

The reference class problem is your problem too

Alan Hájek

Published online: 24 March 2007
© Springer Science+Business Media B.V. 2007

Abstract The reference class problem arises when we want to assign a probability to a proposition (or sentence, or event) X , which may be classified in various ways, yet its probability can change depending on how it is classified. The problem is usually regarded as one specifically for the frequentist interpretation of probability and is often considered fatal to it. I argue that versions of the classical, logical, propensity and subjectivist interpretations also fall prey to their own variants of the reference class problem. Other versions of these interpretations apparently evade the problem. But I contend that they are all “no-theory” theories of probability - accounts that leave quite obscure why probability should function as a guide to life, a suitable basis for rational inference and action. The reference class problem besets those theories that are genuinely informative and that plausibly constrain our inductive reasonings and decisions.

I distinguish a “metaphysical” and an “epistemological” reference class problem. I submit that we can dissolve the former problem by recognizing that probability is fundamentally a two-place notion: conditional probability is the proper primitive of probability theory. However, I concede that the epistemological problem remains.

Keywords Probability · Conditional probability · Reference class problem · Frequentist · Classical · Logical · Propensity · Subjectivist interpretations of probability · Kolmogorov · Popper

Well may we say, with Bishop Butler, that “probability is the very guide of life”. But ‘probability’ is a multifarious notion, while Butler’s aphorism implies that there is exactly one such guide. What sort of probability, then, is this guide?

A. Hájek (✉)
Research School of the Social Sciences,
Australian National University,
Canberra, ACT 0200, Australia
e-mail: alanh@coombs.anu.edu.au

We may think of probability theory at two levels: axiomatization and interpretation. At the level of axiomatization Kolmogorov's theory clearly reigns. He began by axiomatizing unconditional, or absolute, probability. He later defined conditional, or relative, probability as a ratio of unconditional probabilities according to the familiar formula:

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

At the level of interpretation we have an embarrassment of riches. Still, for better or for worse, some version of frequentism—the view that probabilities are suitably defined relative frequencies—continues to have the ascendancy among scientists. To be sure, among philosophers it is moribund—rightly so, in my opinion, and for many reasons (see Hájek, 1997). However, I will revisit one of the best-known arguments against frequentism, one that many consider fatal to it: the so-called *reference class problem*. For if it is fatal to frequentism, it is also fatal to most of the leading interpretations of probability. I will argue that versions of the classical, logical, propensity and subjectivist interpretations also fall prey to their own variants of the reference class problem. Other versions of these interpretations apparently evade the reference class problem. But I contend that they are all 'no-theory' theories of probability, accounts that leave quite obscure why probability should function as a guide to life, a suitable basis for rational inference and action. The reference class problem besets those theories that are genuinely informative and that plausibly constrain our inductive reasonings and decisions.

I distinguish a 'metaphysical' and an 'epistemological' reference class problem. I submit that we can dissolve the former problem by recognizing that probability is fundamentally a two-place notion: as I have argued elsewhere (2003a, b), conditional probability is the proper primitive of probability theory. However, I concede that the epistemological problem remains.

1 What is the reference class problem?

It is not surprising that the reference class problem originates, as far as I am aware, with Venn (1876)—the Venn of 'diagram' fame. After all, the problem is generated by the fact that any particular event belongs to various sets. He observes: "It is obvious that every individual thing or event has an indefinite number of properties or attributes observable in it, and might therefore be considered as belonging to an indefinite number of different classes of things..." (194). Then, he discusses how this leads to a problem in assigning probabilities to individuals, such as the probability that John Smith, a consumptive Englishman aged 50, will live to 61. He concludes: "This variety of classes to which the individual may be referred owing to his possession of a multiplicity of attributes, has an important bearing on the process of inference..." (196). An important bearing, indeed.

Reichenbach (1949) uses the term for this problem that has now become standard:

If we are asked to find the probability holding for an individual future event, we must first incorporate the case in a suitable reference class. An individual thing or event may be incorporated in many reference classes, from which different probabilities will result. This ambiguity has been called the *problem of the reference class*. (374)

With such clear statements of the problem by two of the most famous frequentists, it is perhaps also not surprising that the reference class problem has traditionally been regarded as a problem for *frequentism*; I think, moreover, that many people consider it to be the most serious problem that frequentism faces. But it *is* surprising that the ubiquity of the problem has not been adequately acknowledged. In one form or other, it strikes versions of *all* of the other leading interpretations of probability: the classical, logical, subjectivist, and propensity interpretations.¹

The reference class problem arises when we want to assign a probability to a single proposition, X , which may be classified in various ways, yet its probability can change depending on *how* it is classified. (X may correspond to a sentence, or event, or an individual's instantiating a given property, or the outcome of a random experiment, or a set of possible worlds, or some other bearer of probability.) X may be classified as belonging to set S_1 , or to set S_2 , and so on. *Qua* member of S_1 , its probability is p_1 ; *qua* member of S_2 , its probability is p_2 , where $p_1 \neq p_2$; and so on. And perhaps *qua* member of some other set, its probability does not exist at all. Now, there would be no problem worth speaking of if one way of classifying X , say as a member of S_k , stood out from the rest as being the 'correct' one; then it seems the probability of X would simply be p_k . The problem grows teeth to the extent that this is not the case—to the extent that there are equally good claimants for the probability of X . For it would seem that X can only have one (unconditional) probability. Nor, perhaps, would there be a problem worth speaking of if all the p_i 's were roughly equal; then at least the probability of X could be confined to a small interval, and in the happiest case, to a single value. The teeth grow sharper to the extent that these probabilities differ significantly from one another.

This is really a special case of a more general problem that I think still deserves to be called 'the reference class problem'. Let X be a proposition. It seems that there is *one unconditional* probability of X ; but all we find are *many conditional* probabilities of the form $P(X, \text{given } A)$, $P(X, \text{given } B)$, $P(X, \text{given } C)$, etc. that differ from each other. Moreover, we cannot recover $P(X)$ from these conditional probabilities by the law of total probability, since we likewise lack unconditional probabilities for A , B , C , etc. (and in any case A , B , C , etc. need not form a partition). Relativized to the condition A , X has one probability; relativized to the condition B , it has another; and so on. Yet none of the conditions stands out as being the *right* one.

Here it is important to distinguish a *metaphysical* problem and an *epistemological* problem. The former problem arises because it seems that there should be a fact of the matter about the probability of X ; what, then, is it? The latter problem arises as an immediate consequence: a rational agent apparently can assign only one (unconditional) probability to X : what, then, should that probability be? The former problem concerns what probabilities are 'out there'; as we will see, it concerns the very nature of probability itself. The latter problem concerns which probabilities should serve us as guides to life: which probabilities form appropriate bases for our inductive inferences and provide proper inputs to our decision-making. So perhaps we strictly should not speak of *the* reference class problem, as if there were only one problem, although obviously the two problems are closely related, and I will happily conflate them until

¹ Hild (in preparation b) does a fine job of arguing that the reference class problem arises for the propensity theorist and for the subjectivist who is constrained by the Principle of Direct Probability (see below). My treatment of those interpretations is inspired by and indebted to him. Gillies (2000) also has a nice discussion of how the reference class problem arises for propensity interpretations.

this paper's final section. I will argue that these problems will not go away simply by jettisoning frequentism.

2 The reference class problem and the leading interpretations of probability

Guidebooks to the interpretations of probability ritually list the following: frequentist, classical, logical, propensity, and subjectivist interpretations. This taxonomy is fine as far as it goes, but I will find it useful to refine it, dividing each of these species into two sub-species:

1. Frequentism: (i) actual and (ii) hypothetical.
2. Classical: (i) finite sample spaces, and (ii) infinite sample spaces.
3. Logical: (i) fully constrained and (ii) less constrained.
4. Propensity: (i) frequency- or symmetry-based and (ii) neither frequency- nor symmetry-based.
5. Subjectivism: (i) radical and (ii) constrained.

