Chance[1]

1. This talk is based on my *Probability: A Philosophical Introduction*, just published by Routledge in a fetching paperback at a not unreasonable price, and which I recommend if only because it says more about chance than I can fit into a 30 minute MSC talk.

2. In particular, it includes a *third* view of chance, as well as the two I'm going to talk about here. This is what I call the *modal* view, since it takes chances to be a measure of physical *possibility*, which neither of the views I'm now going to discuss can do, for a reason I may come to at the end if I have time.

3. What do I mean by chance? By a *chance* I mean an *objective physical probability*, like the probability of a cigarette smoker getting cancer or of a radioactive atom decaying in the next ten years.

4. Chances provide one of the *three* principal applications of numerical probability, the other two being *subjective* probabilities, or *credences*, and *epistemic* probabilities.

5. *Credences* measure how strongly we believe contingent propositions; while *epistemic* probabilities measure how strongly we *should* believe them given whatever evidence about them is available to us.

6. A complete theory of probability should of course cover all these applications, and if we could take the same view of all of them, I should not need to distinguish them from each other.

7. But we can't, for reasons we can go into in discussion; which is why we may need different theories of each of them and why, in this talk, I am only going to discuss views of chance, not of credence or of epistemic probability.

---

[1] Earlier versions of this talk were given to the Bristol University Philosophy Department's staff seminar on 16 January 2004, to the Australasian Association of Philosophy Annual Conference in Queensland on 5 July 2004 and to the Jowett Society in Oxford on 22 October 2004.

8. The two basic views of chance that I want to discuss are simple and familiar enough. They are that chances are *frequencies*, and that they are *dispositions* of a certain kind, often called *propensities*.

9. The devil, of course, as usual in all serious analytic subjects, lies in the detail, some of which is far from simple, and less familiar than perhaps it should be.

10. Nor indeed are the details of these views universally agreed, which is why there are not two views here but rather two *families* of views, whose members differ in details that raise different problems for them – which is what makes these details worth discussing.

11. Let me explain, starting with the frequency views of chance. On the simplest frequency view, chances are frequencies that are both *actual* and *finite*: the chance of smokers getting cancer is simply the frequency, i.e. the *fraction*, of a finite number of actual smokers who actually get cancer.

12. These fractions automatically satisfy the main rules of numerical probability: for example that all probabilities lie between 0 and 1, and that the probability of any proposition A and its negation, ¬A, must add up to 1.

13. Actual finite frequencies are also obviously objective, and not epistemic, since they are not relative to evidence. The fraction of actual smokers who get cancer is as objective and physical as the fact that any one smoker gets cancer.

14. So far so good for the idea that chances are frequencies. But there are snags, some more serious than others. First, what the chances of smokers getting cancer are depends on which smokers we are talking about: analytic smokers, continental smokers, men or women smokers, and so on.

15. This view therefore makes chances relative to *reference classes*, such as the classes of British smokers, women smokers, and so on. And this is true for *all* frequency views, which is why the relevant frequencies are called *relative* frequencies.

16. The relativity of frequencies to reference classes creates problems for all frequency views. But at least, if the classes are actual and finite, there can be no doubt that the relevant frequencies *exist*.

17. And this at once raises the question why, if actual finite relative frequencies are what chances *are*, their existence should be as controversial as in fact it is? There are, I'm sorry to say, many *subjectivist* philosophers of probability, so-called precisely because they deny that there are any objective chances. (Or objective epistemic probabilities, come to that; but that's another matter.)

18. How can this be, when no one can deny that there are actual finite frequencies, such as the fraction of British smokers alive today who either have cancer or will get it in the future?

19. Well, one reason is that not *all* actual finite frequencies are chances. Take the class of men born in June 1950 or July 1981 who are swimming in the Pacific at this very moment: no one thinks the fraction of *these* men who get cancer deserves to be called a 'chance' of getting cancer.

20. So even on an actual finite frequency view, we need a criterion for when a relative frequency *is* a chance – and it is controversial what that criterion should be.

21. Here, however, that doesn't matter: for whatever the criterion, any actual finite frequencies which satisfy it will certainly exist. So when subjectivist philosophers of probability deny that *any* chances exist, as they do, what they mean by 'chances' must be something other than actual finite frequencies.

22. So to give subjectivism a run for its money, let us call actual finite frequencies *statistical* probabilities, leaving it open whether chances *are* statistical probabilities, or whether they are something else, which may or may not exist even though statistical probabilities certainly do.

