[7] ———. "Logikk og Sannsynlighet" (in Norwegian), Nordisk Mat. Tidskrift 17 (1969), 71–81.
[8] ———. "Remarks on Logic and Probability", in *Logical Theory and Semantic Analysis*. Ed. Stenlund, Dordrecht: D. Reidel, 1974. Pp. 193–197.
[9] Gaifman, H. "Concerning measures on first-order calculi", *Israel Journal Math.* 2 (1964), 1–18.
[10] Hintikka, J., and Niiniluoto, I. "An Axiomatic Foundation for the Logic of Inductive Generalization". Art. 7 above.
[11] Ionescu Tulcea, A., and Ionescu Tulcea, C. *Topics in the Theory of Lifting*. Ergebniss d. Math. und ihrer Grenzgebiete. Vol. 48. Heidelberg: Springer Verlag, 1969.
[12] Jeffreys, H. *Theory of Probability*. 3d ed. Oxford: Oxford University Press, 1961.
[13] Łos, J., "Remarks on Foundation of Probability", in *Proceedings of the International Congress of Mathematicians*. Stockholm, 1962. Pp. 225–229.
[14] Martin-Löf, P., "Repetitive Structures", in *Proc. Conference on Foundational Questions on Statistical Inference*, Department of Theoretical Statistics, University of Aarhus, ed. P. Bluesild and G. Schon, Barndorf-Nielsen.
[15] Renyi, A. "On a New Axiomatic Theory of Probability," *Acta Math. Acad. Sci. Hungarica* 6 (1955), 285–335.
[16] Scott, D., and Krauss, P. "Assigning probabilities to logical formulas", in *Aspects of Inductive Logic*. Ed. Hintikka and Suppes, Amsterdam: North-Holland, 1966. Pp. 219–264.

13

A Subjectivist's Guide to Objective Chance

BY DAVID LEWIS

INTRODUCTION

We subjectivists conceive of probability as the measure of reasonable partial belief. But we need not make war against other conceptions of probability, declaring that where subjective credence leaves off, there nonsense begins. Along with subjective credence we should believe also in objective chance. The practice and the analysis of science require both concepts. Neither can replace the other. Among the propositions that deserve our credence we find, for instance, the proposition that (as a matter of contingent fact about our world) any tritium atom that now exists has a certain chance of decaying within a year. Why should we subjectivists be less able than other folk to make sense of that?

Carnap (1945) did well to distinguish two concepts of probability, insisting that both were legitimate and useful and that neither was at fault because it was not the other. I do not think Carnap chose quite the right two concepts, however. In place of his "degree of confirmation", I would put *credence* or *degree of belief;* in place of his "relative frequency in the long run", I would put *chance* or *propensity*, understood as making sense in the single case. The division of labor between the two concepts will be little changed by these replacements. Credence is well suited to play the role of Carnap's probability$_1$, and chance to play the role of probability$_2$.

Given two kinds of probability, credence and chance, we can have hybrid probabilities of probabilities. (Not "second order probabilities", which suggests one kind of probability self-applied.) Chance of credence need not detain us. It may be partly a matter of chance what one comes to believe, but what of it? Credence about chance is more important. To the believer in chance, chance is a proper subject to have beliefs about. Propositions about chance will enjoy various degrees of belief, and other propositions will be believed to various degrees conditionally upon them.

As I hope the following questionnaire will show, we have some very firm and definite opinions concerning reasonable credence about chance. These opinions seem to me to afford the best grip we have on the concept of chance. Indeed, I am led to wonder whether anyone *but* a subjectivist is in a position to understand objective chance!

QUESTIONNAIRE

First question. A certain coin is scheduled to be tossed at noon today. You are sure that this chosen coin is fair: it has a 50% chance of falling heads and a 50% chance of falling tails. You have no other relevant information. Consider the proposition that the coin tossed at noon today falls heads. To what degree should you now believe that proposition?

Answer. 50%, of course.

(Two comments. (1) It is abbreviation to speak of the coin as fair. Strictly speaking, what you are sure of is that the entire "chance setup" is fair: coin, tosser, landing surface, air, and surroundings together are such as to make it so that the chance of heads is 50%. (2) Is it reasonable to think of coin-tossing as a genuine chance process, given present-day scientific knowledge? I think so: consider, for instance, that air resistance depends partly on the chance making and breaking of chemical bonds between the coin and the air molecules it encounters. What is less clear is that the toss could be designed so that you could reasonably be sure that the chance of heads is 50% exactly. If you doubt that such a toss could be designed, you may substitute an example involving radioactive decay.)

Next question. As before, except that you have plenty of seemingly relevant evidence tending to lead you to expect that the coin will fall heads. This coin is known to have a displaced center of mass, it has been tossed 100 times before with 86 heads, and many duplicates of it have been tossed thousands of times with about 90% heads. Yet you remain quite sure, despite all this evidence, that the chance of heads this time is 50%. To what degree should you believe the proposition that the coin falls heads this time?

Answer. Still 50%. Such evidence is relevant to the outcome by way of its relevance to the proposition that the chance of heads is 50%, not in any other way. If the evidence somehow fails to diminish your certainty that the coin is fair, then it should have no effect on the distribution of credence about outcomes that accords with that certainty about chance. To the extent that uncertainty about outcomes is based on certainty about their chances, it is a stable, resilient sort of uncertainty—new evidence won't get rid of it. (The term "resiliency" comes from Skyrms [1977]; see also Jeffrey [1965], §12.5.)

Someone might object that you could not reasonably remain sure that the coin was fair, given such evidence as I described and no contrary evidence that I failed to mention. That may be so, but it doesn't matter. Canons of reasonable belief need not be counsels of perfection. A moral code that forbids all robbery may also prescribe that if one nevertheless robs, one should rob only the rich. Likewise it is a sensible question what it is reasonable to believe about outcomes if one is unreasonably stubborn in clinging to one's certainty about chances.

Next question. As before, except that now it is afternoon and you have evidence that became available after the coin was tossed at noon. Maybe you know for certain that it fell heads; maybe some fairly reliable witness has told you that it fell heads; maybe the witness has told you that it fell heads in nine out of ten tosses of which the noon toss was one. You remain as sure as ever that the chance of heads, just before noon, was 50% To what degree should you believe that the coin tossed at noon fell heads?

Answer. Not 50%, but something not far short of 100%. Resiliency has its limits. If evidence bears in a direct enough way on the outcome—a way that may nevertheless fall short of outright implication—then it may bear on your beliefs about outcomes otherwise than by way of your beliefs about the chances of the outcomes. Resiliency under all evidence whatever would be extremely unreasonable. We can only say that degrees of belief about outcomes that are based on certainty about chances are resilient under *admissible* evidence. The previous question gave examples of admissible evidence; this question gave examples of inadmissible evidence.

Last question. You have no inadmissible evidence; if you have any relevant admissible evidence, it already has had its proper effect on your credence about the chance of heads. But this time, suppose you are not sure that the coin is fair. You divide your belief among three alternative hypotheses about the chance of heads, as follows.

You believe to degree 27% that the chance of heads is 50%.
You believe to degree 22% that the chance of heads is 35%.
You believe to degree 51% that the chance of heads is 80%.

Then to what degree should you believe that the coin falls heads?

Answer. $(27\% \times 50\%) + (22\% \times 35\%) + (51\% \times 80\%)$; that is, 62%. Your degree of belief that the coin falls heads, conditionally on any one of the hypotheses about the chance of heads, should equal your unconditional degree of belief if you were sure of that hypothesis. That in turn should equal the chance of heads according to the hypothesis: 50% for the first hypothesis, 35% for the second, and 80% for the third. Given your degrees of belief that the coin falls heads, conditionally on the hypotheses, we need only apply the standard multiplicative and additive principles to obtain our answer.