We will see that most of these accounts face their own version of the reference class problem. However, those that do not, achieve a hollow victory—they say precious little about what probability *is*, or leave mysterious why it should guide us. I will call them ‘no-theory theories’ to convey my dissatisfaction with them.²

2.1 Frequentism

Let us begin where the reference class problem supposedly begins, with frequentism. (Again, it is underappreciated that it does not *end* there.) Versions of frequentism, as I have said, identify probability with relative frequency, differing in the details of how this is to be done. But the word ‘relative’ is already a warning that they will all face a reference class problem.

2.1.1 Actual frequentism

Actual frequentists such as Venn (1876) in at least some passages and, apparently, various scientists even today,³ identify the probability of an attribute or event A in a reference class B with the relative frequency of actual occurrences of A within B. Note well: *in a reference class B*. By changing the reference class we can typically change the relative frequency of A, and thus the probability of A. In Venn's example,⁴ the probability that John Smith, a consumptive Englishman aged 50, will live to 61, is the frequency of people like him who live to 61, relative to the frequency of all such people. But who are the people “like him”? It seems there are indefinitely many ways of classifying him, and many of these ways will yield conflicting verdicts as to the relative frequency.

² I borrow this term from Sober (2000) and Hild (in preparation b), although I think that my usage differs slightly from each of theirs.

³ Witness Frieden (1991): “The word ‘probability’ is but a mathematical abstraction for the intuitively more meaningful term ‘frequency of occurrence’” (10).

⁴ See also Peirce (1867).

2.1.2 Hypothetical frequentism

Hypothetical frequentists such as Reichenbach (1949) and von Mises (1957) are inspired by the dictum that probability is *long-run* relative frequency, and they are well aware that the actual world may not deliver a long run of trials of the required sort. No matter—we can always consider instead a *hypothetical* sequence of trials that is as long as we want, and the longer, the better. In particular, an infinite sequence is as good as it gets.

Consider Reichenbach's formulation. We begin with two sequences of event-tokens: $\{x_i\}$, some members of which may belong to a class A , and $\{y_i\}$, some members of which may belong to a class B . Let $F^n(A, B)$ be shorthand for the ratio of two frequencies: the denominator is the frequency of x 's out of the first n that belong to A , while the numerator is frequency of (x_i, y_i) pairs out of the first n pairs for which x_i belongs to A and y_i belongs to B . We can now state Reichenbach's definition:

If for a sequence pair $x_i y_i$ the relative frequency $F^n(A, B)$ goes toward a limit p for $n \rightarrow \infty$, the limit p is called the probability from A to B within the sequence pair. (69)

What of the probability of an event-token? As we have seen, Reichenbach says that we must incorporate it in "a suitable reference class". Suppose that we are interested in the probability that a given coin lands heads on a given toss. Suppose further, following Reichenbach, that we toss our coin repeatedly, interspersing it with tosses of another coin. All the tosses, in order, constitute our sequence $\{x_i\}$, some members of which belong to the class A of all tosses of our coin. Let $\{y_i\}$ be the sequence of outcomes of all the tosses, some members of which belong to the class B of 'heads' outcomes. Given this specification of x_i , y_i , A , and B , we may suppose that the probability from 'all tosses of our coin' to 'heads' is well-defined (non-trivial though the supposition is). But we could have specified our event differently—for example, as a toss of our coin with such-and-such angular momentum, or within a certain time period. That is, there are various candidates for A . *Qua* member of one candidate for A , we get one answer for our event's probability; *qua* member of another, we get another. What, then, is *the* probability of our event?

This is already enough of a reference class problem, but there is more to come. For a sequence has more structure than a set: its members are *ordered*. So even fixing the set of outcomes (our first reference class problem), there is the further problem of choosing among infinitely many orderings of its members. Probabilities must be relativized not merely to a reference *class* (a set), but to a sequence within the reference class. We might call this the *reference sequence problem*.

The beginnings of a solution to this problem would be to restrict our attention to sequences of a certain kind, those with certain desirable properties. von Mises restricts his to what he calls *collectives*—hypothetical infinite sequences of attributes (possible outcomes) of specified experiments that meet certain requirements. Call a *place-selection* an effectively specifiable method of selecting indices of members of the sequence, 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.

Church (1940) renders precise the notion of a place selection as a recursive function. The probability of an attribute A , relative to a collective ω , is then defined as the limiting relative frequency of A in ω . Nevertheless, the reference sequence problem remains: probabilities must always be *relativized to a collective*, and for a given attribute such as ‘heads’, or ‘dying by age 61’ there are infinitely many. Von Mises embraces this consequence, insisting that the notion of probability only makes sense relative to a collective. In particular, he regards single-case probabilities as “nonsense”. For example: “We can say nothing about the probability of death of an individual even if we know his condition of life and health in detail. The phrase ‘probability of death’, when it refers to a single person, has no meaning at all for us” (11).

Note that von Mises understates his theoretical commitments in two ways. First, he should also say that the phrase ‘probability of death’ has no meaning at all even when it refers to a million, or a billion, or any finite number of people, since they do not a collective make. While his skepticism about single-case probabilities will be shared by many others (for example, most ‘classical’ statisticians), his skepticism about finite-case probabilities, however large the finite numbers involved, is rather more radical, and it raises serious doubts about the applicability of his theory. Second, *even granting him infinitely many cases*—infinitely many persons, or what have you—probability statements still have no meaning at all for him. They only acquire meaning when the cases are *ordered*. Yet such an ordering may apparently be extrinsic to the cases themselves, imposed on them from the outside. If there is no ‘natural’ ordering (whatever that may mean), or if there are multiple equally ‘natural’ orderings (whatever that may mean), the choice of ordering presumably is imposed by us. Subjectivism threatens, in virtue of the reference sequence problem (and perhaps also in the judgment of what is ‘natural’)—and I doubt that von Mises would have welcomed *this* commitment.

On the other hand, Venn and Reichenbach face the reference class problem head on, and they go on to give similar prescriptions for choosing a privileged reference class, and thus a privileged probability. Reichenbach puts it this way: “We then proceed by considering the narrowest class for which reliable statistics can be compiled” (374).

But this prescription is patently inadequate. When are statistics “reliable”? This suggests more than just sufficiently large sample size (it had better include, for example, unbiasedness)—and even that notion is all too vague. It is surely also context-dependent, sensitive to pragmatic considerations such as the weighing of utilities. (We may, for example, tolerate smaller reference classes in studying the formation of white dwarves than in testing the safety of a new drug.) Worse, as various philosophers have observed⁵ there may be multiple classes that *prima facie* are equally narrow, and for which “reliable” statistics can be compiled. Suppose that there are reliable statistics on the deaths of Englishmen who visit Madeira, and of consumptives who visit Madeira, but not on consumptive Englishmen who visit Madeira. John Smith is a consumptive Englishman visiting Madeira. In which class should we place him? Worst, the last objection may still have conceded too much to the prescription, because it conceded that reference classes can be *totally ordered* according to their narrowness, so that we can judge when various classes are equally narrow. But even this is far from obvious (hence the hedge “*prima facie*” above). Can we even compare the narrowness of the class of Englishmen who visit Madeira and that of the class of consumptives who visit

⁵ For example, Gillies (2000, 816), who attributes the point to David Corfield and to Jon Williamson.

Madeira? To be sure, they are both narrower than the class of all people who visit Madeira. But the mere fact that they each refine that class through the application of one further predicate ('Englishman', 'consumptive') is by itself no reason to judge them as *equally* narrow. At best, it seems that 'narrowness of reference class' induces a *partial* ordering; we may not be able even to compare the narrowness of various reference classes when they only partially overlap with each other.

We thus see that even in finite cases we may have no idea how to follow Venn's and Reichenbach's advice. Perhaps we should define the 'narrowness' of a class as simply the size of the class: the fewer members a class has, the narrower it is. We might stipulate that a reference class of 20, say, is the narrowest reference class that we will consider in our assessments of probability (though, why?). I doubt that this is what Reichenbach meant: appending the words "for which reliable statistics can be compiled" then adds nothing, for we have already stipulated that 20 data points, and no fewer, count as reliable. In infinite cases, one wonders what the advice *could* mean. Probabilities are supposedly determined by limiting relative frequencies in denumerable sets of cases; but one denumerable set is exactly as 'narrow' as another. And that's before we even worry about *ordering* the cases!