23. What then is wrong with the statistical view of chance? One common objection to it is that it makes no sense of *single-case* chances, like *my* chance of getting cancer. The reason is that in a reference class with only one member, relative frequencies can only be 1 or 0: in my case, 1 if I get cancer and 0 if I don't. Yet if there is such a thing as my chance of getting cancer, it must surely be able to have values other than 1 or 0.

24. However, this objection is not quite right. The statistical view *need* not identify my chance of getting cancer with a frequency of cancer in a class that only contains *me*: it

can identify it with the frequency of cancer in a larger class of people who resemble me in all relevant respects.

25. This admittedly leaves the problem of saying what respects are 'relevant'; but that is just the problem of finding a criterion to tell us *which* relative frequencies are chances. And while that problem may be difficult to solve, it's by no means impossible.

26. The *real* objection to the statistical view is that every finite reference class rules out *possible* values of chances. In one-member reference classes chances can only be 1 or 0, in two-member classes 1, 1/2 or 0, in three-member classes 1, 2/3, 1/3 and 0, and so on.

27. So however large a finite reference class is, the statistical view will rule out many apparently possible values of chance. Worse still, as we have just seen, *which* values it can let chances have will depend on how large the reference class is.

28. Yet the *possible* chances of getting cancer which smokers of various kinds – British, French, men, women, etc. – *can* have can surely not depend on how many British, French, etc. smokers there happen to be.

29. To meet this objection to the statistical view, most frequency theorists now take chances to be relative to *infinite* reference classes. But that generates two *new* problems, neither of which is easy to solve.

30. First, in *infinite* classes, there are no frequencies other than 0 or 1. Thus suppose there are infinitely many smokers, past present and future. Then if only a *finite* number N, however large, get cancer, their relative frequency N/∞ will, if defined at all, be zero.

31. While if infinitely many of this infinite class of smokers get cancer, there will be no such thing as their relative frequency, since ∞/∞ is certainly not defined. How then can a frequency view of chance be applied in such cases?

32. The usual answer is by extending the mathematical concept of a *limit*, like the limit 1 of the infinite sequence of fractions : 1/2, 2/3, 3/4, 4/5, etc., as on the handout. Similarly, imagine an infinite sequence of increasingly large finite classes of smokers, and of the frequencies of the members of those classes who get cancer.

33. This sequence of frequencies may *also* have a limit. That is to say, there may be a number between 0 and 1 to which later and later frequencies in the sequence get closer and closer as the sequence gets longer.

34. Thus suppose the frequencies of those who get cancer among 1000, 10,000, 100,000, 1 million, etc. smokers are 0.22, 0.17, 0.21, 0.18, and so on – again as on the handout; and that this sequence does in fact have a limit, say 0.20. Then that limit, on the *limiting* frequency view of chance, *is* the chance of smokers getting cancer.

35. It's not quite as simple as this, but the extra complications don't matter here. What matters here is that most *actual* reference classes are *not* infinite. For even if our universe *does* contain infinitely many past, present and future smokers, all the reference classes that matter to us – of British, or French, or men, or women, smokers who are still alive and on this planet – are *finite*. How can the limiting frequency view apply to the chances of *these* smokers getting cancer?

36. The only way it can do so is by applying not just to *actual* but to *hypothetical* smokers. Thus if by 'smokers' we mean *living* smokers, the chance of smokers getting cancer is not the frequency of cancer in the finite class of *actual* smokers but what the limiting frequency of cancer *would* be if there were infinitely many smokers.

37. This frequency view is therefore called the *hypothetical* limiting frequency view, and since I think it is the best, as well as the most widely held, frequency view, from now on it is what I shall mean by 'the frequency view of chance'.

38. But even this view faces hard questions. The main one is this: what fixes the value of the limiting frequencies that it identifies with chances? When the members of a reference class are *actual*, that question has a simple answer: in the smoking and cancer example, the frequency – or the limiting frequency – of smokers getting cancer is fixed by how many *actual* smokers get cancer.

39. But most of the members of a *hypothetical* class of smokers are *not* actual. They are merely *possible*, and any number of merely possible smokers *can* get cancer, and any number of them can also *not* get cancer. So the limiting frequency of cancer in a hypothetical class of smokers *could* be anything from 0 to 1. What then gives that

limiting frequency a definite value which can be identified with a smoker's chance of getting cancer?

40. The answer to this question lies in the familiar distinction between what *could* happen in a hypothetical situation and what *would* happen in it. How for example would you have gone to London if you had gone there today? There are several ways you *could* go; but there may also be *one* way that you *would* go: by train, say.

