A Philosopher's Guide to Probability

Alan Hájek

Autobiographical and pedagogical prelude

Once upon a time I was an undergraduate majoring in mathematics and statistics. I attended many lectures on probability theory, and my lecturers taught me many nice theorems involving probability: 'P of this equals P of that', and so on. One day I approached one of them after a lecture and asked him: "What is this 'P' that you keep on writing on the blackboard? *What is probability?*" He looked at me like I needed medication, and he told me to go the philosophy department. In the interests of pedagogy, in retrospect I think that *he* could have benefited from some discussions with philosophers. For when I now teach those same theorems to my students, I hope that I can imbue them with deeper meaning and motivation when I point out what is at stake philosophically.

Anyway, *I* did go to the philosophy department. (Admittedly, my route there was long and circuitous.) There I found a number of philosophers asking the very same question: what is probability? All these years later, it's still one of the main questions that I am working on. I still don't feel that I have a completely satisfactory answer, although I like to think that I've made some progress on it. For starters, I know many things that probability is *not*, namely various highly influential analyses of it that cannot be right—we will look at them shortly. As to promising directions regarding what probability *is*, I will offer my best bets at the end, concluding with some further autobiographical and pedagogical thoughts.

Introduction

Bishop Butler's dictum that "Probability is the very guide of life" is as true today as it was when he wrote it in 1736. It is almost platitudinous to point out the importance of probability in statistics, physics, biology, chemistry, computer science, medicine, law, meteorology, psychology, economics, and so on. Probability is crucial to any discipline that deals with indeterministic processes, any discipline in which evidence has a non-deductive bearing on hypotheses, indeed any discipline in which our ability to predict outcomes is imperfect—which is to say virtually any serious empirical discipline. Probability is also seemingly ubiquitous outside the academy. Probabilistic judgments of the efficacy and side-effects of a pharmaceutical drug determine whether or not it is approved for release to the public. The fate of a defendant on trial for murder hinges on the jurors' opinions about the probabilistic weight of evidence. Geologists calculate the probability that an earthquake of a certain intensity will hit a given city, and engineers accordingly build skyscrapers with specified probabilities of withstanding such earthquakes. Probability undergirds even measurement itself, since the error bounds that accompany measurements are essentially probabilistic confidence intervals. We find probability wherever we find uncertainty—that is, almost everywhere in our lives.

It is surprising, then, that probability is a comparative latecomer on the intellectual scene. To be sure, inchoate ideas about chance date back to antiquity—Epicurus, and later Lucretius, believed that atoms occasionally underwent indeterministic swerves. In the middle ages, Averroes had a notion of 'equipotency' that might be regarded as a precursor to probabilistic notions. But probability theory was not conceived until the 17th century, when the study of gambling games motivated the first serious mathematical study of chance by Pascal and Fermat in the mid-17th century, culminating in the *Port-Royal Logic.* Over the next three centuries, the theory was developed by such authors as Huygens, Bernoulli, Bayes, Laplace, Condorcet, de Moivre, Venn, Johnson, and Keynes. Arguably, the crowning achievement was Kolmogorov's axiomatization in 1933, which put probability on rigorous mathematical footing.

When I asked my professor "What is probability?", there are two ways to understand that question, and thus two kinds of answer that could be given (apart from bemused advice to seek attention from a doctor, or at least a doctor of philosophy). First, the question may mean: *what are the formal features of probability*? That is a mathematical question, to which Kolmogorov's axiomatization is the widely (though not universally) agreed upon answer. I review this answer in the next section, and it was given to me at great length in my undergraduate statistics courses. Second, the question may mean: *what sorts of things are probabilities*—what, that is, is the subject matter of probability theory? This is a philosophical question, and while the mathematical theory of probability certainly bears on it, the answer must come from elsewhere—in my case, from the philosophy department.

The formal theory of probability

Unconditional probability

Kolmogorov begins his classic book (1933, 1950) with what he calls the "elementary theory of probability": the part of the theory that applies when there are only finitely many events in question. Let Ω be a set (the 'universal set'). A *field* on Ω is a set of subsets of Ω that has Ω as a member, and that is closed under complementation (with respect to Ω) and finite union. Let Ω be given, and let F be a field on Ω. Kolmogorov's axioms constrain the possible assignments of numbers, thought of as *probabilities*, to the members of F. Let P be a function from F to the real numbers obeying:

1. (Non-negativity) $P(A) \geq 0$ for all $A \in F$,
2. (Normalization) $P(\Omega) = 1$.
3. (Finite additivity) $P(A \cup B) = P(A) + P(B)$ for all $A, B \in F$ such that $A \cap B = \emptyset$.

Such a triple (Ω, F, P) is called a *probability space*.

Here the arguments of the probability function are sets, often called *events*. (Note that this is a technical sense of the term that may not neatly align with ordinary usage—for example, it is not clear that 'events' in the latter sense have the required closure properties.) Kolmogorov’s probability theory is thus parasitic on set theory. We could instead attach real-valued probabilities to members of a collection S of *sentences* of a language, closed under finite truth-functional combinations, with the following counterpart axiomatization:

I. $P(A) \geq 0$ for $A \in S$.

II. If T is a tautology, then $P(T) = 1$.

III. $P(A \vee B) = P(A) + P(B)$ for all $A \in S$ and $B \in S$ such that A and B are logically incompatible.

Note how these axioms take the notions of 'tautology', 'logical incompatibility' and 'implication' as antecedently understood. To this extent we may regard probability theory, so formulated, as parasitic on deductive logic.

