Lizard" and "Spock" to the set, for example) that reduce the chances of draws and avoid cycles with three or four players, but these tend to be harder to remember.

8 Some personal opinions

Here, I will discuss a few personal preferences and pet peeves. Unlike the above, where I have (mostly) tried to be neutral and objective, rampant subjectivity will abound. So you are warned.

8.1 Rerolls

I don't like rerolls. They take extra time that you can't decrease much by experience – the physical action of rolling again can't really be sped up. This is in contrast to time used for calculations based on a single roll, such as adding up the dice or counting successes, which you can decrease almost arbitrarily with training. Experienced players can, for example, add up a dozen six-sided dice in a few seconds.

Additionally, repeated rerolls remove upper limits on rolls – anyone can conceivably achieve fantastical results, just not very often. It can spoil any game if a street kid kills the dragon that menaces the town by throwing a stone at it, just like it will spoil the game if a high-level PC is killed fighting a mouse. Also, the absence of an upper limit will make players insist on rolling even when they are hugely overpowered, on the off chance that they will reroll a dozen times. Sure, the GM can forbid such silliness, but then why have unlimited results at all? You can prevent players from insisting on rolling in ridiculous situations if doing so not only carries a tiny chance of success but also a significant risk. This does not avoid the potentially ridiculously high degrees of success that unlimited rerolls allow, though.

Some argue that heroic fiction abounds with cases where the hero wins over a vastly more powerful foe (such as Bard the bowman in "The Hobbit" killing Smaug with a single arrow), but this is usually because the hero has outrageous skill or magical equipment rather than outrageous luck.

A third problem I see with rerolls is that they make the probability distribution uneven: You can get holes in the distribution (i.e., impossible results) or places where the probability drops sharply but then stays nearly constant for a while.

8.2 Methods that are different just to be different

Many games, especially Indie or homebrew, feature "innovative" dice-roll mechanisms that seem to be different just for the sake of being new. Like *haut couture* fashion, these are often overly complex and don't seem to add anything other than colour and strangeness to what they replace. Such design can serve a legitimate purpose – it can get your game noticed where yet another d20 game won't. But it is better to combine innovation with purpose – make the new system do something that can't (as easily) be achieved with existing systems, without sacrificing the good properties of tried-and-tested methods.

Though I may get flak for this, I find Godlike's One Roll System (as described in section 4.4) to be an example of *haut couture* dice mechanisms – it is quite complex, the probabilities are weird and it doesn't seem to do much that you couldn't achieve by simpler methods.

9 Sounding off

I haven't covered all resolution mechanisms – partly because of space, partly because of defective memory and just not knowing them all. And I'm sure many new dice-roll mechanisms will be invented, some to be deservedly forgotten again and others to be copied over and over with minor variations.

Whether you plan to use an existing method or invent your own, I hope this paper has given you something to think about when doing so. And stay tuned – I will at uneven intervals (possibly years) release new versions of this document to the public (compare the dates to see if you have the newest version). An earlier version was published as a column called "Roll the Bones" on RPG.net.

Appendix: Dice as physical objects

What dice are fair?

In the above, we have assume that dice of any size exist, for example d7 or d13, though these are not in the common mix of gaming dice. Some companies sell dice with non-standard sizes, but it is questionable if they are all fair, i.e., that all sides are equally likely.

When can we be sure a die is fair? It is not enough that every face has equal area. You can construct a polyhedron where every face has the same area, but where the polyhedron can only rest on a subset of these, so this is obviously not a fair die. It is not even enough that the faces all have the same shape and size, as you can make a similar construction even then by varying the angles that connect the faces. So the only way to be sure that all faces are equally likely is to require symmetry: No matter what face it rests on, the entire polyhedron should look the same to an observer (possibly excepting mirroring). In other words, if there is no labeling on the faces, you should not be able to distinguish them, even by looking at the whole polyhedron, again modulo mirror images.