41. The question then is this: given all the ways you *could* have gone to London today, what makes it the case that you *would* have gone by train? Similarly, given all the values that the limiting frequency of cancer *could* have in an infinite class of hypothetical smokers, what makes one of them – 20% say – the value that this limiting frequency *would* have?

42. The best answer to these questions is given by the concept of a *disposition*, like fragility. Given all the things a glass *could* do if dropped onto a hard floor, what makes breaking what it *would* do is its fragility. Similarly, what makes going by train the way you *would* go to London is that you are *disposed* to go to London by train.

43. Thus: what fixes what *would* happen in a hypothetical situation is an *actual* disposition of an *actual* thing or person: the actual fragility of an actual glass, or your actual disposition to go to London by train.

44. Similarly with limiting frequencies. The limiting frequency of cancer, which the class of smokers *would* have if it were infinite, is fixed by how disposed *actual* smokers *actually* are to get cancer.

45. Dispositions of this kind, which fix the values of hypothetical limiting frequencies, are what I mean by *propensities*; which brings me to the second view of what chances are.

46. A *propensity* view of chance is one that identifies chances not with limiting frequencies, actual or hypothetical, but with *dispositions* to produce limiting frequencies.

47. The advantage of this view is that there need only be *one* actual smoker for that smoker to have a 20% chance of getting cancer, only one actual coin toss for that toss to have a 50/50 chance of landing heads, and so on. If propensities exist at all, they are properties of actual individuals, just as other dispositions are.

48. However, whether the propensity view *really* differs from the hypothetical limiting frequency view depends on what dispositions are. On one view of dispositions, there is *no* difference. Suppose we think, for example, that for something to be fragile is just for it to break if dropped – or otherwise suddenly stressed – in the right way.

49. Now suppose we have two glasses, *a* and *b*, where *a* is fragile and *b* is not. They will of course differ when they are dropped, since *a* will then break and *b* will not break. But if there is no more to fragility than breaking if dropped, then when *a* and *b* are *not* being dropped they need not differ *at all*, and in particular they need not differ by *a* having a property of fragility which *b* lacks. (See handout.)

50. This *anti-realist* view of dispositions is now less widely held than it used to be. And more to the present point, it is quite useless as a reading of the propensity view that, for example, a smoker's chance *p* of getting cancer is a disposition to make *p* the limiting frequency of cancer in a hypothetical infinite class of smokers.

51. For if all this view says is that *if* there were such a class of smokers, it *would* have a limiting frequency *p* of cancer, then it says no more than the hypothetical limiting frequency view. If a propensity view of chance is to differ from that, it cannot incorporate an anti-realist view of dispositions and hence of propensities.

52. So what do *realists* about dispositions say? What they usually say is that if a glass *a* is fragile and a glass *b* is not, then *a* and *b* must differ, even when they are not being dropped, in some real property, which is the so-called *categorical basis* of this difference in their dispositions. (See handout.)

53. This categorical basis may be an intrinsic property which *a* has and *b* lacks, or a property which *b* has and *a* lacks. It may be a single property, or a combination of single properties, either of the glass itself or of its parts, or even of its surroundings and how it is related to them. For our purposes it doesn't matter which it is: in our example it will doubtless be some complex structure of a glass's molecules that determines whether that glass is fragile.

54. And as for deterministic dispositions, so for propensities. For suppose that *a* and *b* are not glasses but people, and that they have different chances of getting cancer, perhaps because *a* smokes and *b* does not. On the best frequency view of chance, all this means

is that if there *were* infinite classes of smokers and of non-smokers, then the limiting frequencies of cancer in those two classes would differ.

55. But as there *are* no such classes, this difference in *a*'s and *b*'s chances of getting cancer need not entail *any* difference between them. In particular, when *a* and *b* are *not* smoking they could, on our frequency view, be identical twins, atom for atom the same, despite having different chances of getting cancer.

56. Not so, on a realist view of dispositions, if these chances are propensities. For then any difference in *a*'s and *b*'s chances of getting cancer must have some categorical basis. This could be some difference in *a*'s and *b*'s properties, or in those of their parts, as it is with fragility. (See handout.)

57. In fact the bases of *a*'s and *b*'s chances of cancer will be different complexes of metabolic properties, which distinguish smokers like *a* from non-smokers like *b*. And if we call *a*'s basis F and *b*'s basis G, then it is by causing *a* to have the basis F rather than the basis G that smoking causes *a*'s chance of getting cancer to differ from *b*'s.