THE PRINCIPAL PRINCIPLE

I have given undefended answers to my four questions. I hope you found them obviously right, so that you will be willing to take them as evidence for what follows. If not, do please reconsider. If so, splendid—now read on.

It is time to formulate a general principle to capture the intuitions that were forthcoming in our questionnaire. It will resemble familiar principles of direct inference except that (1) it will concern chance, not some sort of actual or hypothetical frequency, and (2) it will incorporate the observation that certainty about chances—or conditionality on propositions about chances—makes for resilient degrees of belief about outcomes. Since this principle seems to me to capture all we know about chance, I call it

The Principal Principle. Let C be any reasonable initial credence function. Let t be any time. Let x be any real number in the unit interval. Let X be the proposition that the chance, at time t, of A's holding equals x. Let E be any proposition compatible with X that is admissible at time t. Then

$$C(A/XE) = x.$$

That will need a good deal of explaining. But first I shall illustrate the principle by applying it to the cases in our questionnaire.

Suppose your present credence function is $C(\text{-}/E)$, the function that comes from some reasonable initial credence function C by conditionalizing on your present total evidence E. Let t be the time of the toss, noon today, and let A be the proposition that the coin tossed today falls heads. Let X be the proposition that the chance at noon (just before the toss) of heads is x. (In our questionnaire, we mostly considered the case that x is 50%.) Suppose that nothing in your total evidence E contradicts X; suppose also that it is not yet noon, and you have no foreknowledge of the outcome, so everything that is included in E is entirely admissible. The conditions of the Principal Principle are met. Therefore $C(A/XE)$ equals x. That is to say that x is your present degree of belief that the coin falls heads, conditionally on the proposition that its chance of falling heads is x. If in addition you are sure that the chance of heads is x—that is, if $C(X/E)$ is one—then it follows also that x is your present unconditional degree of belief that the coin falls heads. More generally, whether or not you are sure about the chance of heads, your unconditional degree of belief that the coin falls heads is given by summing over alternative hypotheses about chance:

$$C(A/E) = \sum_x C(X_x/E)C(A/X_xE) = \sum_x C(X_x/E)x,$$

where X_x, for any value of x, is the proposition that the chance at t of A equals x.

Several parts of the formulation of the Principal Principle call for explanation and comment. Let us take them in turn.

THE INITIAL CREDENCE FUNCTION C

I said: let C be any reasonable initial credence function. By that I meant, in part, that C was to be a probability distribution over (at least) the space whose points are possible worlds and whose regions (sets of worlds) are propositions. C is a nonnegative, normalized, finitely additive measure defined on all propositions.

The corresponding conditional credence function is defined simply as a quotient of unconditional credences:

$$C(A/B) =_{df} C(AB)/C(B).$$

I should like to assume that it makes sense to conditionalize on any but the empty proposition. Therefore I require that C is *regular:* $C(B)$ is zero, and $C(A/B)$ is undefined, only if B is the empty proposition, true at no worlds. You may protest that there are too many alternative possible worlds to permit regularity. But that is so only if we suppose, as I do not, that the values of the function C are restricted to the standard reals. Many propositions must have infinitesimal C-values, and $C(A/B)$ often will be defined as a quotient of infinitesimals, each infinitely close but not equal to zero. (See Bernstein and Wattenberg [1969].) The assumption that C is regular will prove convenient, but it is not justified only as a convenience. Also it is required as a condition of reasonableness: one who started out with an irregular credence function (and who then learned from experience by conditionalizing) would stubbornly refuse to believe some propositions no matter what the evidence in their favor.

In general, C is to be reasonable in the sense that if you started out with it as your initial credence function, and if you always learned from experience by conditionalizing on your total evidence, then no matter what course of experience you might undergo your beliefs would be reasonable for one who had undergone that course of experience. I do not say what distinguishes a reasonable from an unreasonable credence function to arrive at after a given course of experience. We do make the distinction, even if we cannot analyze it; and therefore I may appeal to it in saying what it means to require that C be a reasonable initial credence function.

I have assumed that the method of conditionalizing is *one* reasonable way to learn from experience, given the right initial credence function. I have not assumed something more controversial: that it is the *only* reasonable way. The latter view may also be right (the cases where it seems wrong to conditionalize may all be cases where one departure from ideal rationality is needed to compensate for another) but I shall not need it here.

(I said that C was to be a probability distribution over *at least* the space of worlds; the reason for that qualification is that sometimes one's credence might be divided between different possibilities within a single world. That is the case for someone who is sure what sort of world he lives in, but not at all sure who and when and where in the world he is. In a fully general treatment of credence it would be well to replace the worlds by something like the "centered worlds" of Quine [1969], and the propositions by something corresponding to properties. But I shall ignore these complications here.)

THE REAL NUMBER x

I said: let x be any real number in the unit interval. I must emphasize that "x" is a quantified variable; it is not a schematic letter that may freely be replaced by terms that designate real numbers in the unit interval. For fixed A and t, "the chance, at t, of A's holding" is such a term; suppose we put it in for the variable x. It might seem that for suitable C and E we have the following: if X is the proposition that the chance, at t, of A's holding equals the chance, at t, of A's holding—in other words, if X is the necessary proposition—then

$$C(A/XE) = \text{the chance, at } t \text{, of } A\text{'s holding.}$$

But that is absurd. It means that if E is your present total evidence and $C(\text{-}/E)$ is your present credence function, then if the coin is in fact fair—whether or not you think it is!—then your degree of belief that it falls heads is 50%. Fortunately, that absurdity is not an instance of the Principal Principle. The term "the chance, at t, of A's holding" is a nonrigid designator; chance being a matter of contingent fact, it designates different numbers at different worlds. The context "the proposition that . . .", within which the variable "x" occurs, is intensional. Universal instantiation into an intensional context with a nonrigid term is a fallacy. It is the fallacy that takes you, for instance, from the true premise "For any number x, the proposition that x is nine is noncontingent" to the false conclusion "The proposition that the number of planets is nine is noncontingent". See Jeffrey (1970) for discussion of this point in connection with a relative of the Principal Principle.

I should note that the values of "x" are not restricted to the standard reals in the unit interval. The Principal Principle may be applied as follows: you are sure that some spinner is fair, hence that it has infinitesimal chance of coming to rest at any particular point; therefore (if your total evidence is admissible) you should believe only to an infinitesimal degree that it will come to rest at any particular point.

THE PROPOSITION X

I said: let X be the proposition that the chance, at time t, of A's holding equals x. I emphasize that I am speaking of objective, single-case chance—not credence, not frequency. Like it or not, we have this concept. We think that a coin about to be tossed has a certain chance of falling heads, or that a radioactive atom has a certain chance of decaying within the year, quite regardless of what anyone may believe about it and quite regardless of whether there are any other similar coins or atoms. As philosophers we may well find the concept of objective chance troublesome, but that is no excuse to deny its existence, its legitimacy, or its indispensability. If we can't understand it, so much the worse for us.

Chance and credence are distinct, but I don't say they are unrelated. What is the Principal Principle but a statement of their relation? Neither do I say that chance and frequency are unrelated, but they are distinct. Suppose we have many coin tosses with the same chance of heads (not zero or one) in each case. Then there is some chance of getting any frequency of heads whatever; and hence some chance that the frequency and the uniform single-case chance of heads may differ, which could not be so if these were one and the same thing. Indeed the chance of difference may be infinitesimal if there are infinitely many tosses, but that is still not zero. Nor do hypothetical frequencies fare any better. There is no such thing as *the* infinite sequence of outcomes, or *the* limiting frequency of heads, that *would* eventuate if some particular coin toss were somehow repeated forever. Rather there are countless sequences, and countless frequencies, that *might* eventuate and would have some chance (perhaps infinitesimal) of eventuating. (See Jeffrey [1977], Skyrms [1977], and the discussion of "might" counterfactuals in Lewis [1973].)