Now let Ω be infinite. A non-empty collection F of subsets of Ω is called a *sigma algebra* (or *sigma field*, or *Borel field*) on Ω iff F is closed under complementation and *countable* union, i.e.

$A_1, A_2, \ldots \in F \Rightarrow A_n \in F$. [NOTE TO EDITOR: I HOPE THESE SYMBOLS SURVIVED ELECTRONIC TRANSMISSION. THE MIDDLE SYMBOL, FOR EXAMPLE, SHOULD BE ‘IMPLIES’.]

Kolmogorov introduces a further ‘infinitary’ axiom.

4. (Continuity) If $E_1, E_2, \ldots$ is a sequence of sets such that $E_i \supseteq E_{i+1}\ \forall i$ and $E_n = \emptyset$.

then $P(E_n) \rightarrow 0$ (where $E_n \in F$ for all n).

.

That is, if $E_1, E_2, \ldots$ is a sequence of non-increasing sets (according to the set-inclusion relation), with empty infinite intersection, then $\lim_{n \rightarrow \infty} P(E_n) = P(\ E_n)$. Now, define a *probability measure* P(–) *on* F as a function from F to [0, 1] satisfying axioms 1-3, as before, and also:

Equivalently, we can replace the conjunction of axioms 3 and 4 with a single axiom:

3'. (Countable additivity) If $\{A_i\}$ is a countable collection of (pairwise) disjoint sets, each $\in$ F, then

$$P(\ A_n) = \ P(A_n).$$

Thanks to the assumption that F is a sigma algebra, we are guaranteed that the probability on the left hand side is defined.

De Finetti (1972 and 1974) marshals a battery of arguments against countable additivity, most of them variations on these:

The infinite lottery: Suppose a positive integer is selected at random—we might think of this as an infinite lottery with each positive integer appearing on exactly one ticket. We would like to reflect this in a uniform distribution over the positive integers (indeed, proponents of the principle of indifference would seem to be committed to it), but if we assume countable additivity this is not possible. For if we assign probability 0 in turn to each number's being picked, then the sum of all these probabilities is again 0; yet the union of all of these events has probability 1 (since it is guaranteed that some number will be picked), and $1 \neq 0$. On the other hand, if we assign some probability $\varepsilon > 0$ to each number being picked, then the sum of these probabilities diverges to ∞, and $1 \neq \infty$. If we drop countable additivity, however, then we may assign 0 to each event and 1 to their union without contradiction.

Biased assignments to denumerable sets: Countable additivity allows one to assign uniform probability 1/n to each member of an n-celled partition (for example, 1/6 to each result of tossing a die). However, it requires one to assign an extremely biased distribution to a denumerable partition of events. Indeed, for any $\varepsilon > 0$, however small, there will be a finite number of events that have a combined probability of at least $1 - \varepsilon$, and thus the lion's share of all the probability.

See Seidenfeld (2001) for further discussion of countable additivity.

It is often thought that the only part of the axiomatization that is not merely conventional stipulation is the third axiom, in either its finite or countable form. (For example, it is tempting to say that it is *purely* conventional to set $P(\Omega) = 1$, rather than $P(\Omega) = 100$, say.) That is too quick. For each of the following involves a substantial mathematical assumption:

(i) Probabilities are defined by *functions* (rather than by one-many or many-many mappings).

(ii) These are functions of *one variable* (unlike primitive conditional probability functions, which are functions of two variables): there is just a single argument for a probability function.

(iii) Such a function is defined on a *field* (rather than a set with different closure conditions).

(iv) Probabilities are to be represented *numerically* (rather than qualitatively, or comparatively).

(v) Their numerical values are *real* numbers (rather than those of some other number system).

(vi) These values are *bounded* (unlike other quantities that are treated measure-theoretically, such as lengths).

(vii) Probability functions attain *maximal* and *minimal* values (thus prohibiting open or half-open ranges, such as (0, 1) or (0, 1]).

For a discussion of rival theories that relax or replace (ii), (iii), (iv) and (vi), see Fine (1973). Complex-valued probabilities are proposed by Feynman and Cox (see Mückenheim et al. 1986); infinitesimal probabilities (of non-standard analysis) by Skyrms (1980) and Lewis (1980) among others; unbounded probabilities by Renyi (1970).

<u>Conditional probability</u>

Kolmogorov also defines the *conditional probability* of *A given B* by the ratio formula:

$$\text{(RATIO)} \qquad P(A|B) = \frac{P(A \quad B)}{P(B)} \qquad (P(B) > 0).$$

Thus, we may say that the probability that the toss of a fair die results in a 6 is 1/6, but the probability that it results in a 6 *given* that it results in an even number, is $\frac{1/6}{1/2}$ = 1/3. In straightforward applications in which the requisite unconditional probabilities are well-defined, and the denominator *P(B)* is greater than 0, this formula seems to be impeccable.