2.2 Classical probability

The classical interpretation (Laplace, 1814) converts information about numbers of *possibilities* into information about *probabilities*. It is assumed that we can partition the space of possible outcomes into a set of 'equipossible' outcomes. When this set is finite, the probability of proposition X is simply the fraction of the total number of these possibilities in which X occurs; when the set is infinite, we may finitely partition it into equipossible sets, which can still be regarded as outcomes. Outcomes are equipossible if there is no evidence to the contrary—no evidence that favors some outcomes over others. This is the infamous 'principle of indifference'. We have two cases here: outcomes for which we have *no evidence at all*, and outcomes for which we have *symmetrically balanced evidence*. We will see that the reference class problem looms either way.

Note the structural resemblance of classical probability to frequentism. A set of outcomes is chosen; probability is identified as a ratio of the number of favorable outcomes to the total number of outcomes. The only significant difference is in the nature of the outcomes: in frequentism they are the outcomes of a repeated experiment, while in the classical interpretation they are the possible outcomes of a single experiment. Small wonder, then, that classical probabilities will face a reference class problem much like frequentism's.

2.2.1 Finite sample spaces

There can be no set of equipossibilities without a set of possibilities. In applications of the classical theory in which the sample space is finite, they are invariably one and the same. Thus, the two possible ways a coin might land or the six possible ways a die might land are identified as exactly the *equipossibilities* of the respective spaces; landing on an edge, for instance, is not even considered as a possibility in the first place. A reference class problem arises for the classical theory in the choice of a sample space—the set of outcomes to which an outcome of interest O belongs. *Qua*

member of one set of outcomes, we get one answer for O 's probability; *qua* member of another set of outcomes, we get another answer; and so on.

If there is such a thing as a situation in which we literally have no evidence at all, then presumably there is nothing to distinguish various competing choices of sample space. We should then be indifferent between an original space and various expansions of that space to include further possibilities (when that space was not logically exhaustive), and indifferent between various refinements of the original space.

For an example of the former case, if we really have *no* evidence regarding coin-tossing, then we should be indifferent between the sample space {heads, tails} and {heads, tails, edge}, and even {heads, tails, edge, heads-edge-of-edge, tails-edge-of-edge}; and so on. *Qua* member of the first set, 'heads' gets probability 1/2; *qua* member of the second set, 'heads' gets probability 1/3; and so on.

For an example of the latter case, we should likewise be indifferent between various spaces that refine the 'heads' outcome according to its final orientation. We could partition 'heads' according to the various angles from north in which that face could end up oriented:

{heads oriented within $[0, 180^\circ)$ of north, heads oriented within $[180^\circ, 360^\circ)$ of north, tails},

{heads within $[0, 120^\circ)$, heads within $[120^\circ, 240^\circ)$, heads within $[240^\circ, 360^\circ)$, tails}

and so on ad infinitum. *Qua* member of the first set, 'heads' gets probability 2/3 (occurring as it does in two of the three possible outcomes); *qua* member of the second set, 'heads' gets probability 3/4; and so on. Thus, if we really have *no* evidence, then we have the reference class problem in spades, for probabilities will be acutely sensitive to an apparently arbitrary choice of sample space.

In practice we choose the {heads, tails} space, of course, because we have *some* evidence: we know enough about the physics of coin-tossing or about the frequencies of outcomes to know, among other things, that 'edge' is far less probable than 'heads' or 'tails', that 'heads' is equally likely to point in each direction, and so on. We rarely enter a random experiment in a state of total epistemic innocence. When we do, however, the reference class problem is unavoidable. To adapt a well-known example from physics, Bose–Einstein statistics, Fermi–Dirac statistics, and Maxwell–Boltzmann statistics each arise by considering the ways in which particles can be assigned to states, and then partitioning the set of alternatives in different ways. (See e.g., Fine, 1973.) Someone ignorant of which statistics apply to a given type of particle—and this state of mind is easy to imagine!—can only make an arbitrary choice and hope for the best. And yet the classical interpretation is purported to be the one that applies in the face of ignorance.

In typical applications of the classical theory, the work is really done by the 'symmetrically balanced evidence' clause of the definition. There are two potential sources of a reference class problem here: in the 'evidence', and in the 'symmetry'. The most obvious characterization of symmetrically balanced evidence is in terms of equality of conditional probabilities: given evidence E and possible outcomes O_1, O_2, \dots, O_n , the evidence is symmetrically balanced iff $P(O_1|E) = P(O_2|E) = \dots = P(O_n|E)$.⁶ (One might reasonably wonder how *these* probabilities are determined.) Be that as it may,

⁶ One might use a non-probabilistic account of evidential support here—for example, *plausibility* theory. (See Hild (in preparation a).) The danger is that we won't recover *probabilities* at any stage. Whatever else the classical theory might be, it is at least supposed to be a theory of *probability*.

it is clear that classical probabilities are acutely sensitive to the evidence—not in the sense that they might *change* if the evidence changes, but in the sense that they might *vanish!* That is, if the evidence becomes *unbalanced*, favoring some outcomes over others, then classical probabilities are not merely revised, they are destroyed. Relativized to one piece of evidence, the classical probability of a particular proposition has one value; relativized to another, the probability fails to exist! This is a particularly interesting form of reference class problem, and it occurs with equal force in both finite and infinite sample spaces. The problem of competing respects of symmetry, for its part, occurs most strikingly in infinite sample spaces, to which we now turn.

2.2.2 Infinite sample spaces

When the space of possibilities is infinite, the equipossibilities must be distinct from them. In fact, they must be a finite partition of equivalence classes of the space—otherwise their probabilities could not all be the same *and* sum to 1. The reference class problem arises immediately, then, in the choice of partition: among the infinitely many possible, one is chosen as the basis for the assignment of classical probabilities. But which?

This would not be a problem if one symmetric partition stood out among all of them. However, as Bertrand's paradoxes (1889) teach us, there are symmetries and then there are symmetries. The paradoxes turn on conflicting applications of the principle of indifference, each of which seems equally compelling. Some presentations of this species of paradox are needlessly arcane: I would like to conduct a poll to determine how many philosophers can correctly define 'specific gravity', as found in famous presentations by von Kries, Keynes and Nagel among others. Length and area suffice.⁷ Suppose that we have a square with side-length between 0 and 1 foot. What is the probability that its side-length is between 0 and 1/2 a foot? You have been told so little that ignorance over the possible side-lengths is guaranteed. In particular, it would seem that the intervals (0, 1/2) and [1/2, 1) are equipossible for containing the side-length. Applying the principle of indifference, then, the answer appears to be 1/2. But the question could be equivalently formulated: we have a square with area between 0 and 1 square foot. What is the probability that its area is between 0 and 1/4 square feet? Now the intervals (0, 1/4), [1/4, 1/2), [1/2, 3/4) and [3/4, 1) appear to be equipossible for containing the area. Applying the principle of indifference, then, the answer appears to be 1/4. *Qua* member of one set of alternatives, we get one answer for our proposition's probability; *qua* member of another set of alternatives, we get another. And we could have asked the question equivalently in terms of the cube of the side-length, and the 17th root of the side-length, and so on.

Note that in continuous cases such as this, the problem does not lie with differently weighting the individual outcomes, the infinitely many possible side-lengths or areas, for if it makes sense to speak of their weight at all, they are uniformly 0 on each of the formulations. Rather, the problem lies with the different ways in which the infinite space can be finitely partitioned, each 'symmetric' by some reasonable criterion. The reference class problem, then, is created by the simultaneous membership of a given outcome in multiple partitions that are not simply refinements of a fixed partition. The outcome that the length lies in the interval [0, 1/2) just *is* the outcome that the area lies in the interval [0, 1/4), but construed the first way it is an equal partner in a

⁷ Cf. van Fraassen (1989).

two-membered partition, whereas construed the second way it is an equal partner in a four-membered partition. And that's just the beginning. Classical probabilities must thus be relativized to a reference class.