Chance is not the same thing as credence or frequency; this is not yet to deny that there might be some roundabout way to analyze chance in terms of credence or frequency. I would only ask that no such analysis be accepted unless it is compatible with the Principal Principle. We shall consider how this requirement bears on the prospects for an analysis of chance, but without settling the question of whether such an analysis is possible.

I think of chance as attaching in the first instance to propositions: the chance of an event, an outcome, etc., is the chance of truth of the proposition that holds at just those worlds where that event, outcome, or whatnot occurs. (Here I ignore the special usage of "event" to simply mean "proposition".) I have foremost in mind the chances of truth of propositions about localized matters of particular fact—a certain toss of a coin, the fate of a certain tritium atom on a certain day—but I do not say that those are the only propositions to which chance applies. Not only does it make sense to speak of the chance that a coin will fall heads on a particular occasion; equally it makes sense to speak of the chance of getting exactly seven heads in a particular sequence of eleven tosses. It is only caution, not any definite reason to think otherwise, that stops me from assuming that chance of truth applies to any proposition whatever. I shall assume, however, that the broad class of propositions to which chance of truth applies is closed under the Boolean operations of conjunction (intersection), disjunction (union), and negation (complementation).

We ordinarily think of chance as time-dependent, and I have made that dependence explicit. Suppose you enter a labyrinth at 11:00 A.M., planning to choose your turn whenever you come to a branch point by tossing a coin. When you enter at 11:00, you may have a 42% chance of reaching the center by noon. But in the first half hour you may stray into a region from which it is hard to reach the center, so that by 11:30 your chance of reaching the center by noon has fallen to 26%. But then you turn lucky; by 11:45 you are not far from the center and your chance of reaching it by noon is 78%. At 11:49 you reach the center; then and forevermore your chance of reaching it by noon is 100%.

Sometimes, to be sure, we omit reference to a time. I do not think this means that we have some timeless notion of chance. Rather, we have other ways to fix the time than by specifying it explicitly. In the case of the labyrinth we might well say (before, after, or during your exploration) that your chance of reaching the center by noon is 42%. The understood time of reference is the time when your exploration begins. Likewise we might speak simply of the chance of a certain atom's decaying within a certain year, meaning the chance at the beginning of that year. In general, if A is the proposition that something or other takes place within a certain interval beginning at time t, then we may take a special interest in what I shall call the *endpoint chance* of A's holding: the chance at t, the beginning of the interval in question. If we speak simply of the chance of A's holding, not mentioning a time, it is this endpoint chance—the chance at t of A's holding—that we are likely to mean.

Chance also is world-dependent. Your chance at 11:00 of reaching the center of the labyrinth by noon depends on all sorts of contingent features of the world: the structure of the labyrinth and the speed with which you can walk through it, for instance. Your chance at 11:30 of reaching the center by noon depends on these things, and also on where in the labyrinth you then are. Since these things vary from world to world, so does your chance (at either time) of reaching the center by noon. Your chance at noon of reaching the center by noon is one at the worlds where you have reached the center; zero at all others, including those worlds where you do not explore the labyrinth at all, perhaps because you or it do not exist. (Here I am speaking loosely, as if I believed that you and the labyrinth could inhabit several worlds at once. See Lewis [1968] for the needed correction.)

We have decided this much about chance, at least: it is a function of three arguments. To a proposition, a time, and a world it assigns a real number. Fixing the proposition A, the time t, and the number x, we have our proposition X; it is the proposition that holds at all and only those worlds w such that this function assigns to A, t, and w the value x. This is the proposition that the chance, at t, of A's holding is x.

THE ADMISSIBLE PROPOSITION E

I said: let E be any proposition that is admissible at time t. Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes. Once the chances are given outright, conditionally or unconditionally, evidence bearing on them no longer matters. (Once it is settled that the suspect fired the gun, the discovery of his fingerprints on the trigger adds nothing to the case against him.) The power of the Principal Principle depends entirely on how much is admissible. If nothing is admissible it is vacuous. If everything is admissible it is inconsistent. Our questionnaire suggested that a great deal is admissible, but we saw examples also of inadmissible information. I have no definition of admissibility to offer, but must be content to suggest sufficient (or almost sufficient) conditions for admissibility. I suggest that two different sorts of information are generally admissible.

The first sort is historical information. If a proposition is entirely about matters of particular fact at times no later than t, then as a rule that proposition is admissible at t. Admissible information just before the toss of a coin, for example, includes the outcomes of all previous tosses of that coin and others like it. It also includes every detail—no matter how hard it might be to discover—of the structure of the coin, the tosser, other parts of the setup, and even anything nearby that might somehow intervene. It also includes a great deal of other information that is completely irrelevant to the outcome of the toss.

A proposition is *about* a subject matter—about history up to a certain time, for instance—if and only if that proposition holds at both or neither of any two worlds that match perfectly with respect to that subject matter. (Or we can go the other way: two worlds match perfectly with respect to a subject matter if and only if every proposition about that subject matter holds at both or neither.) If our world and another are alike point for point, atom for atom, field for field, even spirit for spirit (if such there be) throughout the past and up until noon today, then any proposition that distinguishes the two cannot be entirely about the respects in which there is no difference. It cannot be entirely about what goes on no later than

noon today. That is so even if its linguistic expression makes no overt mention of later times; we must beware lest information about the future is hidden in the predicates, as in "Fred was mortally wounded at 11:58". I doubt that any linguistic test of aboutness will work without circular restrictions on the language used. Hence it seems best to take either "about" or "perfect match with respect to" as a primitive.

Time-dependent chance and time-dependent admissibility go together. Suppose the proposition A is about matters of particular fact at some moment or interval t_A, and suppose we are concerned with chance at time t. If t is later than t_A, then A is admissible at t. The Principal Principle applies with A for E. If X is the proposition that the chance at t of A equals x, and if A and X are compatible, then

$$1 = C(A/XA) = x.$$

Put contrapositively, this means that if the chance at t of A, according to X, is anything but one, then A and X are incompatible. A implies that the chance at t of A, unless undefined, equals one. What's past is no longer chancy. The past, unlike the future, has no chance of being any other way than the way it actually is. This temporal asymmetry of chance falls into place as part of our conception of the past as "fixed" and the future as "open"—whatever that may mean. The asymmetry of fixity and of chance may be pictured by a tree. The single trunk is the one possible past that has any present chance of being actual. The many branches are the many possible futures that have some present chance of being actual. I shall not try to say here what features of the world justify our discriminatory attitude toward past and future possibilities,

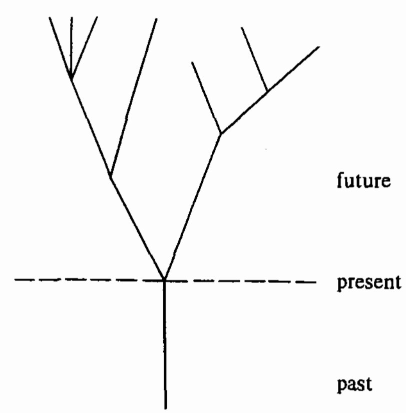

reflected for instance in the judgment that historical information is admissible and similar information about the future is not. But I think they are contingent features, subject to exception and absent altogether from some possible worlds.

That possibility calls into question my thesis that historical information is invariably admissible. What if the commonplace de facto asymmetries between past and future break down? If the past lies far in the future, as we are far to the west of ourselves, then it cannot simply be that propositions about the past are admissible and propositions about the future are not. And if the past contains seers with foreknowledge of what chance will bring, or time travelers who have witnessed the outcome of coin tosses to come, then patches of the past are enough tainted with futurity so that historical information about them may well seem inadmissible. That is why I qualified my claim that historical information is admissible, saying only that it is so "as a rule". Perhaps it is fair to ignore this problem in building a case that the Principal Principle captures our common opinions about chance, since those opinions may rest on a naive faith that past and future cannot possibly get mixed up. Any serious physicist, if he remains at least open-minded both about the shape of the cosmos and about the existence of chance processes, ought to do better. But I shall not; I shall carry on as if historical information is admissible without exception.