But not all applications are straightforward, and in some these conditions are not met. Consider first the proviso that *P(B)* > 0. As probability textbooks repeatedly drum into their readers, probability zero events need not be impossible, and indeed they can be of real significance. It is curious, then, that some of the same textbooks glide over (RATIO)'s proviso without missing a beat. In fact, interesting cases of conditional probabilities with probability-zero conditions are manifold. Consider an example due to Borel: A point is chosen at random from the surface of the earth (thought of as a perfect sphere); what is the probability that it lies in the Western hemisphere, given that it lies on the equator? 1/2, surely. Yet the probability of the condition is 0, since a uniform probability measure over a sphere must award probabilities to regions in proportion to their area, and the equator has area 0. The ratio formula thus cannot deliver the intuitively correct answer. Obviously there are uncountably many problem cases of this form for the sphere. (For more discussion see Hájek 2003b).

Probability theory and statistics are shot through with cases of non-trivial zero-probability events. Witness the probabilities of continuous random variables taking particular values (such as a normally distributed random variable taking the value 0). Witness the various 'almost sure' results—the strong law of large numbers, the law of the iterated logarithm, the martingale convergence theorem, and so on. They assert that certain convergences take place, not with certainty, but 'almost surely'. A fair coin may land tails forever. But despite this event's having probability 0, various probabilities conditional on it are intuitively well defined—for example, the probability that the coin lands tails forever, *given* that it lands tails forever, is surely 1. More generally, it is surely platitudinous that the probability of any possible outcome, *given itself,* is 1. This is about as fundamental a fact about conditional probability as there could be. The fact that the ratio formula cannot respect this platitude is a major strike against it.

The difficulties that probability zero conditions pose for the ratio formula for conditional probability are well known (which is not to say that they are unimportant). Indeed, Kolmogorov himself was well aware of them, and he offered a more sophisticated account of conditional probability as a random variable conditional on a sigma algebra, appealing to the Radon-Nikodym theorem to guarantee the existence of such a random variable. But my complaints about the ratio analysis are hardly aimed at a straw man, since (RATIO) is by far the most commonly used analysis of conditional probability (especially in philosophical applications of probability). Moreover, the move to Kolmogorov's more sophisticated theory of conditional probability does not lay to rest the problem of zero-probability conditions. In particular, even Kolmogorov's more sophisticated account of conditional probability does not respect the platitude above concerning conditional probability, as evidenced by the existence of so-called *improper* conditional probability random variables. Seidenfeld et al. (2001) show just how extreme and how widespread violations of the platitude can be.

Hájek (2003b) goes on to consider further problems for the ratio formula: cases in which the unconditional probabilities that figure in the ratio are *imprecise* or are *undefined*, and yet the corresponding conditional probabilities are defined. Consider first the problem of imprecise unconditional probabilities: P(A|B) does not have a sharp value when P(B) is imprecise. Example: Presumably, your probability that it rains tomorrow is not a sharp value, such as 0.3. After all, that value is *infinitely* sharp, precise to infinitely many decimal places: 0.30000… Rather, your probability is spread out over a range of values—say, the interval [0.25, 0.35]. Philosophers commonly represent such imprecision with a *set* of probability functions—in this case, the set of all such functions that assign some value in the interval [0.25, 0.35] to it raining tomorrow, and that agree with your opinions in all other respects. (See van Fraassen (1990), which develops proposals by Levi 1980 and Jeffrey 1983, among others.) Nevertheless, you assign various perfectly sharp conditional probabilities, *given* that it rains. For example, the probability that a particular fair coin toss lands heads, given that it rains tomorrow, is a sharp 1/2.

The problem of undefined conditional probabilities is even more widespread: P(A|B) is undefined when either or both of $P(A \cap B)$ and P(B) are undefined. Suppose that P(Sam has just watered the garden) is undefined. Still, the probability that the garden is dry, given that Sam has just watered the garden, is 0. Hájek (2003b) discusses a plethora of such cases, ranging from quantum mechanics to decision theory, from non-measurable sets to probabilistic causation. The ratio formula goes silent where our intuitions cry out what the answers should be. Moreover, these prove to be problematic also for the more sophisticated account, so therein lies no solution either.

I believe that the right response is to turn the tables on Kolmogorov's analysis: rather than regarding unconditional probabilities as fundamental, and later defining conditional probabilities in terms of them, we should regard conditional probability as the fundamental notion. And there are various ways to define primitive conditional probabilities as total functions from $\mathsf{F} \times \mathsf{F}$ to [0, 1] (for a given sigma field F). Popper (1959a) presents a general account of such conditional probabilities. Among other things, any such account codifies the platitude that the conditional probability of anything, given *itself,* is 1. See Hájek

(2003b) for further discussion.

.

Independence

If $P(X \mid Y) = P(X)$, then X and Y are said to be *independent*. Intuitively, the occurrence of one of the events is completely uninformative about the occurrence of the other. Thus, successive spins of a roulette wheel are typically regarded as independent. When $P(X) > 0$ and $P(Y) > 0$, the definition of independence is equivalent to

$$P(X \mid Y) = P(X)$$

and to

$$P(X \cap Y) = P(X)P(Y).$$

The latter formulation is used even when $P(X)$ or $P(Y)$ are 0.

The locution '*X* is independent of Y' is somewhat careless, encouraging one to forget that independence is a 3-place relation that events or sentences bear *to a probability function*. Furthermore, this technical sense of 'independence' should not be identified unreflectively with causal independence, or any other pre-theoretical sense of the word, even though such identifications are often made in practice.