2.3 Logical probability

Logical theories of probability retain the classical interpretation's idea that probabilities can be determined a priori by an examination of the space of possibilities. However, they generalize that interpretation in two important ways: the possibilities may be assigned *unequal* weights, and probabilities can be computed whatever the evidence may be, symmetrically balanced or not. Indeed, the logical interpretation, in its various guises, seeks to codify in full generality the degree of support or confirmation that a piece of evidence E confers upon a given hypothesis H , which we may write as $c(H, E)$.

Keynes (1921), Johnson (1932), Jeffreys (1939), and Carnap (1950, 1963) all consider probability theory to be a generalization of logic. They regard statements of the form $c(H, E) = x$ as being either logically true or logically false. Moreover, Keynes, Johnson, Jeffreys, and the early Carnap think that there is exactly one correct measure of such support, one 'confirmation function'. The later Carnap gives up on this idea, allowing a family of such measures. We thus distinguish two versions of logical probability.

2.3.1 Fully constrained logical probability.

Let us concentrate on the early Carnap, since his is the most complete development of a fully constrained logical probability. While he allows logical probabilities to be determined even when the principle of indifference does not apply, symmetries are still essential to the determination of probabilities. This time, the objects to which probabilities are assigned are sentences in a formal language, and it will be symmetries among them that will hold the key to the assignment of logical probabilities. The language contains countably many names, denoting individuals, finitely many one-place predicates, denoting properties that the individuals may or may not have, and the usual logical connectives. The strongest statements that can be made in a given language are called *state descriptions*: each individual is described in as much detail as the language allows (that is, the application of each predicate to each individual is either affirmed or denied). Equivalence classes of such state descriptions can be formed by permuting the names; these equivalence classes are called *structure descriptions*. We can then determine a unique measure m^* over the state descriptions, which awards equal measure to each structure description, and divides this in turn within a structure description equally among its state descriptions. This induces a confirmation function, c^* , defined by:

$$c^*(H, E) = \frac{m^*(H \ \& \ E)}{m^*(E)}, \quad \text{where } m^*(E) > 0.$$

This is the confirmation function that Carnap favors.

And thus is born a reference class problem. There is no such thing as the logical probability of H , simpliciter, but only the probability of H evaluated in the light of this or that evidence. The relativity to an evidence statement is essential: change the evidence, and the degree of confirmation of H typically changes. So we can putatively

determine, say, the logical probability that the next emerald observed is green, *given* the evidence that a hundred (observed) emeralds were all green, or *given* the evidence that three (observed) emeralds are purple and one is vermilion, and so on. But what about the probability that the next emerald is green, *sans qualification*?

Given this plurality of evidence statements to which the degree of confirmation of a hypothesis may be relativized, Carnap's recommendation is to use one's *total evidence*: the maximally specific information at one's disposal, the strongest proposition of which one is certain. This corresponds to the frequentist's dictum to use the narrowest reference class. It is at best a pragmatic, methodological point, for all of the various conditional probabilities, with their conditions of varying strengths, are all well-defined, and logic/probability theory is indifferent among them. Logic, of course, cannot dictate what your total evidence is, so it cannot dictate the probability of H . Nor can logic fault one for basing one's probability judgment of H on this piece of evidence rather than that; by Carnap's lights, it can only fault one for getting the conditional probabilities of H , *given* these respective pieces of evidence, wrong. (Ayer, 1963 makes similar observations in an excellent discussion.)

In any case, when we go beyond toy examples it is unclear whether there really is such a thing as 'the strongest proposition' of which one is certain. Suppose you are playing a real life version of the Monty Hall problem. A prize lies behind one of three doors; you guess that it is behind door 1. Monty Hall, who knows where the prize is and is careful not to reveal the prize, reveals door 2 to be empty. It is a familiar point that you did not just learn that door 2 was empty; you also learned that Monty chose to reveal door 2 to be empty. But you also learned a host of other things: as it might be, that Monty opened the door with his right hand AND at a particular time on a particular date, AND with the audience gasping in a particular way. . . . Call this long conjunction X . Moreover, it seems that you also learned a potentially infinite set of *de se* propositions: 'I learned that X ', 'I learned that I learned that X ' and so on. Perhaps, then, your total evidence is the infinite intersection of all these propositions, although this is still not obvious—and it is certainly not something that can be represented by a sentence in one of Carnap's languages, which is finite in length.

But even this grants too much to the total evidence criterion. It goes hand in hand with positivism, and a foundationalist epistemology according to which there are such determinate, ultimate deliverances of experience. But perhaps learning does not come in the form of such 'bedrock' propositions, as Jeffrey (1992) has argued—maybe it rather involves a shift in one's subjective probabilities across a partition, without any cell of the partition becoming certain. Or perhaps learning is even less determinate than Jeffrey would have it—maybe one's probabilities across such a partition can remain vague, or some of them even undefined. In any of these cases, the strongest proposition of which one is certain is expressed by a tautology T . (Indeed, it might be tempting to think that in any case, the unconditional logical probability of a hypothesis, H , is just the probability of H , given T .) That is hardly an interesting notion of 'total evidence'.

And still a reference class problem remains. For H and T will have to be formulated in some language or other. Which one will it be? Notoriously, Carnap's logical probabilities are acutely sensitive to the choice of language: change the language, and you change the values of the confirmation function. So there is no single value of $c^*(H, T)$, but rather a host of different values corresponding to different choices of the language. And more generally, the notation $c^*(H, E)$ is somewhat misleading,

suppressing as it does the dependence of c^* on the language L for all H and E . This is a second reference class problem: a hypothesis and evidence statement must be incorporated into a *set* of statements (the set of state descriptions and disjunctions thereof), and the logical probability linking them can only be evaluated relative to that set. *Qua* members of one set of statements, we get one answer for the probability of interest; *qua* members of another set, we get another answer; and so on. Thus Carnap's logical probabilities are doubly relativized: first to the specification of the evidence proposition, and second to the choice of language.

2.3.2 Less constrained logical probability

And they eventually become triply relativized. Carnap later generalizes his confirmation function to a continuum of functions c_λ . Define a *family* of predicates to be a set of predicates such that, for each individual, exactly one member of the set applies, and consider first-order languages containing a finite number of families. Carnap (1963) focuses on the special case of a language containing only one-place predicates. He lays down a number of axioms concerning the confirmation function mostly symmetry principles. They imply that, for a family $\{P_n\}$, $n = 1, \dots, k, k > 2$:

$$c_\lambda(\text{individual } s + 1 \text{ is } P_j, s_j \text{ of the first } s \text{ individuals are } P_j) = \frac{s_j + \lambda/k}{s + \lambda},$$

where λ is a positive real number.

The higher the value of λ , the less impact evidence has: induction from what is observed becomes progressively more swamped by a classical-style equal assignment to each of the k possibilities regarding individual $s + 1$.

A new source of relativization thus appears: logical probabilities now depend also on λ , on how 'cautious' is our inductive system. But nothing in logic, probability theory, or anything else for that matter seems to dictate a unique setting of λ . When John Smith asks us which value of λ he should use as he considers whether or not to buy life insurance, what should we tell him? More generally, three things now determine the reference class of a hypothesis H whose probability we might seek: as before, the language in which H is formulated and the evidence relative to which it is evaluated, and now λ .

2.4 Propensity interpretations

Propensity theorists think of probability as a physical propensity, or disposition, or tendency of a given type of physical situation to yield an outcome of a certain kind, or to yield a long run relative frequency of such an outcome. Popper (1959a), for example, regards a probability p of an outcome of a certain type to be a propensity of a repeatable experiment to produce outcomes of that type with limiting relative frequency p . Giere (1973), on the other hand, attributes propensities to single events. We may thus usefully distinguish *frequency-based* propensity theories from *non-frequency-based* propensity theories: the former appeal to relative frequencies in the characterization of propensities, while the latter do not. Frequency-based propensity theorists include Popper, and Gillies (2000); non-frequency-based propensity theorists include Giere, Mellor (1971), Miller (1996) and Fetzer (1977).

Frequency-based propensity theories will immediately inherit frequentism's reference class problem. After all, they are *relative* frequency-based. The relativization of

frequencies to a set of trials, or to an ordered sequence of trials, or to a collective, will transfer to whatever propensities are based on these relative frequencies.

