1 The concept of chance

1 An unlucky gamble

Suppose you were offered the chance to play a simple gambling game, in which you are invited to bet on the outcome of a die-roll. There are only two bets allowed. You can wager that the die will land 6, or you can wager that it will land any of 1, 2, 3, 4, or 5. In either case, if your wager is successful, you will win the same prize: one dollar.

So the bets are:

Die lands 1–5 Pays $1.
Die lands 6 Pays $1.

Assume that you know, moreover, that the die has no significant asymmetry in its construction. It does not have a physical bias to one or more sides.

Which bet ought you to take? Assuming you would prefer more money to less, it is obvious that you ought to take the bet on 1–5, rather than the bet on 6.

Now suppose that you *really do* play this game, and you play it at the same time as a friend. You sensibly choose to bet on 1–5. Your friend, bizarrely, insists that she has a hunch that the die will land 6; so that is the bet she takes. The die lands 6. Your friend wins.

In a sense, this is rather unfair. After all, you took the more sensible bet. You took the bet that you *ought* to have taken. Your friend took the less sensible bet. But it was your friend who won, while you did not. On the other hand, calling this 'unfair' is a touch histrionic. Life is full of chance events, and it is simply the *nature* of chance events that, sometimes, unlikely things happen.

The unfortunate gamble that I just asked you to imagine is a good example of *chance* at work. In this book, I will be attempting to give a general account of what chance is. Before beginning on that task, it will pay to make explicit what I take to be the distinctive features of chance.

2 The hallmarks of chance

Chances are something like 'physical probabilities'

Chances are *chancy*. They are intimately associated with concepts such as likelihood, probability, propensity, and possibility. But not every probability is a chance. What makes chance somewhat different from other probability-like phenomena is that chance is in some sense *physical*.

In emphasising that chances are physical, I mean to highlight two things. The first is that chances are objective. Chances do not depend upon what people believe. At a given time, the chance of a given type of event is the same, no matter who evaluates the chance. Chances are not human creations and in this way they are unlike poems, fictional characters, or lounge suites.

The second point to be highlighted is that chance is the sort of thing that is properly studied by the natural sciences. Physicists, chemists, biologists, and others are all in a good position to give authoritative advice on the chances of various events. (That said, it is often possible to form very well-informed opinions about chances, without any special expertise, special equipment, or lengthy investigation. In particular, most of us are able to make very reliable and rapid inferences from the design of gambling devices about the chance of getting particular outcomes with those devices. Here are some chances that we seem well acquainted with. The chance of: drawing an ace from a well-shuffled pack of cards; a roulette wheel landing on black; getting heads with a toss of a well-made coin; rolling a 6 on a die; etc. In all these cases, there are readily observed physical features of the gambling apparatus which – if the device is used correctly – generate particular chances for the possible outcomes.)

If this all seems too obvious to be worth mentioning, that's all to the good. But just to make sure the point has been hammered home, return to the unlucky gamble in which you bet on the die landing anything but 6. Compare the *chance* that the die would land 1–5 with your and your friend's various *degrees of confidence* that the die would land 1–5. The chance was one and the same for both you and your friend. The chance was the sort of matter on which an applied physicist might have given us authoritative advice, if he or she had the opportunity to make various measurements of the die to check its shape, distribution of mass, and so forth. Neither of these things could be said for your degrees of confidence that the die would land 1–5. First, your degrees of confidence in this outcome were manifestly not the same: you invested most of your confidence in the die landing 1–5. Your friend invested most of hers in it landing 6. So, correspondingly, she must have been less confident than you that it would land 1–5. Second, if we wished to make a more accurate determination of your degrees of confidence, we would not need to consult a physicist. We would not need to study the die. If anything, we would need to examine our respective brains, and we would need the help of a psychologist.

As will be discussed below, degrees of confidence appear to be a species of probability also; but they are very different from chances, in that they are not objective properties of mind-independent processes or events. Rather, they are properties of our minds.

In addition to these ideas about the objectivity of chance and the way we find out about chance, there are two further connotations of the idea that chance is a physical probability. I am less confident that these are strictly part of our concept of chance, so I do not wish to include them as 'hallmarks' proper. Rather I note them here as attractive ideas about chance which might end up being more or less integral to the concept.

The first of these connotations is the thought that chances are fixed by the intrinsic properties of a physical system. Roughly, the idea here is that any two systems that are duplicates – alike in all their intrinsic properties – should also be alike in the sorts of chances that they manifest. Two coins that are physically alike should have the same chance of falling heads. Two atoms of a radioactive element that have the same constitution should have the same chance of decaying. And so on. (There are some difficulties with formulating this requirement precisely, since two qualitatively identical coins, tossed in very different ways, may manifest different chances of landing heads. What we want, in cases like this, is to include the person tossing the coin in the physical system. But then there is a danger that, in order to successfully anticipate all the potential influences from 'outside', we have to treat the *entire universe* as the only relevant system. That would seems to strip this idea of much of its interest.)