Independence plays a central role in probability theory—indeed, it is that theory's distinctive add-on to the more general measure theory on which it is based. Many of those theorems that my statistics professors taught me, and which I now teach, hinge on it—witness again the various laws of large numbers, for instance. It should come as no surprise that my misgivings about the ratio analysis of conditional probability carry over to the present definitions of independence, which presupposes that analysis.

The next section turns to the so-called *interpretations* of probability: attempts to answer the central *philosophical* question: what sorts of things are probabilities? The term 'interpretation' is misleading here. Various quantities that intuitively have nothing to do with 'probability' obey Kolmogorov's axioms—for example, length, volume, and mass, each suitably normalized—and are thus 'interpretations' of it, but not in the intended sense. Nevertheless, we will silence our scruples and follow common usage in our quick survey of the 'interpretations' of probability. (See Hájek 2003a for a far more detailed survey.)

Interpretations of probability

Classical interpretation

The *classical* interpretation, historically the first, can be found in the works of Pascal, Huygens, Bernoulli, and Leibniz, and it was famously presented by Laplace (1814). Cardano, Galileo, and Fermat also anticipated this interpretation. Suppose that our evidence does not discriminate among the members of some set of possibilities—either because that evidence provides equal support for each of them, or because it has no bearing on them at all. Then the probability of an event is simply the fraction of the total number of possibilities in which the event occurs—this is sometimes called the *principle of indifference*. We may think of this as the *rational epistemic probability* appropriate for an agent in the evidential situation described. This interpretation was inspired by, and typically applied to, games of chance that by their very design create such circumstances—for example, the classical probability of a fair die landing with an even number showing up is 3/6. Probability puzzles typically take this means of calculating probabilities for granted.

Unless more is said, the interpretation yields contradictory results: you have a one-in-a-million chance of winning the lottery; but either you win or you don't, so each of these possibilities has probability ?!. We might look for a "privileged" partition of the possibilities, but we will not always find one. For example, in this case, the million-celled partition corresponding to each of the possible lottery outcomes seems more natural than the win/don't win partition, if only because the former is more fine-grained. But Bertrand's paradoxes (1889) show that a particular problem may have competing, equally natural, partitions. They all turn on alternative parametrizations of a given problem that are non-linearly related to each other. The following example (adapted from van Fraassen 1989) nicely illustrates how Bertrand-style paradoxes work. A factory produces cubes with side-length between 0 and 1 foot; what is the probability that a randomly chosen cube has side-length between 0 and 1/2 a foot? The tempting answer is 1/2, as we imagine a process of production that is uniformly distributed over side-length. But the question could have been given an equivalent restatement: A factory produces cubes with face-area between 0 and 1 square-feet; what is the probability that a randomly chosen cube has face-area between 0 and 1/4 square-feet? Now the tempting answer is 1/4, as we imagine a process of production that is uniformly distributed over face-area. And it could have been restated equivalently again: A factory produces cubes with volume between 0 and 1 cubic feet; what is the probability that a randomly chosen cube has volume between 0 and 1/8 cubic-feet? Now the tempting answer is 1/8, as we imagine a process of production that is uniformly distributed over volume. What, then, is *the* probability of the event in question?

Logical interpretation

The *logical* interpretation of probability, developed most extensively by Carnap (1950), sees probability as an extension of logic. Traditionally, logic aims to distinguish valid from invalid arguments by virtue of the syntactic form of the premises and conclusion. (E.g., any argument that has the form

p

If p then q

Therefore, q

is valid in virtue of this form.) But the distinction between valid and invalid arguments is not fine enough: many invalid arguments are compelling, in the sense that the premises strongly support the conclusion—we will see an example of such an argument shortly. Carnap described this relation of "support' or "confirmation" as the logical probability that an argument's conclusion is true, given that its premises are true. He had faith that logic, more broadly conceived, could also give it a syntactic analysis. So according to this program, probability is a measure of the degree to which a sentence supports another sentence, where this could be determined by the syntactic forms of the sentences themselves.

The program did not succeed. A central problem is that changing the language in which items of evidence and hypotheses are expressed will typically change the confirmation relations between them—for example, adding further predicates or names to a given language will typically revise how probabilities are shared around individual sentences. Moreover, Goodman showed that inductive logic must be sensitive to the meanings of words, for syntactically parallel inferences can differ wildly in their inductive strength. For example,

All observed snow is white.

Therefore, all snow is white.

is an inductively strong argument: its premise gives strong support to its conclusion. However,

All observed snow is observed.

Therefore, all snow is observed.

is inductively weak, its premise providing minimal support for its conclusion. It is quite unclear how a notion of logical probability can respect these intuitions.

Frequency interpretations

Frequency interpretations date back to Venn (1876). Gamblers, actuaries and scientists have long understood that relative frequencies bear an intimate relationship to probabilities. Frequency interpretations posit the most intimate relationship of all: identity. In a sound bite, *probabilities are relative frequencies* on this view. Thus, the probability of '6' on a die that lands '6' 3 times out of 10 tosses is, according to the frequentist, 3/10. In general:

> the probability of an outcome *A* in a reference class *B* is the proportion of occurrences of *A* within *B*.