The Platonic solids (the traditional d4, d6, d8, d12 and d20) have this property, but so do the Catalan (or dual Archimedean) solids, see
`https://en.wikipedia.org/wiki/Catalan_solid`.

There 13 of these with 12, 24, 30, 48, 60 and 120 faces. The 30-sided Catalan solid called “the rhombic triacontahedron” is found as a d30 in many game stores, and I have seen a d24 based on the tetrakis hexahedron. For a d24, I would have preferred a deltoidal icositetrahedron or pentagonal Icositetrahedron, as the faces have lower aspect ratio, so larger symbols can fit inside the faces). They also look cooler.

The astute reader will have noticed that the common d10 is not among the above, and there are indeed more fair dice than the Platonic and Catalan solids. The d10 is constructed by joining two “pyramids” constructed by kite-shaped sides in such a way that the convex vertices of one pyramid fits the concave vertices of the other. This construction, called a “trapezohedron”, see `https://en.wikipedia.org/wiki/Trapezohedron` can be generalized by joining two “pyramids” of $N > 2$ kite-shaped sides in a similar way to make a polyhedron with $2N$ sides. If N is odd, there will be a face facing up when the polyhedron rests on a flat surface, which is why this shape is chosen for the d10. Note that for $N = 3$, the construction yields a normal cubical d6. I have seen a 34-sided die constructed this way, used for an earlier version of the Danisk national lottery, which used to have 34 different numbers (it now has 36).

You can also join two “normal” pyramids each made of $N > 2$ triangles to get a polyhedron with $2N$ sides called a “bipyramid”, see `https://en.wikipedia.org/wiki/Bipyramid`. This will have an upwards-facing face if N is even, so it can used for, e.g., a d14. I have seen a d36 constructed this way (probably made for lotteries). This construction includes the traditional Platonic octahedral d8.

Assuming all faces are stable (i.e., the die will not topple over when resting on a face), these are all that are possible, see `dicephysics.info/thesis7.doc`. This, for example, proves that you can not make a fair d5 as a triagular prism, as the probability of landing on one of the triangular ends instead of one of the rectangular sides depend not only on the length of the prism but also on the energy given to the die when thrown, so you can skew the results by throwing the die more or less energetically.

If we allow non-flat faces and faces that the die can not (stably) rest on, we have more options, a coin being the most commonly used of these options.

You can make a die as a prism with rounded or tapered ends. Some ancient Roman games, for example, used four-sided sticks as randomizers, and you can similarly make any dN as an N-sided prism with rounded or tapered ends. Only if N is even will this have a face on top when resting, so you need either labelleing on edges or a different labelling scheme, which we will discuss below. I have heard of students in schools that forbid bringing dice to school use pencils with hexagonal cross-section as a replacement for dice.

You can also use antiprisms (see `https://en.wikipedia.org/wiki/Antiprism`) with rounded or tapered ends as dice. These will also have a face facing up if N is even. These are often called “barrel dice”.

You can also make a dice as a pyramid with a rounded base. If the triangular faces are steep enough, this can not rest on the rounded base. You can make any dN this way. If N is odd, an edge will face up, but you can put the numbers on the base to make it clear which side is up (which makes this a kind of spinner). This would make a fine way to make dN for odd N, but only if N is relatively small.

Though there are a few other (and rather strange) ways to make provably fair dice (see `http://loki3.com/poly/fair-dice.html`), for example by labelling pairs of faces with the same number, these do not really add anything useful.

You can find other shapes of dice in some game stores, the most common being:

- A d5 made as a triangular prism with numbers on the triangular ends as well as on the rectangular sides.
- A d7 made as a pentagonal prism with numbers on the pentagonal ends as well as on the rectangular sides.
- A d100 that looks like a golf ball (the “Zocchihedron”).

These lack the symmetry between faces, so they are not obviously fair. But could they be? If we look at the prismatic d5, it is easy to see that the longer the rectangular sides are, the more likely it is that the die will land on one of these. Furthermore, we can make this likelihood arbitrarily small or big. So it would seem reasonable to assume that there exists a length where the probability is the same for landing on each end as for landing one each side. But, as mentioned above, it has been proven that the energy of the throw influences the probability, so I would not consider this fair.