Besides historical information, there is at least one other sort of admissible information: hypothetical information about chance itself. Let us return briefly to our questionnaire and add one further supposition to each case. Suppose you have various opinions about what the chance of heads would be under various hypotheses about the detailed nature and history of the chance setup under consideration. Suppose further that you have similar hypothetical opinions about other chance setups, past, present, and future. (Assume that these opinions are consistent with your admissible historical information and your opinions about chance in the present case.) It seems quite clear to me—and I hope it does to you also—that these added opinions do not change anything. The correct answers to the questionnaire are just as before. The added opinions do not bear in any overly direct way on the future outcomes of chance processes. Therefore they are admissible.

We must take care, though. Some propositions about future chances do reveal inadmissible information about future history, and these are inadmissible. Recall the case of the labyrinth: you enter at 11:00, choosing your turns by chance, and hope to reach the center by noon. Your subsequent chance of success depends on the point

you have reached. The proposition that at 11:30 your chance of success has fallen to 26% is not admissible information at 11:00; it is a giveaway about your bad luck in the first half hour. What is admissible at 11:00 is a conditional version: if you were to reach a certain point at 11:30, your chance of success would then be 26%. But even some conditionals are tainted: for instance, any conditional that could yield inadmissible information about future chances by *modus ponens* from admissible historical propositions. Consider also the truth-functional conditional that if history up to 11:30 follows a certain course, then you will have a 98% chance of becoming a monkey's uncle before the year is out. This conditional closely resembles the denial of its antecedent, and is inadmissible at 11:00 for the same reason.

I suggest that conditionals of the following sort, however, are admissible; and indeed admissible at all times. (1) The consequent is a proposition about chance at a certain time. (2) The antecedent is a proposition about history up to that time; and further, it is a complete proposition about history up to that time, so that it either implies or else is incompatible with any other proposition about history up to that time. It fully specifies a segment, up to the given time, of some possible course of history. (3) The conditional is made from its consequent and antecedent not truth-functionally, but rather by means of a strong conditional operation of some sort. This might well be the counterfactual conditional of Lewis (1973); but various rival versions would serve as well, since many differences do not matter for the case at hand. One feature of my treatment will be needed, however: if the antecedent of one of our conditionals holds at a world, then both or neither of the conditional and its consequent hold there.

These admissible conditionals are propositions about how chance depends (or fails to depend) on history. They say nothing, however, about how history chances to go. A set of them is a theory about the way chance works. It may or may not be a complete theory, a consistent theory, a systematic theory, or a credible theory. It might be a miscellany of unrelated propositions about what the chances would be after various fully specified particular courses of events. Or it might be systematic, compressible into generalizations to the effect that after any course of history with property J there would follow a chance distribution with property K. (For instance, it might say that any coin with a certain structure would be fair.) These generalizations are universally quantified conditionals about single-case chance; if lawful, they are probabilistic laws in the sense of Railton (1978). (I shall not consider here what would make them lawful; but see Lewis [1973], §3.3, for a treatment that could cover laws about chance along with other laws.) Systematic theories of chance are the ones we can express in language, think about, and believe to substantial degrees. But a reasonable initial credence function does not reject any possibility out of hand. It assigns some nonzero credence to any consistent theory of chance, no matter how unsystematic and incompressible it is.

Historical propositions are admissible; so are propositions about the dependence of chance on history. Combinations of the two, of course, are also admissible. More generally, we may assume that any Boolean combination of propositions admissible at a time also is admissible at that time. Admissibility consists in keeping out of a forbidden subject matter—how the chance processes turned out—and there is no way to break into a subject matter by making Boolean combinations of propositions that lie outside it.

There may be sorts of admissible propositions besides those I have considered. If so, we shall have no need of them in what follows.

This completes an exposition of the Principal Principle. We turn next to an examination of its consequences. I maintain that they include all that we take ourselves to know about chance.

THE PRINCIPLE REFORMULATED

Given a time t and world w, let us write P_{tw} for the *chance distribution* that obtains at t and w. For any proposition A, $P_{tw}(A)$ is the chance, at time t and world w, of A's holding. (The domain of P_{tw} comprises those propositions for which this chance is defined.)

Let us also write H_{tw} for the *complete history* of world w up to time t: the conjunction of all propositions that hold at w about matters of particular fact no later than t. H_{tw} is the proposition that holds at exactly those worlds that perfectly match w, in matters of particular fact, up to time t.

Let us also write T_w for the *complete theory of chance* for world w: the conjunction of all the conditionals from history to chance, of the sort just considered, that hold at w. Thus T_w is a full specification, for world w, of the way chances at any time depend on history up to that time.

Taking the conjunction $H_{tw}T_w$, we have a proposition that tells us a great deal about the world w. It is nevertheless admissible at time t, being simply a giant conjunction of historical propositions that are admissible at t and conditionals from history to chance that are admissible at any time. Hence the Principal Principle applies:

$$C(A/XH_{tw}T_w) = x$$

when C is a reasonable initial credence function, X is the proposition that the chance at t of A is x, and $H_{tw}T_w$ is compatible with X.

Suppose X holds at w. That is so if and only if x equals $P_{tw}(A)$. Hence we can choose such an X whenever A is in the domain of P_{tw}. $H_{tw}T_w$ and X both hold at w, therefore they are compatible. But further, $H_{tw}T_w$ implies X. The theory T_w and the history H_{tw} together are enough to imply all that is true (and contradict all that is false) at world w about chances at time t. For consider the strong conditional with antecedent H_{tw} and consequent X. This conditional holds at w, since by hypothesis its antecedent and consequent hold there. Hence it is implied by T_w, which is the conjunction of all conditionals of its sort that hold at w; and this conditional and H_{tw} yield X by *modus ponens*. Consequently the conjunction $XH_{tw}T_w$ simplifies to $H_{tw}T_w$. Provided that A is in the domain of P_{tw}, so that we can make a suitable choice of X, we can substitute $P_{tw}(A)$ for x, and $H_{tw}T_w$ for $XH_{tw}T_w$, in our instance of the Principal Principle. Therefore we have

The Principal Principle Reformulated. Let C be any reasonable initial credence function. Then for any time t, world w, and proposition A in the domain of P_{tw}

$$P_{tw}(A) = C(A/H_{tw}T_w).$$

In words: the chance distribution at a time and a world comes from any reasonable initial credence function by conditionalizing on the complete history of the world up to the time, together with the complete theory of chance for the world.

This reformulation enjoys less direct intuitive support than the original formulation, but it will prove easier to use. It will serve as our point of departure in examining further consequences of the Principal Principle.

CHANCE AND THE PROBABILITY CALCULUS

A reasonable initial credence function is, among other things, a probability distribution: a nonnegative, normalized, finitely additive measure. It obeys the laws of mathematical probability theory. There are well-known reasons why that must be so if credence is to rationalize courses of action that would not seem blatantly unreasonable in some circumstances.

Whatever comes by conditionalizing from a probability distribution is itself a probability distribution. Therefore a chance distribution is a probability distribution. For any time t and world w, P_{tw} obeys the laws of mathematical probability theory. These laws carry over from credence to chance via the Principal Principle. We have no need of any independent assumption that chance is a kind of probability.

Observe that although the Principal Principle concerns the relationship between chance and credence, some of its consequences concern chance alone. We have seen two such consequences. (1) The thesis that the past has no present chance of being otherwise than it actually is. (2) The thesis that chance obeys the laws of probability. More such consequences will appear later.