What about non-frequency-based propensity theories? Some of them clearly invite a reference class problem in their very formulation. Consider Giere's (1973, 471) formulation, which takes as given a chance setup, *CSU*. He interprets the statement " $P(E) = r$ " as follows: "The strength of the *propensity* of *CSU* to produce outcome *E* on trial *L* is *r*".⁸ Propensities, then, are relativized to a chance setup (and a trial). There are other terms for much the same idea: the relativization of propensities to 'experimental arrangements', or 'test conditions', or what-not. It goes hand-in-hand with the view that "a propensity must be a propensity of something (*X*) to produce something else (*Y*)" (*ibid*, 472). On this picture, asking what the propensity of *Y* is, simpliciter, is like asking whether Renée is taller than, simpliciter.⁹

Some other non-frequency-based propensity theories face the reference class problem less obviously, but face it nonetheless. Here another distinction is useful: between *symmetry-based* propensity theories, which ground propensities in more basic physical symmetries, and *non-symmetry-based* theories, which don't. The former will meet the reference class problem in the specification of the appropriate respect of symmetry; the latter may avoid it, but I suspect only at the cost of being 'dormative virtue' theories in which the nature of probability remains obscure.¹⁰

I take Mellor (1971), for example, to be offering a symmetry-based propensity theory. A coin may be judged to be fair in virtue of its *physical* symmetries—its symmetric shape, its symmetric mass distribution, or what have you—with no regard for what the results of tossing it happen to be, nor even for whether it is tossed at all. More complicated systems, such as dice and roulette wheels, function as well as they do as gambling devices in virtue of other symmetries (or near-symmetries), albeit more complicated. Strevens (1998, and more fully in 2003), following Poincaré, Keller, and others, beautifully shows how stable objective probabilities arise out of low-level complexity, *in virtue of symmetries*.

Symmetry-based propensities look a fair bit structurally like classical probabilities, and so it should come as no surprise that the reference class problem arises for the former, much as it arose for the latter. Indeed, Mellor's "principle of connectivity" can be thought of as a counterpart to the principle of indifference. The idea is simple: propensities of outcomes are the same unless there is a difference in their causes. But this leads to a reference class problem, much as the principle of indifference did. If the outcomes of coin-tossing are genuinely indeterministic, with the outcomes 'heads' and 'tails' *uncaused*, then the principle of connectivity applies, and they must have the same propensity, 1/2. So far, so good. (Make the example quantum mechanical if you think that the results of coin tosses are caused by initial conditions.) But presumably various refinements of the outcomes are also uncaused. Recall the example of the various possible final orientations of 'heads'. Suppose that both heads-oriented-within $[0, 180^\circ)$ of north and heads-oriented-within $[180^\circ, 360^\circ)$ of north are also uncaused. So by the principle of connectivity, these too should get the same probability as heads, namely 1/2. But then we have a violation of additivity

⁸ It is unclear to me how the two statements can be equivalent when the former depends on just one variable, *E*, whereas the latter depends on three, *CSU*, *E*, and *L*.

⁹ See Levi (1990) for a valuable discussion of some of the motivations for relativizing chances to kinds of trials.

¹⁰ Sober (2000) also speaks of 'dormative virtue' theories of objective probability.

$(1/2 + 1/2 \neq 1/2)$, and propensities are not really probabilities at all. We could drop the principle of connectivity to avoid such unwelcome results, but then we would be left with a no-theory theory. Better to relativize the principle's application. Relative to one partition of outcomes, all of which have the same causes (or lack thereof), we have one set of propensities; relative to another, we have another. Propensities may be based on symmetries; but there are symmetries, and then there are symmetries. A reference class problem arises out of the sensitivity of propensities to the respect of symmetry.

Then there are propensity theories that do not appeal either to frequencies or to symmetries. Moreover, it appears that some of them do not face a reference class problem. For example, Miller (1994, 56) relativizes propensities at a given time to “the complete situation of the universe (or the light-cone) at the time”. Fetzer (1982 195) relativizes propensities in a world at a time to “a complete set of (nomically and/or causally) relevant conditions [...] which happens to be instantiated in that world at that time.” On these accounts, propensities must still be relativized—to a reference class, as I like to say—but perhaps there is only one relevant reference class. (Perhaps not, if there is *more than one* complete set of relevant conditions—Fetzer does say ‘*a*’, not ‘*the* complete set’—but let that pass.) In that case there seems to be no reference class problem. But do we truly know what such propensities are? Consider again poor old John Smith’s predicament at this moment. Imagine him having knowledge of the complete situation of the universe at this moment; or knowledge of a complete set of relevant conditions. We tell him that propensities are dependent on these things, but we do not tell him *how*. Granting him all the computational power that he might need, does he have any idea what is his propensity for living to 61, or even what this means? If not, I suspect that we have a no-theory theory of propensities.

I believe that this worry generalizes to other non-frequency-based, non-symmetry-based propensity theories. We may speak of “intrinsic properties of chance set-ups”, or “inherent dispositions”, or “tendencies,” or “graded modalities”, or . . . But unless we say more, we do not really know what is being said. This is ‘no-theory theory’ territory. In fact, it is then not clear why propensities have numerical values, still less why they should be generated by numerical *functions*, still less why they should be generated by numerical *probability* functions, still less why they are probabilities that serve as guides to life. I fear then that “propensity” is just a resonant name for something that we do not really understand. “Aleative virtue” would be equally resonant.¹¹

2.5 Subjectivism

Subjectivists regard probabilities as degrees of belief, and see the theorems of probability as rationality constraints on degrees of belief.

2.5.1 Radical subjectivism

Radical subjectivists such as de Finetti (1937) regard them as the *only* such constraints. Your degrees of belief can be whatever you like, as long as they remain probabilistically coherent. It thus appears that there is not any interesting reference class problem for the radical subjectivist. The probability that you assign to *E* is whatever it is. *Qua* nothing.

¹¹ I owe the term “aleative virtue” to Peter Godfrey-Smith.

This is a benefit, if that's the right word for it, of the radical subjectivist's permissive epistemology. But it comes at a cost. The epistemology is so *spectacularly* permissive that it sanctions opinions that we would normally call ridiculous. For example, you may with no insult to rationality assign probability 0.999 to George Bush turning into a prairie dog, provided that you assign 0.001 to this not being the case (and that your other assignments also obey the probability calculus). And you are no more or less worthy of praise if you assign it 0,¹² or 0.17485, or 1/e or whatever you like. Your probability assignments can be completely at odds with the way the world is, and thus are 'guides' in name only. Here some radical subjectivists may reply: "it is not *irrational* to have such assignments, so in that sense you should not be faulted". But you may also see dependences among events in ways that are properly regarded as irrational without setting off the radical subjectivist alarm. For example, having seen a long run of tosses of a coin land 'heads', you may be confident that the next toss will land 'tails' because you think that it is *due*. It is for good reason that we call this kind of counter-inductive reasoning the 'gambler's *fallacy*'. But the radical subjectivist is ill placed to use such an epithet, for there are coherent priors without number that license just such reasoning.

Being saddled with such unwelcome results is the price that the radical subjectivist pays for offering a no-theory theory of probability: there is so little constraining probability assignments that I wonder what interest is left in them. If you want exclusively to assign extremely high probabilities to contingent propositions that are in fact false, or to undergo perverse courses of 'learning' such as that of the fallacious gambler, you have the radical subjectivist's blessing. (Just stay coherent!) Probability theory becomes autobiography rather than epistemology.