The second connotation is the idea that chances can feature in physical – maybe even causal – explanations. For instance: Why did that man get lung cancer? *Because his smoking increased his chances of getting cancer.* Why did that substance emit radiation? *Because each atom of that isotope has a high chance of decaying every few minutes.* These look like good explanations which cite chance.

This is a somewhat controversial idea about chance, because there are quite different views about what explanation involves and what causation involves. Some of these views would suggest that the explanations I have gestured at above are strictly misleading, and that chances are never directly involved in causal explanation.

We ought to believe in accordance with the chances

Your friend, who bet on 6, appeared to be behaving irrationally. Assuming she shared your concern to win the $1 prize, and had no other competing concerns, she chose the suboptimal bet.

What is the *source* of your friend's irrationality? On the information given, we cannot answer this question. There are many different ways in which someone can fail to be rational. Moreover, it is possible that your friend – despite appearances – was in fact rational, because she had access to very different information about the die. (Perhaps a very reliable source had told her that, despite appearing fair, the die is very heavily biased towards 6.)

There is one particular way your friend might have been irrational, however, that is important in understanding what chance is. Suppose your friend agreed with you that the die was fair. So she believed that the die had a $1/6$ chance of landing 6. Despite this, her confidence in the proposition 'The die will land 6 on the next roll' was much higher than her confidence that the die would not land 6. In this scenario, your friend has an apparent mismatch between her *belief about the chance* of getting a 6 and her *degree of confidence* that the die will land 6.

We cannot conclusively say whether your friend should have had a different belief about the chance or if she should have had a different degree of confidence. That depends upon the sort of evidence which she has available to her. But we can still conclude that your friend is being irrational because her degree of confidence does not align with her belief about the chances. This is the way chances constrain rational beliefs in general: other things being equal, we ought to apportion our degrees of confidence so as to correspond with what we believe about the chances.

Believing in accordance with the chances is no guarantee of success

You did the rational thing. You bet on 1–5, because that had a much greater chance of paying off than the alternative. But, sadly, you were unlucky, and lost. This is entirely normal for chance events. Even if you adopt sensible beliefs about the chances, that does not mean that all your bets will succeed. Similarly with other aspects of life that are not explicit betting games: if chance is involved, your best-laid plans may still go awry. (The forecast says the chance of rain is low. You leave your umbrella at home . . .)

This gives rise to something of a puzzle, because this feature of chance seems to be in tension with the previous one. *Why* should we believe in accordance with the chances if it does not *guarantee* that we will be better off? Since the most likely outcome is not guaranteed to occur, might we not receive evidence that something particularly unlikely may occur?

There are two replies that are likely to occur to you:

1. *In the long run*, we expect to be more successful if we believe in accordance with the chances than if we adopt any other strategy for forming beliefs. So in some way that is derivative from this claim about the long run, we can draw a conclusion about what is best to believe now, in the single instance.
2. It may be that believing in accordance with the chances is no guarantee, but it is also the case that there is *no better strategy available*. Beggars can't be choosers.

There is some truth in both of these replies, but it will be a task of later chapters to explicate them more clearly, and to assess whether they resolve the puzzle.

Chances of propositions and chances of events

Some philosophers speak of *events* as having chances, others speak of *propositions* having chances. I prefer, in general, to ascribe chances to propositions, since a proposition can refer to an event that does not actually happen, whereas if we speak of events having chances, we need always to qualify that not only actual, but merely possible events can have chances.

It does not seem likely that anything of great import turns on this distinction. There is generally a straightforward translation from a claim that an event of type E has chance x to the claim that ascribes chance x to a *proposition*, which states that an event of type E will occur.

Something interesting does turn on whether we ascribe chances to *propositions* or to *sentences*, however. The same proposition can be expressed by different sentences, and it has been a matter of ongoing interest to philosophers that two sentences which express the same proposition can have very different cognitive significance. For example:

(1) Jack the Ripper killed at least five women in 1888

is a relatively uninformative sentence. But if Jack the Ripper is in fact the name of a man otherwise known as Tom Smith, then the following sentence is also true:

(2) Tom Smith killed at least five women in 1888.

Arguably, these two sentences express the same proposition. The names 'Jack the Ripper' and 'Tom Smith' both denote the same man and the same property is ascribed in both sentences. On commonly held views about propositions, these facts suffice to determine that the propositions are the same.

But, clearly, although they might express the same proposition, the cognitive significance of these sentences is very different. No well-informed person would be surprised to hear the first, but the second would be very big news.