CHANCE AS OBJECTIFIED CREDENCE

Chance is objectified subjective probability in the sense of Jeffrey (1965), §12.7. Jeffrey's construction (omitting his use of sequences of partitions, which is unnecessary if we allow infinitesimal credences) works as follows. Suppose given a partition of logical space: a set of mutually exclusive and jointly exhaustive propositions. Then we can define the *objectification* of a credence function, with respect to this partition, at a certain world, as the probability distribution that comes from the given credence function by conditionalizing on the member of the given partition that holds at the given world. Objectified credence is credence conditional on the truth—not the whole truth, however, but exactly as much of it as can be captured by a member of the partition without further subdivision of logical space. The member of the partition that holds depends on matters of contingent fact, varying from one world to another; it does not depend on what we think (except insofar as our thoughts are relevant matters of fact) and we may well be ignorant or mistaken about it. The same goes for objectified credence.

Now consider one particular way of partitioning. For any time t, consider the partition consisting of the propositions $H_{tw}T_w$ for all worlds w. Call this the *history-theory partition* for time t. A member of this partition is an equivalence class of worlds with respect to the relation of being exactly alike both in respect of matters of particular fact up to time t and in respect of the dependence of chance on history. The Principal Principle tells us that the chance distribution, at any time t and world w, is the objectification of any reasonable credence function, with respect to the history-theory partition for time t, at world w. Chance is credence conditional on the truth—*if* the truth is subject to censorship along the lines of the history-theory partition, and *if* the credence is reasonable.

Any historical proposition admissible at time t, or any admissible conditional from history to chance, or any admissible Boolean combination of propositions of these two kinds—in short, any sort of admissible proposition we have considered—is a disjunction of members of the history-theory partition for t. Its borders follow the lines of the partition, never cutting between two worlds that the partition does not distinguish. Likewise for any proposition about chances at t. Let X be the proposition that the chance at t of A is x, let Y be any member of the history-theory partition for t, and let C be any reasonable initial credence function. Then, according to our reformulation of the Principal Principle, X holds at all worlds in Y if $C(A/Y)$ equals x, and at no worlds in Y otherwise. Therefore X is the disjunction of all members Y of the partition such that $C(A/Y)$ equals x.

We may picture the situation as follows. The partition divides space into countless tiny squares. In each square there is a black region where A holds and a white region where it does not. Now blur the focus, so that divisions within the squares disappear from view. Each square becomes a gray patch in a broad expanse covered with varying shades of gray. Any maximal region of uniform shade is a proposition specifying the chance of A. The darker the shade, the higher is the uniform chance of A at the worlds in the region. The worlds themselves are not gray—they are black or white, worlds where A holds or where it doesn't—but we cannot focus on single worlds, so they all seem to be the shade of gray that covers their region. Admissible propositions, of the sorts we have considered, are regions that may cut across the contours of the shades of gray. The conjunction of one of these admissible propositions and a proposition about the chance of A is a region of uniform shade, but not in general a maximal uniform region. It consists of some, but perhaps not all, the members Y of the partition for which $C(A/Y)$ takes a certain value.

We derived our reformulation of the Principal Principle from the original formulation, but have not given a reverse derivation to show the two formulations equivalent. In fact the reformulation may be weaker, but not in any way that is likely to matter. Let C be a reasonable initial credence function; let X be the proposition that the chance at t of A is x; let E be admissible at t (in one of the ways we have considered) and compatible with X. According to the reformulation, as we have seen, XE is a disjunction of incompatible propositions Y, for each of which $C(A/Y)$ equals x. If there are only finitely many Y's, it would follow that $C(A/XE)$ also equals x. But the implication fails in certain cases with infinitely many Y's (and indeed we would expect the history-theory partition to be infinite) so we cannot quite recover the original formulation in this way. The cases of failure are peculiar, however, so the extra strength of the original formulation in ruling them out seems unimportant.

KINEMATICS OF CHANCE

Chance being a kind of probability, we may define conditional chance in the usual way as a quotient (leaving it undefined if the denominator is zero):

$$P_{tw}(A/B) =_{df} P_{tw}(AB)/P_{tw}(B).$$

To simplify notation, let us fix on a particular world—ours, as it might be—and omit the subscript 'w'; let us fix on some particular reasonable initial credence function C, it doesn't matter which; and let us fix on a sequence of times, in order from earlier to later, to be called $1, 2, 3, \ldots$. (I do not assume they are equally spaced.) For any time t in our sequence, let the proposition I_t be the complete history of our chosen world in the interval from time t to time $t+1$ (including $t+1$ but not t). Thus I_t is the set of worlds that match the chosen world perfectly in matters of particular fact throughout the given interval.

A complete history up to some time may be extended by conjoining complete histories of subsequent intervals. H_2 is H_1I_1, H_3 is $H_1I_1I_2$, and so on. Then by the Principal Principle we have:

$$P_1(A) = C(A/H_1T),$$

$$P_2(A) = C(A/H_2T) = C(A/H_1I_1T) = P_1(A/I_1),$$

$$P_3(A) = C(A/H_3T) = C(A/H_1I_1I_2T) = P_2(A/I_2) = P_1(A/I_1I_2),$$

$$\vdots$$

and in general

$$P_{t+n+1}(A) = P_t(A/I_t \ldots I_{t+n}).$$

In words: a later chance distribution comes from an earlier one by conditionalizing on the complete history of the interval in between.

The evolution of chance is parallel to the evolution of credence for an agent who learns from experience, as he reasonably might, by conditionalizing. In that case a later credence function comes from an earlier one by conditionalizing on the total increment of evidence gained in the interval in between. For the evolution of chance we simply put the world's chance distribution in place of the agent's credence function, and the totality of particular fact about a time in place of the totality of evidence gained at that time.

In the interval from t to $t+1$ there is a certain way that the world will in fact develop: namely, the way given by I_t. And at t, the last moment before the interval begins, there is a certain chance that the world will develop in that way: $P_t(I_t)$, the endpoint chance of I_t. Likewise for a longer interval, say from time 1 to time 18. The world will in fact develop in the way given by $I_1 \ldots I_{17}$, and the endpoint chance of its doing so is $P_1(I_1 \ldots I_{17})$. By definition of conditional chance

$$P_1(I_1 \ldots I_{17}) = P_1(I_1) \cdot P_1(I_2/I_1) \cdot P_1(I_3/I_1I_2) \ldots P_1(I_{17}/I_1 \ldots I_{16}),$$

and by the Principal Principle, applied as above,

$$P_1(I_1 \ldots I_{17}) = P_1(I_1) \cdot P_2(I_2) \cdot P_3(I_3) \ldots P_{17}(I_{17}).$$

In general, if an interval is divided into subintervals, then the endpoint chance of the complete history of the interval is the product of the endpoint chances of the complete histories of the subintervals.

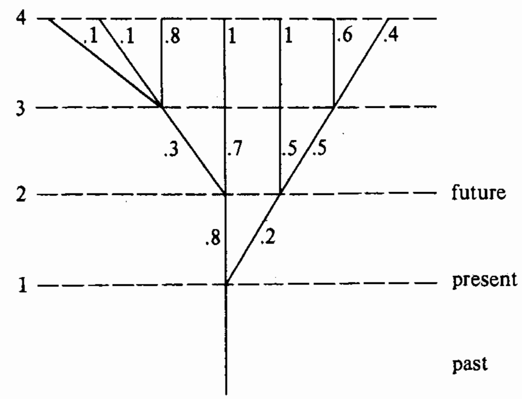

Earlier we drew a tree to represent the temporal asymmetry of chance. Now we can embellish the tree with numbers to represent the kinematics of chance. Take time 1 as the present. Worlds—those of them that are compatible with a certain common past and a certain common theory of chance—lie along paths through the tree. The numbers on each segment give the endpoint chance of the course of history represented by that segment, for any world that passes through that segment. Likewise, for any path consisting of several segments, the product of numbers along the path gives the endpoint chance of the course of history represented by the entire path.