Frequentism is still the dominant interpretation among scientists who seek to capture an objective notion of probability, heedless of anyone's beliefs. It is also the philosophical position that lies in the background of the classical Fisher/Neyman-Pearson approach that is used in most statistics textbooks. Frequentism does, however, face some major objections. For example, a coin that is tossed exactly once yields a relative frequency of heads of either 0 or 1, whatever its true bias—an instance of the infamous ‘problem of the single case’. A coin that is tossed twice can only yield relative frequencies of 0, 1/2, and 1. And in general, a finite number *n* of tosses can only yield relative frequencies that are multiples of 1/*n*. Yet it seems that probabilities can often fall between these values. Quantum mechanics, for example, posits irrational-valued probabilities such as $1/\sqrt{2}$.

Some frequentists (notably Reichenbach, and von Mises 1957) address this problem by considering infinite reference classes of hypothetical occurrences. Probabilities are then defined as limiting relative frequencies in suitable infinite sequences of trials. Von Mises offers a sophisticated formulation based on the notion of a *collective:* an (hypothetical, or virtual) infinite sequence of 'attributes' (possible outcomes) of a specified experiment that is performed infinitely often. He goes on to lay down two requirements for such an infinite sequence ω to be a collective. Call a *place-selection* an effectively specifiable method of selecting indices of members of ω, such that the selection or not of the index i depends at most on the first i – 1 attributes. The axioms are:

Axiom of Convergence: the limiting relative frequency of any attribute exists.

Axiom of Randomness: the limiting relative frequency of each attribute in a collective ω is the same in any infinite subsequence of ω which is determined by a place selection.

The probability of an attribute A, relative to a collective ω, is then defined as the limiting relative frequency of A in ω.

Collectives are abstract mathematical objects that are not empirically instantiated, but that are nonetheless posited by von Mises to explain the stabilities of relative frequencies in the behaviour of actual sequences of outcomes of a repeatable random experiment. Church (1940) renders precise the notion of a place selection as a recursive function.

If there are in fact only finitely many trials of the relevant type, then this kind of frequentism requires the actual sequence of outcomes to be extended to a hypothetical or 'virtual' sequence. This creates new difficulties. For instance, there is apparently no fact of the matter of how the coin in my pocket would have landed if it had been tossed indefinitely—it *could* yield any hypothetical limiting relative frequency that you like. Moreover, a well-known problem for any version of frequentism is the *reference class problem*: relative frequencies must be relativized to a *reference class*. Suppose that you are interested in the probability that Collingwood will win its next match. Which reference class should you consult? The class of all matches in Collingwood's history? Presumably not. The class of all recent Collingwood matches? That's also unsatisfactory: it is somewhat arbitrary what counts as 'recent', and some recent matches are more informative than others regarding Collingwood's prospects. The only match that resembles Collingwood's next match in every respect is that match itself. But then we are saddled again with the problem of the single case, and we have no guidance to its probability in advance. The reference class problem can also be a very practical problem—insurance companies face it on a daily basis. After all, the premiums that they set for a given individual are based on frequencies of claims of people of that type; but the individual is a member of many classes of people, whose relevant frequencies may differ wildly.

Propensity interpretations

Propensity interpretations, like frequency interpretations, regard probability as an objective feature of the world. Probability is thought of as a physical propensity, or disposition, or tendency of a system to produce given outcomes. This view, which originated with Popper (1959b), was motivated by the desire to make sense of single-case probability attributions on which frequentism apparently foundered, particularly those found in quantum mechanics. Propensity theories fall into two broad categories. According to *single-case* propensity theories, propensities measure a system's tendencies to produce given outcomes; according to *long-run* propensity theories, propensities are tendencies to produce long-run outcome frequencies over repeated trials. See Gillies (2000) for a useful survey.

Single-case propensity attributions face the charge of being untestable. Long-run propensity attributions may be considered to be verified if the long-run statistics agree sufficiently well with those expected, and falsified otherwise; however, then the view risks collapsing into frequentism, with its attendant problems. A prevalent objection to any propensity interpretation is that it is uninformative to be told that probabilities are 'propensities'. For example, what exactly is the property in virtue of which this coin has a 'propensity' of 1/2 of landing heads (when suitably tossed)? Indeed, some authors regard it as mysterious whether propensities even obey the axioms of probability in the first place. (See Hájek 2003a.)

Subjectivist interpretions

Subjectivist interpretations—sometimes called 'Bayesian'—pioneered by Ramsey (1926) and de Finetti (1937), see probabilities as *degrees of belief,* or *credences* of appropriate agents. These agents cannot be actual people since, as psychologists have repeatedly shown, people typically violate probability theory in various ways, often spectacularly so. Instead, we imagine the agents to be ideally rational.

But what are credences? De Finetti identifies an agent's subjective probabilities with his or her betting behavior. For example,

your probability for the coin landing heads is

if and only if

you are prepared to buy or sell for 50 cents a ticket that pays: $1 if the coin lands heads, nothing otherwise.

All of your other degrees of belief are analyzed similarly.

The analysis has met with many objections. Taken literally, it assumes that opinions would not exist without money, and moreover that you must value money linearly; but if it is just a metaphor, then we are owed an account of the literal truth. Even if we allow other prizes that you value linearly, problems remain. For your behavior in general, and your betting behavior in particular, is the result of your beliefs and desires working in tandem; any such proposal fails to resolve these respective components. Even an ideally rational agent may wish to misrepresent her true opinion; or she may particularly enjoy or abhor gambling; or, like a Zen master, she may lack a desire for worldly goods altogether. In each case, her betting behavior is a highly misleading guide to her true probabilities.