2.5.2 Constrained subjectivism

But many subjectivists are more demanding of their subjects—and their further demands will bring reference class problems in their train. There are various proposals for extra constraints on rational opinion. I find it most perspicuous to present them all as instances of a certain canonical form. Gaifman (1988) coins the terms "expert assignment" and "expert probability" for a probability assignment that a given agent strives to track: "The mere knowledge of the [expert] assignment will make the agent adopt it as his subjective probability" (193). The guiding idea is captured by the equation

$$(\text{Expert}) \quad P(A|pr(A) = x) = x$$

where ' P ' is the agent's subjective probability function, and ' $pr(A)$ ' is the assignment that the agent regards as expert. For example, if you regard the local weather forecaster as an expert on matters meteorological, and she assigns probability 0.1 to it raining tomorrow, then you may well follow suit:

$$P(\text{rain} | pr(\text{rain}) = 0.1) = 0.1.$$

More generally, we might speak of an entire probability function as being such a guide for an agent. van Fraassen (1989), extending Gaifman's usage, calls pr an

¹² Some otherwise radical subjectivists impose the requirement of *regularity*—only logical impossibilities get assigned probability 0—in order to allow learning by repeated conditionalization on evidence. But then, staggeringly, an assignment of 0 to Bush's turning into a prairie dog is judged *irrational*, while an assignment of 0.999 is judged *rational*.

“expert function” for P if (Expert) holds for all x such that $P(\text{pr}(A) = x) > 0$, so that the conditional probability is defined. We should keep in mind the distinction between an expert function and an expert assignment, because an agent may not want to track all the assignments of her ‘expert’. (If your forecaster gives probability 0 to it raining in Los Angeles tomorrow, you may think that she’s gone too far, and you may not want to follow her there.)

An entire theory may provide an expert function for an agent. Quantum mechanics, for example, serves as my expert over all of the propositions that fall under its purview—I strive to track its probability assignments to those propositions as per (Expert). More generally, van Fraassen (1989) appeals to (Expert) to explicate what it means for an agent to accept a probabilistic theory.

Various other candidates for expert functions for rational agents have been proposed:

The *Principle of Direct Probability* regards relative frequencies that way. (See Hacking, 1965 for a presentation of it.) Let A be an event-type, and let $\text{relfreq}(A)$ be the relative frequency of A . Then for any rational agent with probability function P , we have

$$P(A|\text{relfreq}(A) = x) = x, \quad \text{for all } A \text{ such that } P(\text{relfreq}(A) = x) > 0.$$

Lewis (1980) posits a similar role for the *objective chance function*, ch , in his Principal Principle:

$$P(A|ch(A) = x) = x, \quad \text{for all } A \text{ such that } P(ch(A) = x) > 0.^{13}$$

A frequentist who thinks that chances just are relative frequencies would presumably think that the Principal Principle just is the Principle of Direct Probability; but Lewis’ principle may well appeal to those who have a very different view about chances—e.g., propensity theorists.

van Fraassen (1984, 1995), following Goldstein (1983), argues that one’s future probability assignments play such a role in constraining one’s present assignments in his Reflection Principle:

$$P_t(A|P_{t+\Delta}(A) = x) = x.$$

The idea is that a certain sort of epistemic integrity requires you to regard your future self as ‘expert’ relative to your current self.

One might also give conditionalized versions of these already-conditional principles, capturing the idea that an agent might want to track certain conditional probability assignments of her expert. (See van Fraassen, 1989, 201–202). For example, the Principal Principle might be amended in such a way:

$$P(A|ch(A|B) = x \cap B) = x.$$

Finally, if Carnap is to be believed, then *logical probability* plays such a role as expert—perhaps the ultimate one.

With various expert assignments in place, the reference class problem is poised to strike. In fact, it can now strike in two different, though closely related, ways: first, if

¹³ I ignore complications due to the time-relativity of chances, and Lewis’ notion of “admissibility”.

the expert assignments *disagree* with one another; second, if the expert assignments themselves are susceptible to a reference class problem. Let's take these points in order.

All is well if all your experts speak in unison—as it might be, your two favorite weather forecasters both assign a probability of 0.1 to it raining tomorrow, a meteorological theory that you accept assigns the same probability, 10% of days 'like' tomorrow were rainy days, your own research (perhaps on meteorological symmetries?) convinces you that the chance of rain is 0.1, your current probability assignments to your future probability assignments concur, and (somewhat fancifully) it turns out that the logical probability of rain tomorrow, given your putative total evidence, is the same again. But all may not be well. Suppose some of these numbers differ. You can't serve all your masters at once, so you have to play favorites. But who trumps whom, which trumps which? You have no difficulty forming a series of conditional probabilities, each of the form (Expert), with different functions playing the role of *Pr* in each case. Your difficulty arises in combining them to arrive at a single unconditional probability assignment. *Qua* proposition that your first expert assigns probability p_1 , you want to assign it p_1 ; *qua* proposition that your second expert assigns probability p_2 , you want to assign it p_2 ; and so on.

Now of course you can simply weight your various experts' assignments and combine them in some way in order to come up with your own assignment. But what are the weights to be? If they are totally unconstrained, then you risk collapsing into radical subjectivism, and its no-theory theory: make a particular weight 0.999, or 0, or 0.17485, or $1/e$, or whatever you like as long as all the weights add up to 1. (Just stay coherent!) But if the weights are constrained by something external—some *expert*—then you find yourself with further conditional probabilities, and no respite. For example, you might give more weight to one of your favorite forecasters than the other because he is better *calibrated*: his probability assignments in the past have been better vindicated by the relevant relative frequency data. But by now I hardly need to point out that any relative frequency data are *relativized*. Moreover, Simpson's paradox teaches us of the perils you can face when you mix conditional probability assignments: correlations that all of the experts see may be washed out or even reversed. Worse still, these correlations may be reversed again when we partition our probability space more finely—which is to say, when we refine our reference classes. Enter the reference class problem again.

A further, related problem arises when the expert functions assignments are susceptible to a reference class problem—and it seems to me that they invariably are. Consider again the Principle of Direct Probability: given its dependence on relative frequencies, it immediately inherits frequentism's reference class problem. Likewise, if the assignment by the chance function '*ch*' is susceptible to a reference class problem, so too will the corresponding assignment by the subjective probability function '*P*'. We simply have the previous case if we identify '*ch*' with relative frequency, à la Venn, or à la Reichenbach. (Not à la von Mises, since he eschews single-case chances, so the Principal Principle is never instantiated according to him.) This point generalizes to any account of '*ch*' that is thus susceptible—e.g., frequency-based propensities, symmetry-based propensities, accounts that appeal to 'chance set-ups', and what not. (There may be no reference class problem if '*ch*' is a no-theory theory propensity; but then we would be left wondering why the Principal Principle should have any claim on us, and how ever to apply the Principle.) And if logical probability is your expert, then its reference class problems are yours.

As for human ‘experts’ (weather forecasters, your future selves) and their subjective probabilities, we have a dilemma: either they are constrained by something external to *them* or they are not. In the former case, the reference class problem looms, for presumably something else is playing the role of ‘expert’ for *them*—some theory, frequencies, symmetries, logical probabilities, or what have you. In the latter case it is dubious whether they have earned their title as ‘experts’; we would be left with a no-theory theory of expertise.

The two problems just discussed—that of conflicting experts, and that of inheriting the reference class problem from your experts—are closely related. In a sense, the reference class problem for non-radical subjectivism just *is* the problem of conflicting experts. When a proposition is classified one way, or relativized to one background assumption, one of your experts—relative frequency information, chance, your weather forecaster, your future self, logical probability—assigns it one probability; when it is typed another way, or relativized to another background assumption, the same expert assigns it another probability. A single expert is conflicted with itself. Or looked at another way, any given expert fissions into *many* experts, one for each reference class. For each way of classifying ‘tomorrow’, we have a relative frequency expert, a chance expert, and so on. The problem of conflicting experts is worse than we might have thought, because we have so *many* of them. And this means that the reference class problem for the non-radical subjectivist is worse than we might have thought.

3 Solving or dissolving the reference class problem?

3.1 The ratio analysis of conditional probability

We have found the reference class problem bobbing up in important versions of every major interpretation of probability. One might conclude that all of the interpretations are somehow incomplete: that they each need to be supplemented with a further theory about what are the ‘right’ reference classes on which probability statements should be based. Yet I believe that the prospects for such theories (e.g., in terms of ‘narrowest classes for which reliable statistics can be compiled’, or ‘total evidence’) are dim. And those interpretations that appear to escape the reference class problem do so by being no-theory theories. I conclude that it is seemingly inescapable among theories that make substantive claims about what probabilities are and how they should be determined—that might be genuine *guides to life*.

