

Probability as a Guide to Life

Author(s): Helen Beebee and David Papineau

Source: *The Journal of Philosophy*, Vol. 94, No. 5 (May, 1997), pp. 217-243

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2564581>

Accessed: 19-07-2016 12:29 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://about.jstor.org/terms>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

THE JOURNAL OF PHILOSOPHY

VOLUME XCIV, NO. 5, MAY 1997

PROBABILITY AS A GUIDE TO LIFE*

Let us assume that people's degrees of belief conform to the probability calculus. This is an idealization, but it will not matter for our purposes here. In line with normal usage, we shall refer to these degrees of belief as *subjective probabilities*. Then we can say that agents are *subjectively rational* if they choose those actions which their subjective probabilities imply are most likely to achieve what they want. More precisely, given that agents normally attach different utilities to a number of different ends (another useful idealization), subjectively rational agents will choose those actions which will maximize their subjective expected utility.

An example will be useful for what follows. Suppose Alma offers you her £2 to your £1, with you to win all the money if a coin lands heads. Should you accept or decline the bet? If you have no aversion to betting, and your utility function varies linearly with money for these small amounts, then subjective rationality dictates that you should accept her bet if your subjective probability for heads is greater than $1/3$ (for example, if you think the coin fair), and you should decline if it is less than $1/3$.

Note that in this example the possible outcomes (heads/tails) are causally and probabilistically independent of the choice of actions (accept/decline). We shall concentrate on examples like this throughout, in order to bypass the debate between causal and evidential decision theory. A fully general account of subjective rationality needs to specify whether decisions should depend on the degrees to which an agent believes that certain actions will *cause* cer-

* We would like to thank Dorothy Edgington, Peter Milne, and especially Scott Sturgeon for comments on earlier versions of this paper.

tain outcomes, rather than on the agent's conditional degrees of belief in those outcomes given those actions. This is an important issue, but it cuts across our present concerns: the points we make will apply equally to the degrees of belief invoked by causal and by evidential decision theories. (In the examples on which we shall focus, where outcomes are independent of actions, both theories agree that choices should depend on agents' simple degrees of belief in outcomes.)

I. OBJECTIVE CORRECTNESS IN DEGREES OF BELIEF

Whatever view is taken on the causal-evidential issue, there is clearly little point in being subjectively rational if your degrees of belief are objectively inaccurate.¹ Consider Betty, who has normal nonprobabilistic beliefs about standard coins, yet believes to degree 0.9 that such coins will land heads when tossed. Betty will put up large sums of money against small ones on any ordinary coin's coming down heads. We take it that this is a bad thing for Betty to do. Yet Betty is still subjectively rational, since she chooses those actions which maximize her subjective expected utility.

Betty is subjectively rational. Moreover, she need not have any false beliefs, since believing to degree 0.9 that the coin will land heads is not the same as (fully) believing that in some objective sense its probability of heads is 0.9, and we can suppose Betty does not believe this. Nevertheless, there is something objectively unsatisfactory about Betty's degrees of belief. They will lead her to act in a way that is subjectively rational yet objectively inadvisable. We shall investigate objective desiderata on subjective probabilities, and consider what can be said about which degrees of belief are objectively *correct* as a basis for decision. After exploring various possible principles governing objective desiderata on degrees of belief, we shall eventually conclude that the correct degree of belief in an outcome is not necessarily equal to its single-case probability, but rather to its *relative probability*, by which we mean the *objective probability* of that outcome *relative* to the features of the situation which the agent knows about.

Let us stress immediately that we are concerned here with the correctness of degrees of belief from a *prudential* point of view. Our question is: Which are the right degrees of belief to act on, given that you want such-and-such results? Other philosophers start with a more epistemological question: Which degrees of belief ought you

¹ Cf. D.H. Mellor, "Chance and Degrees of Belief," in R. McLaughlin, ed., *What? Where? When? Why?* (Dordrecht: Reidel, 1982), pp. 49-68.

to adopt, given that you have such-and-such information that bears on the occurrence of the outcome?² We are interested here in the first, prudential question. Moreover, we take this question to be different from the evidential question. Degrees of belief are made prudentially correct by objective features of the world, and are correct in this sense even for agents who are in no epistemological position to adopt those degrees of belief. Of course, only those agents who are fully informed about the relevant objective features of the world can *ensure* they have the correct degrees of belief. But, even so, you do not need to have such knowledge in order for the relevant notion of correctness to apply to you. It would be correct, in our sense, for you to give up smoking, if you want to avoid cancer, even if you have never heard about the smoking-cancer link.