A more sophisticated approach, championed by Ramsey, seeks to fix an agent's utilities and probabilities simultaneously by appeal to her preferences. Suppose that you have a preference ranking over various possible states of affairs and gambles among them, meeting certain conditions required by rationality (for example, if you prefer A to B, and B to C, then you prefer A to C). Then we can prove a 'representation' theorem: these preferences can be represented as resulting from an underlying probability distribution and utility function. This approach avoids some of the objections to the betting interpretation, but not all of them. Ramsey's method essentially appeals to preferences over gambles, raising again the concern that the wrong quantities are being measured. And notice that the representation theorem does not show that rational agents' opinions *must* be represented as probabilities; it merely shows that they *can* be, leaving open that they can also be represented in *other,* radically different ways.

Radical subjectivists such as de Finetti recognize no constraints on initial, or 'prior' subjective probabilities beyond their conformity to Kolmogorov's axioms. But they typically advocate a learning rule for updating probabilities in the light of new evidence. Suppose that you initially have a probability function $P_{initial}$, and that you become certain of E (and of nothing more). What should be your new probability function P_{new}? The favoured updating rule among Bayesians is conditionalization; P_{new} is related to $P_{initial}$ as follows:

(Conditionalization) $P_{new}(X) = P_{initial}(X \mid E)$ (provided $P_{initial}(E) > 0$).

Radical subjectivism has been charged with being too permissive. It apparently licenses credences that we would ordinarily regard as crazy. For example, you can assign without its censure probability 0.999 to your navel ruling the universe—provided that you remain coherent (and update by conditionalization). Radical subjectivism also seems to allow inferences that are normally considered fallacious, such as 'the gambler's fallacy' (believing, for instance, that after a surprisingly long run of heads, a fair coin is more likely to land tails). Rationality, the objection goes, is not so ecumenical.

A standard defence (e.g., Savage 1954, Howson and Urbach 1993) appeals to famous 'convergence-to-truth', and 'merger-of-opinion' results. Roughly, they say that in the long run, the effect of choosing one prior probability function rather than another is washed out: successive conditionalizations on the evidence will, with probability one, make a given agent eventually converge to the truth, and thus initially discrepant agents eventually come to agreement. Unfortunately, these theorems tell us nothing about how quickly the convergence occurs. In particular, they do not explain the unanimity that we in fact often reach, and often rather rapidly. We will apparently reach the truth 'in the long run'; but as Keynes quipped, "in the long run, we shall all be dead".

Conclusion

In this limited space I have tried to convey how tendentious the mathematical and philosophical foundations of probability remain, despite some 350 years of research in the area. The interested reader will find more discussion of some of the liveliest current debates, trends, and prospects for the future in Hájek and Hall (2002) and Fitelson, Hájek, and Hall (2005).

Feller (1957, 19) writes: "All possible definitions of probability fall short of the actual practice." Certainly, a lot is asked of the concept of probability. In a suitably reflexive postmodern moment, I will complete this survey with some of my own bets on the uncertain future of the field.

I wager that we will continue to appeal to some *quasi-logical* notion of probability—for the evidential relations between various sentences or propositions are hardly exhausted by 'entailment' and 'refutation', the stock-in-trade of deductive logic. Confirmation theory, pioneered by Hempel (1945) and Carnap (1950), is making a big comeback in philosophy. Arguably, the leading approach is probabilistic—sometimes called *Bayesian confirmation theory*. Its central idea is simple: confirmation relations are identified with dependence relations. Thus, we may say that, relative to probability function *P:*

- *E* confirms *H* iff $P(H \mid E) > P(H)$.

Note that this notion of confirmation is *incremental* in the sense that *E increases* the amount of evidence for *H,* without necessarily leaving *H* highly supported. Thus, a coin's landing heads on the first toss confirms its landing heads 100 times in a row. Similarly,

- *E* disconfirms *H* iff $P(H \mid E) < P(H)$.
- *E* is evidentially irrelevant to *H* iff $P(H \mid E) = P(H)$.

See Hájek and Joyce (forthcoming) for a survey of confirmation theory.

I see a healthy future for *objective probability*, or *chance*, underpinning the aleatoric, indeterministic aspects of the mind-independent world, such as we apparently find in radioactive decay. I find especially promising approaches that ground chance in physical *symmetries*—see e.g. Strevens (1998).

And we will need the notion of *degrees of belief* or *credences* as long as there is uncertainty—which is to say, as long as there is human thought. But radical subjectivism is, to my mind, *too* radical—remember the navel ruling the universe! It needs to be constrained by something objective. For example, Lewis's (1980) *Principal Principle* says roughly that rational credences strive to align with chances, so that if a rational agent knows the chance of a given outcome, her degree of belief will be the same. More generally, where 'P' is the subjective probability function of a rational agent, and 'ch' is the chance function,

$$P(A \mid ch(A) = x) = x, \text{ for all } A \text{ and for all } x \text{ such that } P(ch(A) = x) > 0. \text{ [1]}$$

For example, my degree of belief that this coin toss lands heads, given that its chance of landing heads is 1/4, is 1/4.