CHANCE OF FREQUENCY

Suppose that there is to be a long sequence of coin tosses under more or less standardized conditions. The first will be in the interval between time 1 and time 2, the second in the interval between 2 and 3, and so on. Our chosen world is such that at time 1 there is no chance, or negligible chance, that the planned sequence of tosses will not take place. And indeed it does take place. The outcomes are given by a sequence of propositions $A_1, A_2, \ldots$. Each A_t states truly whether the toss between t and $t+1$ fell heads or tails. A conjunction $A_1 \ldots A_n$ then gives the history of outcomes for an initial segment of the sequence.

The endpoint chance $P_1(A_1 \ldots A_n)$ of such a sequence of outcomes is given by a product of conditional chances. By definition of conditional chance,

$$P_1(A_1 \ldots A_n) = P_1(A_1) \cdot P_1(A_2/A_1) \cdot P_1(A_3/A_1A_2) \ldots P_1(A_n/A_1 \ldots A_{n-1}).$$

Since we are dealing with propositions that give only incomplete histories of intervals, there is no general guarantee that these factors equal the endpoint chances of the A's. The endpoint chance of A_2, $P_2(A_2)$, is given by $P_1(A_2/I_1)$; this may differ from $P_1(A_2/A_1)$ because the complete history I_1 includes some relevant information that the incomplete history A_1 omits about chance occurrences in the first interval. Likewise for the conditional and endpoint chances pertaining to later intervals.

Even though there is no general guarantee that the endpoint chance of a sequence of outcomes equals the product of the endpoint chances of the individual outcomes, yet it may be so if the world is right. It may be, for instance, that the endpoint chance of A_2 does not depend on those aspects of the history of the first interval that are omitted from A_1—it would be the same regardless. Consider the class of all possible complete histories up to time 2 that are compatible both with the previous history H_1 and with the outcome A_1 of the first toss. These give all the ways the omitted aspects of the first interval might be. For each of these histories, some strong conditional holds at our chosen world that tells what the chance at 2 of A_2 would be if that history were to come about. Suppose all these conditionals have the same consequent: whichever one of the alternative histories were to come about, it would be that X, where X is the proposition that the chance at 2 of A_2 equals x. Then the conditionals taken together tell us that the endpoint chance of A_2 is independent of all aspects of the history of the first interval except the outcome of the first toss.

In that case we can equate the conditional chance $P_1(A_2/A_1)$ and the endpoint chance $P_2(A_2)$. Note that our conditionals are of the sort implied by T, the complete theory of chance for our chosen world. Hence A_1, H_1, and T jointly imply X. It follows that A_1H_1T and XA_1H_1T are the same proposition. It also follows that X holds at our chosen world, and hence that x equals $P_2(A_2)$. Note also that A_1H_1T is admissible at time 2. Now, using the Principal Principle first as reformulated and then in the original formulation, we have

$$P_1(A_2/A_1) = C(A_2/A_1H_1T) = C(A_2/XA_1H_1T) = x = P_2(A_2).$$

If we also have another such battery of conditionals to the effect that the endpoint chance of A_3 is independent of all aspects of the history of the first two intervals except the outcomes A_1 and A_2 of the first two tosses, and another battery for A_4, and so on, then the multiplicative rule for endpoint chances follows:

$$P_1(A_1 \ldots A_n) = P_1(A_1) \cdot P_2(A_2) \cdot P_3(A_3) \ldots P_n(A_n).$$

The conditionals that constitute the independence of endpoint chances mean that the incompleteness of the histories $A_1, A_2, \ldots$ doesn't matter. The missing part wouldn't make any difference.

We might have a stronger form of independence. The endpoint chances might not depend on *any* aspects of history after time 1, not even the outcomes of previous tosses. Then conditionals would hold at our chosen world to the effect that if any complete history up to time 2 which is compatible with H_1 were to come about, it would be that X (where X is again the proposition that the chance at 2 of A_2 equals x). We argue as before, leaving out A_1: T implies the conditionals, H_1 and T jointly imply X, H_1T and XH_1T are the same, X holds, x equals $P_2(A_2)$, H_1T is admissible at 2; so, using the Principal Principle in both formulations, we have

$$P_1(A_2) = C(A_2/H_1T) = C(A_2/XH_1T) = x = P_2(A_2).$$

Our strengthened independence assumption implies the weaker independence assumption of the previous case, wherefore

$$P_1(A_2/A_1) = P_2(A_2) = P_1(A_2).$$

If the later outcomes are likewise independent of history after time 1, then we have a multiplicative rule not only for endpoint chances but also for unconditional chances of outcomes at time 1:

$$P_1(A_1 \ldots A_n) = P_1(A_1) \cdot P_1(A_2) \cdot P_1(A_3) \ldots P_1(A_n).$$

Two conceptions of independence are in play together. One is the familiar probabilistic conception: A_2 is independent of A_1, with respect to the chance distribution P_1, if the conditional chance $P_1(A_2/A_1)$ equals the unconditional chance $P_1(A_2)$; equivalently, if the chance $P_1(A_1A_2)$ of the conjunction equals the product $P_1(A_1) \cdot P_1(A_2)$ of the chances of the conjuncts. The other conception involves batteries of strong conditionals with different antecedents and the same consequent. (I consider this to be *causal* independence, but that's another story.) The conditionals need not have anything to do with probability; for instance, my beard does not depend on my politics since I would have such a beard whether I were Republican, Democrat, Prohibitionist, Libertarian, Socialist Labor, or whatever. But one sort of consequent that can be independent of a range of alternatives, as we have seen, is a consequent about single-case chance. What I have done is to use the Principal Principle to parlay battery-of-conditionals independence into ordinary probabilistic independence.

If the world is right, the situation might be still simpler; and this is the case we hope to achieve in a well-conducted sequence of chance trials. Suppose the history-to-chance conditionals and the previous history of our chosen world give us not only independence (of the stronger sort) but also uniformity of chances: for any toss in our sequence, the endpoint chance of heads on that toss would be h (and the endpoint chance of tails would be $1-h$) no matter which of the possible previous histories compatible with H_1 might have come to pass. Then each of the A_i's has an endpoint chance of h if it specifies an outcome of heads, $1-h$ if it specifies an outcome of tails. By the multiplicative rule for endpoint chances,

$$P_1(A_1 \ldots A_n) = h^{fn} \cdot (1-h)^{(n-fn)}$$

where f is the frequency of heads in the first n tosses according to $A_1 \ldots A_n$.

Now consider any other world that matches our chosen world in its history up to time 1 and in its complete theory of chance, but not in its sequence of outcomes. By the Principal Principle, the chance distribution at time 1 is the same for both worlds. Our assumptions of independence and uniformity apply to both worlds, being built into the shared history and theory. So all goes through for this other world as it did for our chosen world. Our calculation of the chance at time 1 of a sequence of outcomes, as a function of the uniform single-case chance of heads and the length and frequency of heads in the sequence, goes for any sequence, not only for the sequence $A_1, A_2, \ldots$ that comes about at our chosen world.

Let F be the proposition that the frequency of heads in the first n tosses is f. F is a disjunction of propositions each specifying a sequence of n outcomes with frequency f of heads; each disjunct has the same chance at time 1, under our assumptions of independence and uniformity; and the disjuncts are incompatible. Multiplying the number of these propositions by the uniform chance of each, we get the chance of obtaining some or other sequence of outcomes with frequency f of heads:

$$P_1(F) = \frac{n! \cdot h^{fn} \cdot (1-h)^{(n-fn)}}{(fn)! \cdot (n-fn)!}.$$

The rest is well known. For fixed h and n, the right-hand side of the equation peaks for f close to h; the greater is n, the sharper is the peak. If there are many tosses, then the chance is close to one that the frequency of heads is close to the uniform single-case chance of heads. The more tosses, the more stringent we can be about what counts as "close". That much of frequentism is true; and that much is a consequence of the Principal Principle, which relates chance not only to credence but also to frequency.