58. This then is how a propensity view of chance differs from the hypothetical limiting frequency view. How do they compare? I think the propensity view is much the better of the two, for the following reasons.

59. First, as we have just seen, the propensity view makes it easy to see how smoking can *affect* our chances of getting cancer, by affecting the metabolic properties that are the bases of those chances. The frequency view makes this causal link mysterious, since it denies that our chances of getting cancer have *any* basis that smoking could affect.

60. Secondly, even setting causation aside, I find anti-realism about propensities and other dispositions equally incredible. In particular, I can't believe that two people with the very same metabolic properties, doing all the same things in identical surroundings, *could* have different chances of getting cancer; any more than I can believe that two glasses with the very same properties could differ in that one is fragile and the other is not.

61. In both cases it seems to me obvious that these dispositions *strongly supervene* on real properties of things and people; by which I mean that things and people *cannot* differ in

their propensities or other dispositions unless they differ in their properties, or in those of their parts, or in their relations to other things or people.

62. Thirdly, realists about propensities and other dispositions are better able than anti-realists to meet a common objection to the usual definition of dispositions in terms of conditionals. The objection, for example to defining '*x* is fragile' as '*x* would break if dropped', is that a glass can be fragile even if it would *not* break if dropped, because dropping it would make it *cease* to be fragile. (See handout.)

63. I know of no good anti-realist answer to this objection. Realists however can meet it by defining '*x* is fragile' as '*x* is F for some F such that *x* would break if dropped *while it is F*', where F is the categorical basis of *x*'s fragility. This definition, which only a realist can give, shows how a glass can be fragile even if it would not break if dropped, because dropping it would make it cease to be F. (See handout.)

64. The problem for *propensities* is slightly different because the conditionals which define a propensity refer not only to whatever thing or person *has* that propensity but to a class of *hypothetical* things or people. The problem here is how to specify the *members* of this class in a way that will give it the *right* limiting frequency.

65. We can see why this is a problem for an anti-realist view of propensities by recalling that, on this view, two people *a* and *b*, with different chances of getting cancer, need not differ in *any* of their real properties or relations. But how then can we distinguish the hypothetical members of the two reference classes with the different limiting frequencies of cancer which we want to identify with *a*'s and *b*'s different chances of getting cancer?

66. Again, I know of no good *anti*-realist answer to that question, whereas propensity theorists have an easy answer to it. For on their *realist* view, *a*'s and *b*'s different propensities must have different categorical bases, F and G, which can then be used to generate two different infinite classes: a class of hypothetical people who are F, whose limiting frequency of cancer equals *a*'s chance of getting cancer; and a class of people who are G, whose different limiting frequency equals *b*'s chance of getting cancer.

67. These are my main reasons for preferring a realist propensity view of chance to the hypothetical limiting frequency view.

68. Note however that the custom of calling the view of dispositions I have been defending *realist* can be misleading, because it suggests that, on this view, dispositions like fragility are real properties of things. In fact, the view implies that most if not all dispositions are *not* real properties, for the following reason.

69. Realism about dispositions asserts that all dispositions have categorical bases, which *are* real properties of things or people. But dispositions themselves can rarely be *identified* with their bases, because most of them have *different* bases in different kinds of things.

70. The molecular structure that makes *glass* fragile, for example, will differ from the molecular structure that make *china* fragile. So fragility itself cannot be identified with *either* of these structures, which are the only properties postulated by a realist view of fragility. This is why, on that view, while all the categorical bases of fragility are real properties, fragility itself is not. (See handout.)

71. Similarly with propensities. Specific chances, say 20%, of getting cancer may well have different metabolic bases in different people, and they will certainly have different bases in different species. So these chances can also not be identified with any of their bases, which is why, on a propensity view of them, while their categorical bases are real properties of people and animals, the chances themselves are not.

72. Subjectivist philosophers of probability who deny the existence of *any* objective chances *could* therefore be right, even on a propensity view of chance. [Option: cut to §87.]

73. However, although these subjectivists *could* be right, I think they are wrong. By this I do *not* mean that *all* dispositions are real properties: fragility is not, and chances of getting cancer are not.

74. But I think some other dispositions *are* real properties, for the simple reason that I think *all* real properties of things are dispositions. So in particular, the categorical bases of dispositions like fragility, and of propensities like our chances of dying of cancer, are themselves dispositions.

75. The difference is that, unlike dispositions such as fragility, dispositions that *are* real properties *are* identical to their categorical bases; which entails of course that they do *not* have different bases in different things or people.