Perhaps one would do better to think of these as distinct *concepts* of probability, to be sure with some important interrelations—we have already seen one such interrelation in the Principal Principle (and see Hájek 2003a for more.) Each of the leading interpretations, then, attempts to illuminate one of these concepts, while leaving the others in the dark. In that sense, the interpretations might be regarded as complementary, although to be sure each will need some further refinement.

Clearly, much work remains to be done on the foundations of probability. Equally clearly, we have come a long way since the *Port Royal Logic* and Bishop Butler.

Autographical and pedagogical epilogue

I began with some brief autobiographical and pedagogical reflections, and so I will end. You have just taken a crash course in the philosophical foundations of probability, a high-speed version of a course that filled 10 weeks at Caltech when I used to teach there. My

[1] There are subtleties that I cannot go into here, including the notion of admissibility, the relativity of chances to times, and Lewis' (1994b) revised version of the Principle.

students were typically budding scientists and engineers, and I tried to bring the material to life for them by emphasizing how ubiquitous probability is, and how often high-stakes decisions are made on the basis of probability judgments. I used to begin with this example:

> On January 28, 1986 at 11:38 a.m., the space shuttle Challenger was launched in Florida. Seventy-three seconds later it exploded, setting back the American manned space program by years. Managers made the decision to launch, against the advice of engineers, on the basis of a superficial and flawed analysis of the probability that the two solid rocket motors would fail at low temperatures, leading to a serious underestimate of that probability (Dalal et al. 1989). Lacking a clear conception of probability—and with it, a well understood, universally accepted methodology for determining probabilities—otherwise careful engineers and managers resorted to ad hoc calculations and dubious rules of thumb that resulted in tragedy. In particular, I believe that none of the parties concerned truly understood the notion of the single-case probability of disaster that was appropriate for Challenger. And yet its launch that day, in exactly the conditions that prevailed, was by its very nature unrepeatable.

Nor have the scientists (even at Caltech!) succeeded in understanding probability. I used to set my students the following question on the final exam for my course:

> In Feynman's *Lectures on Physics*, Volume 1, we find the following "definition" of probability:
>
> > By the "probability" of a particular outcome of an observation we mean our estimate for the most likely fraction of a number of repeated observations that will yield that particular outcome.
>
> There are many problems with this definition. Briefly indicate several of them.

Now that you, dear reader, have seen some of the problems with frequentism, you should be able to make a good start on this question. Here are some further hints:

- the definition is circular;
- it is easy to come up with cases in which there is more than one "most likely fraction";
- irrational probabilities, such as $1/\sqrt{2}$ are excluded—yet our best physical theory,

quantum mechanics, freely assigns such probabilities!

Finally, to bring home the subjective interpretation of probability in a way that I hope the students will *never* forget, I used to give them a multiple-choice test with a twist. Rather than choosing a correct answer, they had to assign *credences* to each possible answer.[2] The test began with the following explanation:

> Each of the following questions has exactly one correct answer among the choices a – d. I would like you to assign subjective probabilities to each of the choices, representing in each case your own probability that *that* choice is correct. For example, suppose you are nearly certain that b. is the correct answer to a given question, and the other choices look about equally implausible to you. Then you might represent your opinion as follows:
>
> a. 0.01
> b. 0.97
> c. 0.01
> d. 0.01
>
> For each question, you will receive a score determined by the probability you gave to the correct answer. Let that probability be p. Your score for that question will be:
>
> Score = $1 + 1/2\log_2 p$
>
> Thus, if you give probability 1 to the correct answer to a question, and 0 to the rest, you get a perfect score of 1 for that question; if you give 0 to the correct answer, you get a score of $-\infty$ for that question. (Totals less than 0 will be rounded up to 0.) Make sure your probabilities for a given question are non-negative, and add up to 1—otherwise you get 0 for that question automatically!

You, dear reader, might like to try your hand at the first question on my test, reprinted below. Good luck!

> Q1. Let Ω be a non-empty set. Which of the following provides a correct characterization of a set F of subsets of Ω being a **sigma algebra on** $\boldsymbol{\Omega}$?
>
> a. $\Omega \in F$; if $A \in F$, then $\neg A \in F$; and if $A_1, A_2, \ldots$ is a sequence of pairwise disjoint sets, each one $\in F$, then their countable union $\cup A_n \in F$.

[2] David Dowe of Monash University devised a similar 'probabilistic football betting' system, and I am grateful to him for suggesting the scoring rule.

b. F is non-empty, closed under complementation (with respect to Ω) and under countable intersection.

c. $\varnothing \in$ F; F is closed under complementation (with respect to Ω) and under finite union.

d. F is the power set of Ω.

Post-script: almost every year at least one student would get a score of $-\infty$.

Philosophy Program

Research School of the Social Sciences

Australian National University

Canberra, ACT 0200

Bibliography

Bertrand, J. (1989): *Calcul des Probabilités*, 1st edition, Gauthier-Villars.

Butler, Joseph (1961): *Analogy of Religion*, 1736; reprinted by Frederick Ungar Publishing Company.

Carnap, Rudolf (1950): *Logical Foundations of Probability*, Chicago: University of Chicago Press; 2nd. edition. Chicago: University of Chicago Press.