How does this bear on matters of chance? Take another example (drawn from Hawthorne and Lasonen-Aarnio 2009), in which we create an artificial name for a lottery ticket, in a lottery that is yet to be drawn. We declare: 'The ticket that will win this lottery is named "Lucky"'. With this name to hand, I can now confidently assert that:

(3) Lucky will win the lottery.

This sentence is guaranteed to be true, given the special way I created the name 'Lucky'. So I should have the highest possible degree of confidence that this sentence is true. If there are chances that attach to *sentences*, this sentence has a chance equal to one.

Now consider, for each of the 10,000 tickets in the lottery, sentences of the form:

(*) Ticket number *N* will win the lottery.

Assuming the lottery is fair, presumably each of these sentences has a chance of 1 in 10,000. But one of these sentences expresses the very same proposition as (3). So we have two sentences that express the same proposition, but they ascribe different chances. This gives rise to some awkward questions. Believing that a proposition is true while also believing that it is false is to be caught in a contradiction. Similarly, it seems contradictory to ascribe to a proposition two different chances. But if chances attach to sentences rather than propositions, it seems that we will have to ascribe different chances to the same proposition.

If, on the other hand, we ascribe chances to propositions rather than to sentences, we need to choose what chance we ascribe to the proposition expressed by (3). Presumably, this proposition cannot have a chance of one, simply because we have created a name like 'Lucky'. Otherwise, all chance claims would risk being trivialised as soon as we had created a name of this sort. So the chance of (3) must be only 1 in 10,000, even though we are absolutely certain that the sentence is true! Now we have a problem formulating our requirement that chances should accord with our degrees of belief. For in this case, our degree of belief in the *sentence* (3) should *not* match the chance of the corresponding *proposition*.

I won't go further here into precisely how this should be handled. I will simply work on the idea that our degrees of confidence should – other things being equal – match the sorts of chances that attach to propositions. The complexities related to our degree of confidence in sentences that involve names like 'Lucky' and 'Jack the Ripper' I leave for another occasion.

Chances change over time

A less obvious feature of chance is that the chances of a proposition can change over time. To bring this out, consider the chance that the die lands 6. Before the toss, the chance of this was something like ⅙. Accordingly, you were not certain that the die would land 6. But now that it has landed 6, your degree of belief has changed: you are certain that it landed 6. Is this because you now believe contrary to the chances? Surely not: rather, it is because the chance has *changed*, from ⅙ to 1.

More generally, there seems to be a pattern in which propositions about future events can have chances that are greater than zero and less than one – what I'll call non-trivial chances – but past events frequently have only trivial chances: zero or one.

That is not to say that all propositions about the past have trivial chances. Consider the claim that *Napoleon had, in his entire life, an even number of meals*. It is very hard to be certain of a claim like this. It is hard to envisage how it could ever be settled. Consequently, depending somewhat upon your views about chance, you might think that this proposition – and other propositions about the past – can have a non-trivial chance also.

For readers who are unconvinced that Napoleon's culinary history can be a matter of chance, there are other examples which, though more contrived, are perhaps more intuitively compelling. For instance, suppose that you have a time machine which can take you back to the moment of Napoleon's birth. As you begin your time-travel journey, you toss a coin a long way in the air. By the time the coin lands, both you and the coin will have arrived at that earlier time. The proposition 'This coin lands heads at the moment of Napoleon's birth' might appear to be a matter of non-trivial chance, even though the event occurs in the past.

All of this invites the inquiry: what *makes* the chances change over time? A satisfactory account of chance should give some idea of how or why this occurs.

Changing chances over time

Note that it is still true to say that the chance the die would land 6 *was*, before the die was rolled, $1/6$. So in what sense has the chance changed?

We can best get the distinction needed clear by using some formal notation. I'll write the proposition in question, such as 'Napoleon consumed an even number of meals in his entire life', as P. For the chance that P at a given time t, I'll write: Chance-at-$t(P)$. So the claim that chances can vary over time is simply the claim that it is not always the case that, for all times, t_1 and t_2, Chance-at-$t_1(P)$ = Chance-at-$t_2(P)$.

The claim that the chance was $1/6$, then, is simply the claim that there is a time t_0, earlier than now, such that Chance-at-$t_0(P) = 1/6$. That chance claim is true now – at t_1 – even though Chance-at-$t_1(P) = 1$.

It is true at t_1 that Chance-at-$t_1(P) = 1$.
It is true at t_1 that Chance-at-$t_0(P) = 1/6$.

The potential for confusion arises because there are two times involved: the time at which we evaluate the claim and the time involved in the chance itself.

3 Beliefs and probabilities

Degrees of belief

Before attempting to characterise chance, it is necessary to introduce an important technical notion: the idea of 'degree of belief' (also referred to as 'credence'). In ordinary practice, we simply talk about believing or disbelieving. Sometimes we might admit talk of 'partial' belief. But it is quite unfamiliar to think that we might assign these partial beliefs numerical grades.[1]

Even if we do not endorse describing our conscious thought in these terms, there are a number of convergent reasons to think that we might possess a mental state that is at least importantly similar to belief, yet which comes in numerical degrees.