A previous time-slice of mine (2003a) argued that we should stop trying to escape the reference class problem. Where we seek unconditional, single-case probabilities we keep finding conditional probabilities instead. I saw a hint there to be taken. For there is a sense in which the relativity to a reference class is not really a *problem* for these interpretations at all, any more than the relativity of simultaneity is a *problem* for time. If the reference class problem is a problem for anything, I argued, it is for Kolmogorov’s treatment of conditional probability.

As I noted in the introduction, the orthodox Kolmogorov theory identifies conditional probability with a ratio of unconditional probabilities:

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

Thus, all probabilities are unconditional probabilities or ratios thereof, and conditional probabilities are undefined whenever their antecedents have zero unconditional probability. Let us call the identification of conditional probabilities with ratios of unconditional probabilities the *ratio analysis* of conditional probability. I have long been campaigning against the ratio analysis. In the next section, I briefly rehearse some of the main kinds of arguments discussed in my (2003b), before returning in the subsequent sections to the reference class problem and its bearing on the ratio analysis.

3.2 Problems with the ratio analysis

Conditional probabilities are undefined whenever their antecedents have zero unconditional probability. It is for good reason that (RATIO) has its proviso ($P(B) > 0$). Now, perhaps the proviso strikes you as innocuous. To be sure, we could reasonably dismiss probability zero conditions as ‘don’t cares’ if we could be assured that all probability functions of any interest are *regular*—that is, they assign zero probability only to logical impossibilities.

Unfortunately, this is not the case. As probability textbooks repeatedly drum into their readers, probability zero events need not be impossible and can be of real interest. Yet some of these very textbooks seem to forget these examples when they adhere to the ratio analysis of conditional probability. Indeed, interesting cases of conditional probabilities with probability zero antecedents are manifold. Consider the perfectly random selection of a real number from the unit interval: the probability that either the point $\frac{1}{4}$ or the point $\frac{3}{4}$ is selected is zero; still, the conditional probability that $\frac{1}{4}$ is selected, *given* that either ($\frac{1}{4}$) or ($\frac{3}{4}$) is selected, is surely ($\frac{1}{2}$). The ratio formula for conditional probability thus cannot deliver the intuitively correct answer. Obviously there are uncountably many problem cases of this form.

Kolmogorov was well aware of the problem, and he went on to provide a refinement of the ratio analysis in terms 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. The difficulties that probability zero antecedents pose for the ratio formula for conditional probability are familiar to the mathematics and statistics communities; but they deserve more attention from philosophers. In any case, less familiar, I submit, are the difficulties posed by *vague* probabilities. Many Bayesians relax the requirement that probabilities are single real numbers, allowing them to be intervals or sets of such numbers. For example, your probability that there is intelligent life elsewhere in our galaxy need not be a sharp number such as 0.7049570000 . . . , but might instead be ‘0.7-ish’, represented by a suitable set of numbers around 0.7. Yet even then, various corresponding conditional probabilities can be sharp—for example, the probability that there is such life, *given* that there is such life, is clearly 1, and more interestingly, the probability that this fair coin lands heads, *given* that there is such life, is 1/2.

The problem of vague unconditional probabilities is bad enough, but the problem of *undefined* unconditional probabilities is worse (and it is trouble even for Kolmogorov’s refinement of the ratio analysis). Conditional probabilities of the form ‘ $P(A, \text{ given } B)$ ’ can be defined even when $P(A \cap B)$ and $P(B)$ are undefined, and hence their ratio is undefined. Here is an urn with 90 red balls and 10 white balls, well mixed. What is the probability that Joe draws a red ball, given that Joe draws a ball at random from the urn? 0.9, of course. According to the ratio analysis, it is:

$$\frac{P(\text{Joe draws a red ball} \cap \text{Joe draws a ball at random from the urn})}{P(\text{Joe draw a ball at random from the urn})}$$

Neither the numerator nor the denominator is defined. For example, there is no fact of the matter of the probability that Joe draws a ball at random from the urn. *Who is this Joe, anyway?* None of that matters, however, to the conditional probability, which is well-defined (and obvious). By analogy, we can determine that the argument:

Joe is a liar
Therefore,
Joe is a liar

is valid, even though there is no fact of the matter of the truth value of the statement ‘Joe is a liar’.¹⁴

I conclude that it is time to rethink the foundations of probability. It is time to question the hegemony of Kolmogorov’s axiomatization, and in particular the conceptual priority it gives to unconditional probability. It is time to follow heterodox probability theorists such as Popper (1959b), Renyi (1970), Spohn (1986), and Roeper and Leblanc (1999), to take conditional probability as the fundamental notion in probability theory, and to axiomatize it directly.

3.3 The ratio analysis and the reference class problem

In my (2003a) I argued that the ubiquity of the reference class problem provides more reason for taking conditional probabilities as primitive. I still think so, although now with some qualifications, for I believe that an important version of the reference class problem will still remain.

Let us return to the sense in which the reference class “problem” is not really a *problem* for the various interpretations of probability at all. Rather, it reveals something important about the fundamental nature of probability: it is essentially a two-place notion. All probability statements of any interest are at least tacitly, and often explicitly, relativized. So rather than try to solve the reference class problem, I proposed that we *dissolve* it: accept the fact that probabilities are essentially reference class-dependent, and honor that fact by taking conditional probabilities as basic. I have argued that conditional probability is the proper primitive of probability theory, and that it should be axiomatized directly. The ubiquity of the reference class problem only drives home the essential relativity of probability assignments.

For we have now seen more examples of conditional probabilities that cannot be identified with ratios of unconditional probabilities, because the required unconditional probabilities simply don’t exist. Various frequentists could tell us the conditional probability that John Smith will live to age 61, *given* that he is a consumptive Englishman aged 50. But they could not identify this with the ratio:

$$\frac{P(\text{John Smith will live to 61} \cap \text{John Smith is a consumptive Englishman aged 50})}{P(\text{John Smith is a consumptive Englishman aged 50})}$$

Neither term in the ratio is defined, but let us focus on the denominator. According to frequentism, this is another relative frequency. But there’s the rub: another *relative* frequency. The reference class problem strikes again! *Qua* one way of classifying John

¹⁴ In my (2003a, b), I canvas—and reject—various proposals for rescuing the ratio analysis in the face of such counterexamples.

Smith, we get one relative frequency for his being a consumptive Englishman aged 50; *qua* another way of classifying him, we get another relative frequency. Indeed, in a universe with infinitely many things—quasars, quarks, space-time points all included in their number—almost all of which are *not* consumptive Englishmen aged 50, this relative frequency could be as small as 0. That's hardly comforting for the ratio analysis! But in any case, we only get further relative frequencies—or as I would prefer to put it, we only get further conditional probabilities. Moreover, the conditions, in turn, have various relative frequencies, but yet again, relative to still further reference classes. And so the regress goes. The process never 'bottoms out' with unconditional probabilities. To paraphrase an old joke, it's conditional probabilities all the way down.

And so it goes for the other interpretations as well. One can assign classical *conditional* probabilities *given* a specification of a set of equipossibilities; but one cannot assign a classical unconditional probability to this *being* the set of equipossibilities. One can assign a logical *conditional* probability to the next emerald being green, *given* the evidence of 10 observed green emeralds; but one cannot assign a logical unconditional probability to this evidence. This coin may have a *conditional* propensity of landing heads, *given* a specification of an experimental set-up or what not, but there is no unconditional propensity for this experimental set-up or what not itself. And various subjectivists could assign various *conditional* probabilities, *given* corresponding expert assignments; but they could only assign unconditional probabilities to these assignments themselves by becoming no-theory theorists.

3.4 The metaphysical and the epistemological reference class problems

I still believe that we should take conditional probabilities as primitive. The formal treatment of probability should mirror the fact that any probability that serves as a guide to life is at base conditional. But I now think that this dissolves only the *metaphysical* reference class problem, while the *epistemological* problem remains.