On the other hand, unless h is zero or one, the right-hand side of the equation is nonzero. So, as already noted, there is always some chance that the frequency and the single-case chance may differ as badly as you please. That objection to frequentist analyses also turns out to be a consequence of the Principal Principle.

EVIDENCE ABOUT CHANCES

To the subjectivist who believes in objective chance, particular or general propositions about chances are nothing special. We believe them to varying degrees. As new evidence arrives, our credence in them should wax and wane in accordance with Bayesian confirmation theory. It is reasonable to believe such a proposition, like any other, to the degree given by a reasonable initial credence function conditionalized on one's present total evidence.

If we look at the matter in closer detail, we find that the calculations of changing reasonable credence involve *likelihoods:* credences of bits of evidence conditionally upon hypotheses. Here the Principal Principle may act as a useful constraint. Sometimes when the hypothesis concerns chance and the bit of evidence concerns the outcome, the reasonable likelihood is fixed, independently of the vagaries of initial credence and previous evidence. What is more, likelihoods are fixed in such a way that observed frequencies tend to confirm hypotheses according to which these frequencies differ not too much from uniform chances.

To illustrate, let us return to our example of the sequence of coin tosses. Think of it as an experiment, designed to provide evidence bearing on various hypotheses about the single-case chances of heads. The sequence begins at time 1 and goes on for at least n tosses. The evidence gained by the end of the experiment is a proposition F to the effect that the frequency of heads in the first n tosses was f. (I assume that we use a mechanical counter that keeps no record of individual tosses. The case in which there is a full record, however, is little different. I also assume, in an unrealistic simplification, that no other evidence whatever arrives during the experiment.) Suppose that at time 1 your credence function is $C(-/E)$, the function that comes from our chosen reasonable initial credence function C by conditionalizing on your total evidence E up to that time. Then if you learn from experience by conditionalizing, your credence function after the experiment is $C(-/FE)$. The impact of your experimental evidence F on your beliefs, about chances or anything else, is given by the difference between these two functions.

Suppose that before the experiment your credence is distributed over a range of alternative hypotheses about the endpoint chances of heads in the experimental tosses. (Your degree of belief that none of these hypotheses is correct may not be zero, but I am supposing it to be negligible and shall accordingly neglect it.) The hypotheses agree that these chances are uniform, and each independent of the previous course of history after time 1; but they disagree about what the uniform chance of heads is. Let us write G_h for the hypothesis that the endpoint chances of heads are uniformly h. Then the credences $C(G_h/E)$, for various h's comprise the *prior distribution* of credence over the hypotheses; the credences $C(G_h/FE)$ comprise the *posterior distribution;* and the credences $C(F/G_hE)$ are the likelihoods. Bayes's Theorem gives the posterior distribution in terms of the prior distribution and the likelihoods:

$$C(G_h/FE) = \frac{C(G_h/E) \cdot C(F/G_hE)}{\sum_h [C(G_h/E) \cdot C(F/G_hE)]}.$$

(Note that "h" is a bound variable of summation in the denominator of the right hand side, but a free variable elsewhere.) In words: to get the posterior distribution, multiply the prior distribution by the likelihood function and renormalize.

In talking only about a single experiment, there is little to say about the prior distribution. That does indeed depend on the vagaries of initial credence and previous evidence.

Not so for the likelihoods. As we saw in the last section, each G_h implies a proposition X_h to the effect that the chance at 1 of F equals x_h, where x_h is given by a certain function of h, n, and f. Hence G_hE and X_hG_hE are the same proposition. Further, G_hE and X are compatible (unless G_hE is itself impossible, in which case G_h might as well be omitted from the range of hypotheses). E is admissible at 1, being about matters of particular fact—your evidence—at times no later than 1. G_h also is admissible at 1. Recall from the last section that what makes such a proposition hold at a world is a certain relationship between that world's complete history up to time 1 and that world's history-to-chance conditionals about the chances that would follow various complete extensions of that history. Hence any member of the history-theory partition for time 1 either implies or contradicts G_h; G_h is therefore a disjunction of conjunctions of admissible historical propositions and admissible history-to-chance conditionals. Finally, we supposed that C is reasonable. So the Principal Principle applies:

$$C(F/G_hE) = C(F/X_hG_hE) = x_h.$$

The likelihoods are the endpoint chances, according to the various hypotheses, of obtaining the frequency of heads that was in fact obtained.

When we carry the calculation through, putting these implied chances for the likelihoods in Bayes's theorem, the results are as we would expect. An observed frequency of f raises the credences of the hypotheses G_h with h close to f at the expense of the others; the more sharply so, the greater is the number of tosses. Unless the prior distribution is irremediably biased, the result after enough tosses is that the lion's share of the posterior credence will go to hypotheses putting the single-case chance of heads close to the observed frequency.

CHANCE AS A GUIDE TO LIFE

It is reasonable to let one's choices be guided in part by one's firm opinions about objective chances or, when firm opinions are lacking, by one's degrees of belief about chances. *Ceteris paribus*, the greater chance you think a lottery ticket has of winning, the more that ticket should be worth to you and the more you should be disposed to chose it over other desirable things. Why so?

There is no great puzzle about why credence should be a guide to life. Roughly speaking, what makes it be so that a certain credence function is *your* credence function is the very fact that you are disposed to act in more or less the ways that it rationalizes. (Better: what makes it be so that a certain reasonable initial credence function and a certain reasonable system of basic intrinsic values are both yours is that you are disposed to act in more or less the ways that are rationalized by the pair of them together, taking into account the modification of credence by conditionalizing on total evidence; and further, you would have been likewise disposed if your life history of experience, and consequent modification of credence, had been different; and further, no other such pair would fit your dispositions more closely.) No wonder your credence function tends to guide your life. If its doing so did not accord to some considerable extent with your dispositions to act, then it would not be your credence function. You would have some other credence function, or none.

If your present degrees of belief are reasonable—or at least if they come from some reasonable initial credence function by conditionalizing on your total evidence—then the Principal Principle applies. Your credences about outcomes conform to your firm beliefs and your partial beliefs about chances. Then the latter guide your life because the former do. The greater chance you think the ticket has of winning, the greater should be your degree of belief that it will win; and the greater is your degree of belief that it will win, the more, *ceteris paribus*, it should be worth to you and the more you should be disposed to choose it over other desirable things.

PROSPECTS FOR AN ANALYSIS OF CHANCE

Consider once more the Principal Principle as reformulated:

$$P_{tw}(A) = C(A/H_{tw}T_w).$$

Or in words: the chance distribution at a time and a world comes from any reasonable initial credence function by conditionalizing on the complete history of the world up to the time, together with the complete theory of chance for the world.

Doubtless it has crossed your mind that this has at least the form of an analysis of chance. But you may well doubt that it is informative as an analysis; that depends on the distance between the analysandum and the concepts employed in the analysans.

Not that it has to be informative *as an analysis* to be informative. I hope I have convinced you that the Principal Principle is indeed informative, being rich in consequences that are central to our ordinary ways of thinking about chance.

There are two different reasons to doubt that the Principal Principle qualifies as an analysis. The first concerns the allusion in the analysans to reasonable initial credence functions. The second concerns the allusion to complete theories of chance. In both cases the challenge is the same: could we possibly get any independent grasp on this concept, otherwise than by way of the concept of chance itself? In both cases my provisional answer is: most likely not, but it would be worth trying. Let us consider the two problems in turn.

It would be natural to think that the Principal Principle tells us nothing at all about chance, but rather tells us something about what makes an initial credence function be a reasonable one. To be reasonable is to conform to objective chances in the way described. Put this strongly, the response is wrong: the Principle has consequences, as we noted, that are about chance and not at all about its relationship to credence. (They would be acceptable, I trust, to a believer in objective single-case chance who rejects the very idea of degree of belief.) It tells us more than nothing about chance. But perhaps it is divisible into two parts: one part that tells us something about chance, another that takes the concept of chance for granted and goes on to lay down a criterion of reasonableness for initial credence.