76. Take, for example, inertial mass, as defined by Newton's laws of motion. To have an inertial mass of M units is to be *disposed* to accelerate, when a net force of F units is applied, at a rate proportional to F/M. An inertial mass is a pure disposition.

77. The inertial mass M of any object must therefore have a categorical basis; but as this basis is the *same* in all things, of whatever kind, this basis can be, and I say is, M itself. In other words, the disposition that is an inertial mass is also a real property of things, as are many other dispositions, like temperatures, pressures, electric charges and the intensities of electromagnetic fields. (See handout.)

78. And similarly with some chances, notably those postulated by modern microphysics, of which the simplest and least contentious are the chances of radioactive decay.

79. Take the chance *p* that an atom of radium-226 will decay into an atom of radon-222 within a time interval *t*. The value of this chance *p* is fixed by a function of *t* which contains a parameter that can be expressed as a so-called *half-life*, i.e. the value of *t* – about 1620 years – for which *p* is 1/2; meaning that any atom of radium-226 has a 50/50 chance of decaying within that period of time.

80. An atom's half-life therefore entails all its chances of decaying in different periods of time, just as its inertial mass entails all its dispositions to accelerate under different applied forces.

81. And just as all these dispositions to accelerate have the same categorical basis in all atoms, so all these chances of decaying into atoms of radon-222 have the same categorical basis in all atoms with those chances.

82. That basis will be some real property of an atom's nucleus, a property which can then be *identified* with the conjunction of all the atom's chances of decaying, just as its inertial mass can be identified with the conjunction of all its dispositions to accelerate.

83. So if inertial mass, which is nothing but a conjunction of dispositions, is a real property of things, then the radioactivity of any atom, as measured by its half-life, which is

nothing but a conjunction of propensities to decay, can also be a real property of things. (See handout.)

84. So if all real properties are dispositions, as I think they are, then I also think that some of them are the indeterministic dispositions that we call chances. In short, I am more of a realist about *some* single-case chances than even a propensity view requires me to be.

85. This ambiguity in the term 'realism' as applied to dispositions can of course cut either way. Some philosophers may find a propensity view of chance *more* acceptable when they realise that it need *not* commit them to accepting chances as ultimate and irreducible features of the world.

86. The claim that there *are* such features is an optional extra, which I am not trying to sell here as part of the propensity view itself. All I am trying to sell here is the claim that the propensity view, with or without this optional *extra*, is a better view of chance than any frequency view.

87. But then, if even a propensity view of chance is compatible with there being no chances, why should subjectivists about chance reject this view, as they do? Well, I dare say one reason is not realising that a propensity view doesn't entail the existence of chances.

88. But there is a better reason, namely that even if a propensity view doesn't entail any real chances, it does entail that there are hypothetical limiting frequencies. These, of course, being merely hypothetical, are not actual and hence – on an actualist view of reality – not real.

89. But that, to many philosophers, only makes them more objectionable, especially as it's hard to see how possible-world semantics can make sense of them. How, for example, could there possibly be an *endless* sequence of tosses of a given coin?

90. And if there couldn't, then there is no possible world in which there is such a sequence, and *a fortiori* no such world closest to ours in which the limiting frequency of heads on tosses of this coin has a particular value, say 1/2.

91. So much the worse, many philosophers will say, for the whole idea of hypothetical limiting frequencies, and therefore for the idea of dispositions to produce them, i.e. for propensities.

92. So there *is* a reason for rejecting not merely the existence but the very idea of objective chances, given that they cannot all be merely actual finite frequencies or even actual limiting frequencies. The reason is that in most if not all cases, postulating chances entails postulating hypothetical limiting frequencies.

93. And this is true not only on the two views of chance I've been discussing, but also on the modal view that I've not discussed, which is that what chances measure are degrees of physical possibility.

94. For although I haven't shown it, this view of chance also entails the existence of hypothetical limiting frequencies. Worse still, the quantitative physical possibilities it postulates are features of single worlds, and as such can also not be explained – or explained away – by a possible-world semantics.

95. So much the worse for possible-world semantics, say I; but that's another story, which I can't tell here.

96. All I can tell you here is what you have to pay for denying the idea of chances. You have to deny that any actual finite frequencies – such as the fractions of smokers and others who actually get cancer, of radium atoms that decay in any given time interval, of people who catch flu in a flu epidemic, and so on – have any serious explanation.

97. But as no one can deny that at least some of these frequencies *do* have serious explanations, I don't see how anyone can seriously deny the idea of chances. That is what is really wrong, not only with possible-world semantics, but with subjectivism about chance.