The Zocchihedron is somewhat different. It is sufficiently close to being a sphere that the energy of the throw probably doesn’t matter, but I still don’t believe it to be fair, as the spacing between the circular "faces" is not constant. This view is supported by a test made by Jason Mills for White Dwarf magazine, which concluded that results under 8 or over 93 are considerably less likely than other results (see `https://en.wikipedia.org/wiki/Zocchihedron`). Also, I see little point in a physical d100, as the markings are too small to read easily and it takes forever to stop rolling. Using two d10s, one for the tens and one for the units, achieves the purpose and is quite easy to read, especially if one of of the dice is labeled 00, 10, 20, 30, ... Most dice sets these days include both a d10 labeled this way and a d10 labeled from 0 to 9.

But do dice have to be fair? You can argue that as long as it all players use the same dice, these need not be fair. Ancient games used knuckle-bones and the more recent game "Pass the Pigs" (`https://en.wikipedia.org/wiki/Pass_the_Pigs`) uses pig-shaped dice that are highly irregular in shape. As mentioned above, the energy used in the throw can affect the probabilities of irregular dice, so you should at least standardise the throwing method if you allow these.

9.1 Dice wish list

Not all of the Catalan solids mentioned above can be bought as dice. Some of these, i.e., the two 12-sided Catalan solids, are redundant, as the Platonic d12 can be used instead. They could be used as novelty dice, similar to the barrel dice mentioned above, though. I quite like the idea of a rhombic dodecahedron as a d12, as I find the shape appealing. A d24 is quite useful (such as for randomly rolling the time of day) but, as mentioned, I would prefer a different shape than the one you can currently get. Catalan solids with more than 30 faces are probably not very useful, but they look impressive.

But what is really missing from the standard RPG dice selection are dice that continue the sequence started by d4, d6, d8, d10 and d12. There is a long gap until d20, which could be filled with a d14, a d16 and a d18. I have seen a d14 (made in the same way as a d10) and a d16 (made as two base-to-base 8-sided pyramids), but the 18-sided dice I have seen are not symmetric and, hence, suspect.

Dice with an odd number of faces are also useful. These are easily made by labelling the faces of a die with twice the sides with each number occurring twice. For example, a d3 can be made as a d6 with the numbers 1 to 3 twice each. Making an odd-sided die using a prism gives the challenge that an edge is facing up when the die is resting. An N-sided pyramid with rounded base and numbers on the base could be a good way of making a dN for odd N.

Labelling dice

The most obvious way is to label the sides of an N-sided die from 1 to N, but there are examples that differ from this norm:

- A d10 is often labeled 0,. . . ,9 or 00,. . . ,90 instead of 1,. . . ,10, as this makes it easier to use two d10 as a d100.
- Early editions of D&D came with 20-sided dice that were labelled with 0 to 9 twice for use as d10s, as the trapezohedral d10 was not available then.
- Fudge-dice are cubical d6s labeled with -1, 0 and 1 twice each.
- The doubling-die used in Backgammon has the numbers 2, 4, 8, 16, 32 and 64 on a cubical d6. This is not used as a randomizer when playing Backgammon, but it could be used as such in an RPG if you want an exponential progression.
- You can get d3s that are cubes labeled with 1, 2 and 3 twice each.
- In some board games, dice are labeled with non-numerical symbols. For example, the "Lord of the Rings" board game by Reiner Knizia uses a d6 that is labeled one one side with The Lidless Eye of Sauron and other sides being blank or showing dots or outlines of cards. The RPG "The One Ring" replace some numbers on d6s and d12s by special symbols, including a Gandalf rune and The Lidless Eye. The Danish board game Ludo (similar to Parcheesi) has replaced the numbers 2 and 5 on a d6 with a star and a globe that have special meaning during the game, and you can also get Poker dice that use card symbols.