Church, A. (1940): "On the Concept of a Random Sequence", *Bulletin of the American Mathematical Society* 46, 130-135.

De Finetti, Bruno (1937): "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré*, 7, 1-68; translated as "Foresight. Its Logical Laws, Its Subjective Sources", in *Studies in Subjective Probability*, eds. H. E. Kyburg, Jr. and H. E. Smokler. New York: Robert E. Krieger Publishing Co., 1980.

De Finetti, B. (1972): *Probability, Induction and Statistics*, Wiley, New York.

De Finetti, Bruno (1974): *Theory of Probability*, John Wiley, New York. Reprinted 1990.

Feller, William (1968): *An Introduction to Probability Theory and Its Applications*, 3rd ed., John Wiley & Sons, Inc..

Fine, Terrence (1973): *Theories of Probability*, New York: Academic Press.

Fitelson, Branden, Alan Hájek, and Ned Hall (2005): "Probability", in *The Routledge Encyclopedia of Philosophy of Science*, eds. Jessica Pfeiffer, Sherri Rausch, and Sahotra Sarkar, Routledge.

Gillies, Donald (2000): "Varieties of Propensity", *British Journal for the Philosophy of Science,* 51, 807-835.

Hájek, Alan and Ned Hall (2002): "Induction and Probability", in *The Blackwell Guide to the Philosophy of Science,* eds. Peter Machamer and Michael Silberstein, Blackwell, 149-172.

Hájek, Alan (2003a): "Probability, Interpretations of", *The Stanford Encyclopedia of Philosophy,* ed. E. Zalta, http://plato.stanford.edu/entries/probability-interpret/.

Hájek, Alan (2003b): "What Conditional Probability Could Not Be", *Synthese*, Vol. 137, No. 3, December 2003, 273-323.

Hájek, Alan and James M. Joyce (forthcoming), "Confirmation", in the *Routledge*

Companion to the Philosophy of Science, eds. Stathis Psillos and Martin Curd.

Hempel, Carl (1945): "Studies in the Logic of Confirmation", *Mind* **54**: 1-26, 97-121.

Howson, Colin and Peter Urbach (1993): *Scientific Reasoning: The Bayesian Approach*, 2nd ed., Open Court, Chicago.

Jeffrey, Richard (1983): "Bayesianism with a Human Face", in *Testing Scientific Theories*, ed. John Earman, *Minnesota Studies in the Philosophy of Science*, vol. **X**, 133-156.

Kolmogorov, Andrei. N. (1933): *Grundbegriffe der Wahrscheinlichkeitrechnung, Ergebnisse Der Mathematik;* translated as *Foundations of Probability*, New York: Chelsea Publishing Company, 1950.

Laplace, Pierre Simon (1814): *A Philosophical Essay on Probabilities,* New York: Dover Publications Inc., English edition 1951.

Levi, Isaac (1980): *The Enterprise of Knowledge*, The MIT Press.

Lewis, David (1980): "A Subjectivist's Guide to Objective Chance", in R. C. Jeffrey (ed.) *Studies in Inductive Logic and Probability*, Vol II., University of California Press, 263-293; reprinted in *Philosophical Papers Volume II,* Oxford University Press.

Li, Ming and Paul Vitanyi (1997): *An Introduction to Kolmogorov Complexity and its Applications*, 2nd ed., New York: Springer-Verlag.

Muckenheim, W., Ludwig, G., Dewdney, C., Holland, P., Kyprianidis A., Vigier, J., Petroni, N., Bartlett, M., and Jaynes. E. (1986): "A Review of Extended Probability". *Physics Reports, 133, 337–401.*

Popper, Karl (1959a): "The Propensity Interpretation of Probability", *British Journal of the Philosophy of Science* 10: 25–42

Popper, Karl (1959b): *The Logic of Scientific Discovery*, London: Hutchinson & Co..

Ramsey, Frank P., (1926): "Truth and Probability." In *Foundations of Mathematics and other Essays*, ed. R. B. Braithwaite, 156-198, Routledge & P. Kegan, 1931; reprinted in *Studies in Subjective Probability*, eds. H. E. Kyburg, Jr. and H. E. Smokler, 2nd ed., 23-52, New York: R. E. Krieger Publishing Co., 1980; reprinted in *Philosophical Papers*, ed. D. H. Mellor. Cambridge: University Press, Cambridge, 1990.

Seidenfeld, Teddy (2001): "Remarks on the Theory of Conditional Probability: Some Issues of Finite Versus Countable Additivity", in *Probability Theory,* V. F. Hendricks et al. (eds.), Kluwer Academic, 167-178. Also available at: http://philsci-archive.pitt.edu/archive/00000092/

Strevens, Michael (1998): "Inferring Probabilities From Symmetries", *Noûs* 32, 231–46.

van Fraassen, Bas (1989): *Laws and Symmetry,* Oxford: Clarendon Press.

van Fraassen, Bas (1990): "Figures in a Probability Landscape", in J.M. Dunn and A. Gupta (eds.), *Truth or Consequences*, Kluwer, 345-356.

Venn, John (1876): *The Logic of Chance*, 2nd ed. Macmillan and co; reprinted, New York, 1962.

von Mises, Richard (1957): *Probability, Statistics and Truth*, revised English edition, New York: Macmillan.