First, the good news. Our search for privileged reference classes that would ground absolute probabilities of the form $P(X)$ was bound to prove futile. We should accept that there are only *relative* probabilities out there—probabilities *conditional* on this condition or that ($P(X, \text{given } A)$, $P(X, \text{given } B)$, etc). The ubiquity of the reference class problem is a reminder that probability assignments are, by their very nature, always *relativized*. By analogy, our search for a privileged reference frame—that of the ether—that would ground absolute distances or absolute temporal intervals, was bound to prove futile. We should accept that there are only *relative* distances or temporal intervals out there—distances and temporal intervals *relative* to this reference frame or that. At base, probability assignments must be relativized to a reference class. I call this the essential *relativity* of probability theory. It dissolves, I contend, the metaphysical reference class problem.

Now, the bad news. Giving primacy to conditional probabilities does not so much rid us the epistemological reference class problem as give us another way of stating it.¹⁵ Which of the many conditional probabilities should guide us, should underpin our inductive reasonings and decisions? Our friend John Smith is still pondering his prospects of living at least 11 more years as he contemplates buying life insurance. It will not help him much to tell him of the many conditional probabilities that apply

¹⁵ Here I am grateful to very incisive remarks by Jacob Rosenthal, which led to substantial revisions of this section.

to him, each relativized to a different reference class: “conditional on your being an Englishman, your probability of living to 60 is x ; conditional on your being consumptive, it is y ; . . .”. (By analogy, when John Smith is pondering how far away is London, it will not help him much to tell him of the many distances that there are, each relative to a different reference frame.) If probability is to serve as a guide to life, it should in principle be possible to designate one of these conditional probabilities as the *right* one. To be sure, we could single out one conditional probability among them, and insist that *that* is the one that should guide him. But that is tantamount to singling out one reference class of the many to which he belongs, and claiming that we have solved the original reference class problem. Life, unfortunately, is not that easy—and neither is our guide to life.

Still, it's better to have one problem than two. I will leave it to others to judge the extent to which I have succeeded in ridding us of the metaphysical reference class problem. But I am aware that I have not solved the epistemological problem. I invite you to join me in the search for a solution for the interpretations of probability that have a genuine claim to being guides to life. After all, whichever interpretation you favor, the epistemological version of the reference class problem is your problem too.

Acknowledgements I thank for very helpful comments and discussions: Mark Colyvan, Adam Elga, Branden Fitelson, Peter Godfrey-Smith, Matthias Hild, Chris Hitchcock, Ralph Miles, Jim Woodward, and audiences at the Australasian Association of Philosophy conference in Hobart, the Singapore Management University, and the Konstanz Summer School on Philosophy and Probability. I am especially grateful to Jacob Rosenthal, who gave me penetrating comments on two earlier drafts. Thanks also to an anonymous referee for *Synthese* for comments on an earlier draft, and to Stephan Hartmann for his help *qua* editor of this volume.

References

- Ayer, A. J. (1963). Two notes on probability. In A. J. Ayer (Ed.), *The concept of a person and other essays* (pp. 188–208). London: MacMillan.
- Bertrand, J. (1889). *Calcul des Probabilités* (1st ed.). Gauthier-Villars.
- Carnap, R. (1950). *Logical foundations of probability*. University of Chicago Press.
- Carnap, R. (1963). Replies and systematic expositions. In P. A. Schilpp (Ed.), *The philosophy of Rudolf Carnap* (pp. 966–998). LaSalle, IL: Open Court.
- Church, A. (1940). On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46, 130–135.
- de Finetti, B. (1937). Foresight: Its logical laws, its subjective sources, translated in Kyburg and Smokler (1964).
- Fetzer, J. (1977). Reichenbach, reference classes, and single case ‘Probabilities’. *Synthese*, 34, 185–217; *Errata*, 37, 113–114.
- Fetzer, J. (1982). Probabilistic explanations. *PSA*, 2, 194–207.
- Fine, T. (1973). *Theories of probability*. Academic Press.
- Frieden, B. R. (1991). *Probability, statistical optics, and data testing*. Springer-Verlag.
- Gaifman, H. (1988). A theory of higher order probabilities. In B. Skyrms, & W. L. Harper (Eds.), *Causation, chance, and credence*. Kluwer.
- Giere, R. N. (1973). Objective single-case probabilities and the foundations of statistics. In P. Suppes et al. (Eds.), *Logic, methodology and philosophy of science IV* (pp. 467–483). North Holland.
- Gillies, D. (2000). Varieties of propensity. *British Journal of Philosophy of Science*, 51, 807–835.
- Goldstein, M. (1983). The prevision of a prevision. *Journal of the American Statistical Association*, 77, 817–819.
- Hacking, I. (1965). *Logic of statistical inference*. Cambridge University Press.
- Hájek, A. (1997). ‘Mises Redux’—Redux: Fifteen arguments against finite frequentism. In *Erkenntnis*, Vol. 45 (pp. 209–227). Reprinted in *Probability, dynamics and causality — Essays in Honor of R. C. Jeffrey, D. Costantini, & M. Galavotti* (Eds.), Kluwer.

- Hájek, A. (2003a). Conditional probability is the very guide of life. In H. Kyburg Jr., & M. Thalos (Eds.), *Probability is the very guide of life: The philosophical uses of chance* (pp. 183–203). Open Court. Abridged version in *Proceedings of the International Society for Bayesian Analysis* 2002.
- Hájek, A. (2003b). What conditional probability could not be. *Synthese*, 137(3), 273–323.
- Hild, M. (in preparation a). *An introduction to induction*.
- Hild, M. (in preparation b). *Introduction to the concept of probability: A reader*.
- Jeffrey, R. (1992). *Probability and the art of judgment*. Cambridge University Press.
- Jeffreys, H. (1939). *Theory of probability*. Reprinted in Oxford Classics in the Physical Sciences series, Oxford University Press, 1998.
- Johnson, W. E. (1932). Probability: The deductive and inductive problems. *Mind*, 49, 409–423.
- Keynes, J. M. (1921). *Treatise on probability*. London. Macmillan. Reprinted 1962, New York: Harper and Row.
- Laplace, P. S. (1814). *Essai Philosophique sur les Probabilités*. Paris. Translated into English as *A philosophical essay on probabilities*, New York, 1952.
- Levi, I. (1990). Chance. *Philosophical Topics*, 18(2), 117–149.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In *Philosophical papers Volume II*. Oxford University Press.
- Mellor, D. H. (1971). *The matter of chance*. Cambridge: Cambridge University Press.
- Miller, D. W. (1994). *Critical rationalism: A restatement and defence*. La Salle, Illinois: Open Court.
- Miller, D. W. (1996). Propensities and indeterminism. In A. O'Hear (Ed.), *Karl Popper: Philosophy and problems* (pp. 121–147). Cambridge University Press.
- Peirce, C. S. (1867): *Review of Venn* (1866), reprinted in *Writings of Charles S. Peirce*, Vol. 2, pp. 98–102.
- Popper, K. (1959a). The propensity interpretation of probability. *British Journal of Philosophy of Science*, 10, 25–42.
- Popper, K. (1959b). *The logic of scientific discovery*. Basic Books.
- Reichenbach, H. (1949). *The theory of probability*. University of California Press.
- Rényi, A. (1970). *Foundations of probability*. Holden-Day, Inc.
- Roeper, P., & Leblanc, H. (1999). *Probability theory and probability semantics*. Toronto Studies in Philosophy.
- Sober, E. (2000). *Philosophy of biology* (2nd ed.). Westview Press.
- Spohn, W. (1986). The representation of Popper measures. *Topoi*, 5, 69–74.
- Strevens, M. (1998). Inferring probabilities from symmetries. *Noûs*, 32(2), 231–246.
- Strevens, M. (2003). *Bigger than chaos*. Harvard University Press.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81, 235–256.
- van Fraassen, B. (1989). *Laws and symmetry*. Oxford: Clarendon Press.
- van Fraassen, B. (1995). Belief and the problem of Ulysses and the Sirens. *Philosophical Studies*, 77, 7–37.
- Venn, J. (1876). *The logic of chance* (2nd ed.). Macmillan and Co; originally published 1866.
- von Mises, R. (1957). *Probability, statistics and truth*, revised English edition, New York.