One reason is that there are scenarios where we have *limited*, but *precise*, information about what is going to happen in the future, and if we are to use that limited information in the best way possible, we will need to 'divide our minds' between different possibilities, in some sense. Consider the mental state which you might have adopted with regard to the gamble introduced at the beginning of this chapter. Suppose you played that gambling game six times, and you were reliably informed that the die would land 6 on only one occasion. If, for each trial, you could adopt only one fixed state of mind with respect to the propositions, 'the die lands 6 on trial 1', '. . . on trial 2', and so on, what would the best such state of mind be? Whatever state of belief that is, we might be able effectively to *define* degree of belief $1/6$ to be the attitude that is best to have in a scenario such as this, where the event in question happens once in six trials.[2]

A second way of getting at the concept of degree of belief is to consider how you would value various possible gambles. This was the sort of strategy I was using in the story of the gamble at the beginning of the chapter. The fact that you preferred the gamble on 1–5 over the gamble on 6 was evidence that you had a higher degree of belief that the die would land 1–5.

In order to develop this strategy so as to determine precise numerical degrees of belief, we would need to vary the relative value of the prizes for the two gambles. If a more valuable prize is attached to a gamble,

1 Psychological evidence confirms that people prefer to use qualitative terms to describe their own psychological states, rather than numerical terms, even when allowed some vagueness for the numerical expressions (Budescu and Wallsten 1995: 297–8), although, interestingly, the same studies suggest that most people prefer to *receive* probabilistic information *from others* in numerical form, rather than in qualitative terms.

Richard Holton (2008: 35–40) makes a helpful philosophical proposal to characterise the ideas of belief and partial belief, as opposed to the concept of credence. Holton suggests that all-out belief involves resolving the unmanageable amount of information involved in our credal state into a single possibility which one takes as a 'live possibility': a basis for deliberation. So one all-out believes that *P* if and only if one takes *P* to be a live possibility, and does not take not-*P* as a live possibility. Partial belief is a relaxation of all-out belief. One partially believes *P* if (and only if) one takes *P* as a live possibility, but also takes not-*P* as a live possibility.

2 F. P. Ramsey used this sort of characterisation when he wrote that 'belief of degree m/n is the sort of belief which leads to the action which would be best if repeated *n* times in *m* of which the proposition is true' (Ramsey 1931: 188). See also Galavotti (2005: 202).

then it becomes more attractive to take it. When the two gambles seem equally desirable, it is possible to translate the ratio of the prize-values into information about relative degrees of belief. For instance – speaking for myself – I value \$5 roughly five times more than I value \$1. And, moreover, if the prize for betting on 6 was \$5 and the prize for betting 1–5 was \$1, then I would be indifferent between the two gambles. This reflects my confidence that it is five times less likely that the die will land 6 than it is that it will land 1–5. Accordingly, my credence that the die will land 1–5 is $5/6$, and my credence that the die will land 6 is $1/6$.

To turn the different prize-values into a degree of belief, you use the formula:

$$\text{Degree of belief(Gamble 1 wins)} = \frac{\text{Value(Prize 2)}}{\text{Value(Prize 1)} + \text{Value(Prize 2)}}$$

This idea has been defended by a number of writers. I am largely following D. H. Mellor (1971: chap. 2) in my presentation of the idea.

Unfortunately, this rather indirect approach to explaining a degree of belief suffers from the problem that it involves psychological assumptions that may be utterly unrealistic. Many perfectly rational people object to gambling on a variety of grounds. Asking them to contemplate what they would think about being offered possible gambles might return the disappointing result that they would reject all such gambles. But simply being inclined to reject such gambling games surely does not suffice to show that these individuals have no credences.[3]

A further problem is that it supposes that our valuations of prizes are mathematically well behaved, so that we can, for instance, meaningfully say that a person values one thing ‘five times more’ than she values something else. It is not clear that this is so.

Without going into the full details of debates in the theory of value, note that while it might seem normal for me to be indifferent between the two gambles on the die when the prizes are \$1 and \$5, it would be surprising

3 Ramsey addressed this point briefly:

> Whenever we go to the station we are betting that a train will really run, and if we had not a sufficient degree of belief in this we should decline the bet and stay at home. The options God gives us are always conditional on our guessing whether a certain proposition is true. (Ramsey 1931: 183)

Though Ramsey is quite correct that our choice of action is almost invariably based upon ‘guesses’ as to what will happen as a result of our choice, this is probably not an entirely adequate reply to the worry. The concern is that our degrees of belief may not be as *tightly linked* to our betting dispositions as is required by the Ramseyan method of identifying credences. So some degree of idealisation remains inevitable.