Odd-numbered dice are often made by labeling a die with twice the required sides with two occurrences of each number, such as the d3 mentioned above.

All of the above have equal occurrences of all the used numbers, but you can also make dice that have different numbers appearing a different number of times. For example, you can have a d6 with three occurrences of 1, two occurrences of 2 and one occurrence of 3 for a d3 that is skewed towards low numbers.

Traditional six-sided dice don't use number symbols, but label each side with a number of "pips" from 1 to 6 in standardized patterns. The standard patterns and the rule that opposing sides sum to 7 have been used since ancient Greece, and it is only after the introduction of non-cubic dice for role-playing games that number symbols have been common for labeling dice. A few non-cubical dice have used pips instead of numbers, but such is not common.

Placement of symbols

The placement of symbols on dice might seem unimportant: Since all sides are of equal probability, just place the symbols in any order. There are, however, a few special cases to consider as well as traditions that, while not important, should be kept in mind for the sake of aesthetics.

The first special case is the d4. Unlike most other dice, this does not have a face opposite the one on which it rests, so you can't read the result off a top face. Early d4s, as supplied with the original Dungeons & Dragons game, used the rule that the result is determined by the face that rests on the table, but since you can't see this, the numbers were put on the edges of all the neighbouring faces, so you would read the value at the bottom edge of any visible face. Most current d4s read the value from the vertex that is on top, putting the symbols near the corners of every triangle (again in three copies each), so you will read the value from the top corner of any showing faces. I prefer the latter. You can also get d4s in other shapes, such as anti-prisms or cubes with two rounded-off faces, where there is a face facing up.

I have seen d5s and d7s that are pentagonal and heptagonal prisms with rounded ends. Like the d4, these do not have a face on top when they rest. The ones I have seen rule that you read the value on the one of the two topmost faces that are nearest to you, which is easily determined if you roll the dice away from you. An alternative is to put the numbers at the edges, but that isn't very aesthetic (though I have seen it done). For the d7, there is another option where you use pips: Add the number of pips on the two topmost faces. If the faces have the following numbers of pips (in sequence): 0, 1, 2, 2, 3, 4, 2, you get 0+1 = 1, 1+2 = 3, 2+2 = 4, 2+3 = 5, 3+4 = 7, 4+2 = 6 and 2+0 = 2. There is no such arrangement for the similar d5 unless we make it produce values from 0 to 4 instead of 1 to 5, in which case 0, 0, 1, 2, 2 will work.

Using a similar system for a tetrahedral d4 will not work for several reasons: You need to look at a d4 almost directly from above to see three sides, and you can't get make a labelling such that you can get all values in the range 1 to 4 by adding the pips on the three visible faces.

As mentioned above, the numbers or pips on a d6 are traditionally placed so opposing faces add up to 7. There are only two ways of doing this, which are mirror images of each other. Most manufacturers of polyhedral dice follow the generalized version of that rule: The sum of opposing faces is constant.

If you use the rule for a dN labeled 1,...,N, the sum of opposite edges should be $N+1$, and for a dN labeled 0,...,$N-1$ (such as a d10), opposing sides should add to $N-1$. For dice larger than d6, there are several non-symmetric ways of obeying the constant-sum rule. All the d10s I have agree on having 0, 8, 2, 6, 4 clockwise around one vertex and the odd numbers around the opposite vertex so opposite faces add up to 9, but for other polyhedral dice there does not seem to be any consensus about which of the possible constant-sum arrangements to use.

I have seen a few dice that don't follow the constant-sum rule, such as a d8 that has 6 opposite of 1 and a d20 that has 12 opposite to 2. While these are equally good randomizers as dice that follow the constant-sum rule, I find that they grate on my sense of aesthetics.

I find dice that are labelled 0 to $N-1$ useful as alternatives to dice labelled 1 to N, so I would like to see such for other than ten-sided dice. Obviously, you can use a normal dN and simply rule that an N is read as 0, but that is not as pleasing as a having a real zero.