Is there any hope that we might leave the Principal Principle in abeyance, lay down other criteria of reasonableness that do not mention chance, and get a good enough grip on the concept that way? It's a lot to ask. For note that just as the Principal Principle yields some consequences that are entirely about chance, so also it yields some that are entirely about reasonable initial credence. One such consequence is as follows. There is a large class of propositions such that if Y is any one of the these, and C_1 and C_2 are any two reasonable initial credence functions, then the functions that come from C_1 and C_2 by conditionalizing on Y are exactly the same. (The large class is, of course, the class of members of history-theory partitions for all times.) That severely limits the ways that reasonable initial credence functions may differ, and so shows that criteria adequate to pick them out must be quite strong. What might we try? A reasonable initial credence function ought to (1) obey the laws of mathematical probability theory; (2) avoid dogmatism, at least by never assigning zero credence to possible propositions and perhaps also by never assigning infinitesimal credence to certain kinds of possible propositions; (3) make it possible to learn from experience by having a built-in bias in favor of worlds where the future in some sense resembles the past; and perhaps (4) obey certain carefully restricted principles of indifference, thereby respecting certain symmetries. Of these, criteria (1)–(3) are all very well, but surely not yet strong enough. Given C_1 satisfying (1)–(3), and given any proposition Y that holds at more than one world, it will be possible to distort C_1 very slightly to produce C_2, such that $C_1(-/Y)$ and $C_2(-/Y)$ differ but C_2 also satisfies (1)–(3). It is less clear what (4) might be able to do for us. Mostly that is because (4) is less clear *simpliciter*, in view of the fact that it is not possible to obey too many different restricted principles of indifference at once and it is hard to give good reasons to prefer some over their competitors. It also remains possible, of course, that some criterion of reasonableness along different lines than any I have mentioned would do the trick.

I now turn to our second problem: the concept of a complete theory of chance. In saying what makes a certain proposition be the complete theory of chance for a world (and for any world where it holds), I gave an explanation in terms of chance. Could these same propositions possibly be picked out in some other way, without mentioning chance?

The question turns on an underlying metaphysical issue. A broadly Humean doctrine (something I would very much like to believe if at all possible) holds that all the facts there are about the world are particular facts, or combinations thereof. This need not be taken as a doctrine of analyzability, since some combinations of particular facts cannot be captured in any finite way. It might be better taken as a doctrine of supervenience: if two worlds match perfectly in all matters of particular fact, they match perfectly in all other ways too—in modal properties, laws, causal connections, chances, and so on. It seems that if this broadly Humean doctrine is false, then chances are a likely candidate to be the fatal counterinstance. And if chances are not supervenient on particular fact, then neither are complete theories of chance. For the chances at a world are jointly determined by its complete theory of chance together with propositions about its history, which latter plainly are supervenient on particular fact.

If chances are not supervenient on particular fact, then neither chance itself nor the concept of a complete theory of chance could possibly be analyzed in terms of particular fact, or of anything supervenient thereon. The only hope for an analysis would be to use something in the analysans which is itself not supervenient on particular fact. I cannot say what that something might be.

How might chance, and complete theories of chance, be supervenient on particular fact? Could something like this be right: the complete theory of chance for a world is that one of all possible

complete theories of chance that somehow best fits the global pattern of outcomes and frequencies of outcomes? It could not. For consider any such global pattern, and consider a time long before the pattern is complete. At that time, the pattern surely has some chance of coming about and some chance of not coming about. There is surely some chance of a very different global pattern coming about; one which, according to the proposal under consideration, would make true some different complete theory of chance. But a complete theory of chance is not something that could have some chance of coming about or not coming about. By the Principal Principle,

$$P_{tw}(T_w) = C(T_w/H_{tw}T_w) = 1.$$

If T_w is something that holds in virtue of some global pattern of particular fact that obtains at world w, this pattern must be one that has no chance at any time (at w) of not obtaining. If w is a world where many matters of particular fact are the outcomes of chance processes, then I fail to see what kind of global pattern this could possibly be.

But there is one more alternative. I have spoken as if I took it for granted that different worlds have different history-to-chance conditionals, and hence different complete theories of chance. Perhaps this is not so: perhaps all worlds are exactly alike in the dependence of chance on history. Then the complete theory of chance for every world, and all the conditionals that comprise it, are necessary. They are supervenient on particular fact in the trivial way that what is noncontingent is supervenient on anything—no two worlds differ with respect to it. Chances are still contingent, but only because they depend on contingent historical propositions (information about the details of the coin and tosser, as it might be) and not also because they depend on a contingent theory of chance. Our theory is much simplified if this is true. Admissible information is simply historical information; the history-theory partition at t is simply the partition of alternative complete histories up to t; for any reasonable initial credence function C

$$P_{tw}(A) = C(A/H_{tw}),$$

so that the chance distribution at t and w comes from C by conditionalizing on the complete history of w up to t. Chance is reasonable credence conditional on the whole truth about history up to a time. The broadly Humean doctrine is upheld, so far as chances are concerned: what makes it true at a time and a world that something has a certain chance of happening is something about matters of particular fact at that time and (perhaps) before.

What's the catch? For one thing, we are no longer safely exploring the consequences of the Principal Principle, but rather engaging in speculation. For another, our broadly Humean speculation that history-to-chance conditionals are necessary solves our second problem by making the first one worse. Reasonable initial credence functions are constrained more narrowly than ever. Any two of them, C_1 and C_2, are now required to yield the same function by conditionalizing on the complete history of any world up to any time. Put it this way: according to our broadly Humean speculation (and the Principal Principle) if I were perfectly reasonable and knew all about the course of history up to now (no matter what that course of history actually is, and no matter what time is now) then there would be only one credence function I could have. Any other would be unreasonable.

It is not very easy to believe that the requirements of reason leave so little leeway as that. Neither is it very easy to believe in features of the world that are not supervenient on particular fact. But if I am right, that seems to be the choice. I shall not attempt to decide between the Humean and the anti-Humean variants of my approach to credence and chance. The Principal Principle doesn't.

BIBLIOGRAPHY

Bernstein, Allen R., and Frank Wattenberg. (1969). "Non-Standard Measure Theory", in *Applications of Model Theory of Algebra, Analysis, and Probability*. Ed. W. Luxemburg. New York: Holt, Reinhart & Winston.

Carnap, Rudolf. (1945). "The Two Concepts of Probability", *Philosophy and Phenomenological Research* **5**, 513–532.

Jeffrey, Richard C. (1965). *The Logic of Decision*. New York: McGraw-Hill.

———. (1970). Review of articles by David Miller et al., *Journal of Symbolic Logic*, **35**, 124–127.

———. (1977). "Mises Redux", in *Basic Problems in Methodology and Linguistics: Proceedings of the Fifth International Congress of Logic, Methodology and Philosophy of Science, Part III*. Ed. R. Butts and J. Hintikka. Dordrecht: D. Reidel.

Lewis, David. (1968). "Counterpart Theory and Quantified Modal Logic", *Journal of Philosophy* **65**, 113–126.

———. (1973). *Counterfactuals*. Oxford: Blackwell.

Mellor, D. H. (1971). *The Matter of Chance*. Cambridge: Cambridge University Press.

Quine, W. V. (1969). "Propositional Objects", in *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Railton, Peter. (1978). "A Deductive-Nomological Model of Probabilistic Explanation", *Philosophy of Science* **45**, 202–226.

Skyrms, Brian. (1977). "Resiliency, Propensities, and Casual Necessity", *Journal of Philosophy* **74**, 704–713.

NOTE

I am grateful to several people for valuable discussions of this material, especially John Burgess, Nancy Cartwright, Richard Jeffrey, Peter Railton, and Brian Skyrms. I am also much indebted to Mellor (1971), which presents a view very close to mine; exactly how close, I am not prepared to say.