This last point marks an important difference between our approach and that of the tradition of writers on “direct inference,” such as Hans Reichenbach, H. E. Kyburg, Isaac Levi,³ and J. L. Pollock (*op. cit.*). These writers, like us, address the question of which degrees of belief agents ought to adopt toward outcomes. They are in agreement with our eventual conclusion, at least to the extent that they hold that agents ought to set their degrees of belief equal to the relevant relative probabilities when they are in an epistemological position to do so. It is not clear, however, that writers in this tradition share our central concern with prudential rationality. Writers on direct inference tend simply to take it for granted that you ought to set your degrees of belief equal to the relative probabilities if you can. Their real focus is rather on the evidential issues that arise when you cannot do this, because you do not know the requisite objective probabilities, but only have incomplete evidence about them. By contrast, according to our sense of “correctness,” the degrees of belief of an agent who has no knowledge of the relevant objective probabilities are still constrained by those objective probabilities, in the sense that the correct degrees of belief for the agent to act on are still those which match the objective probabilities. Evidential issues that arise for agents who lack full knowledge of the relevant objective probabilities are important, but they lie beyond the scope of this paper. Our concern is to show that certain objective features of the

² For this distinction, cf. J.L. Pollock, *Nomic Probability and the Foundations of Induction* (New York: Oxford, 1990), pp. 22-23.

³ See, for example, Reichenbach, *The Theory of Probability* (Berkeley: California UP, 1949); Kyburg, *The Logical Foundations of Statistical Inference* (Dordrecht: Reidel, 1974); Levi, “Direct Inference,” in this JOURNAL, LXXIV, 1 (January 1977): 5-29.

world—namely, relative probabilities—make it prudentially correct to adopt corresponding degrees of belief. How far agents can achieve this, and what they should do if their lack of information hampers them, are further questions on which we shall comment only in passing.

II. THE PERFECT PRINCIPLE

Our aim, then, is to specify what makes certain degrees of belief correct from a prudential point of view. The first idea we want briefly to consider is the principle that the correct degrees of belief are those which match the truth. According to this principle, if outcome o_k is going to occur, then the right degree of belief in o_k is 1, and if o_k is not going to occur, the right degree of belief in o_k is 0. Thus, we would have:

- (1) The *perfect principle*: if the k^{th} of mutually exclusive outcomes $\{o_i\}$ is going to occur, then the correct subjective probability for o_i is 1 for $i=k$, and 0 otherwise.

We take it, however, that this is not a satisfactory account of which degrees of belief are objectively correct from a prudential point of view. If determinism is false, it is not in general fixed at the time of decision what the outcome will be. So even an agent who is omniscient about the laws of nature and current state of the universe would not be able to implement the perfect principle. In many decision situations, then, the prudentially correct degrees of belief must be fixed by the objective probabilities of outcomes rather than by what is actually going to happen.

III. THE OBJECTIVE PRINCIPLE

The obvious alternative to the perfect principle is to say that the correct degrees of belief are those which match the objective probabilities:

- (2) The *objective principle*: the correct subjective probabilities for outcomes $\{o_i\}$ are those which match the objective probabilities of those outcomes.

Note that agents who have the right degrees of belief in the sense of principle (2) and are also subjectively rational⁴ will act in such a way as to maximize not just subjective expected utility, but also objective expected utility.

⁴ We shall take this qualification about subjective rationality as read henceforth. Note that this warrants our using phrases like 'setting your degrees of belief equal to such-and-such objective probabilities' interchangeably with 'acting on such-and-such objective probabilities'.

The obvious question about the objective principle is how to understand its reference to "objective probabilities." This will be our focus here. We shall argue that the best interpretation of the objective principle requires us to read 'objective probability', not in the most obvious sense of *single-case probability*, but rather as *probability relative to a description* of the decision situation.

IV. SINGLE-CASE PROBABILITIES

The nature of single-case probabilities will be discussed further in what follows. For the moment, let us simply think of them as the rock-bottom metaphysical probabilities fixed by all probabilistic laws and all particular facts about specific situations. Quantum events provide the paradigm. The 0.5 probability of a radium atom decaying by A.D. 3615 is a single-case probability, in that there are no further contemporary facts about the radium atom or its surroundings that imply that the probability of its decaying by A.D. 3615 is different from 0.5.

Note that, if determinism were true, then single-case probabilities would never be different from 0 or 1. Nondegenerate single-case probabilities require outcomes whose occurrence or nonoccurrence is not fixed by present circumstances and laws.

Note also that such nondegenerate single-case probabilities change over time. This happens when circumstances relevant to the outcome, but not themselves previously determined, either occur or fail to occur. Thus, the single-case probability of a given radium atom decaying by a particular time will decrease as time passes and the atom does not decay, but will suddenly jump to one if the atom does decay. To take a more everyday example, due to David Lewis,⁵ imagine people who move through a multipath, multi-exit maze, deciding at each branch point whether to turn left or right by some quantum-mechanically random procedure. When they start, their reaching a given exit on the other side of the maze will have a given single-case probability different from 0 or 1. But this single-case probability will increase or decrease, depending on which choices they actually make as they pass through the maze, until it is determined where they will come out, by which time the single-case probability will have become 1 or 0.

V. THE SINGLE-CASE PRINCIPLE

One obvious way to read 'objective probabilities' in principle (2) is as the single-case probabilities at the time of the agent's decision. If you are betting on where someone will come out of the maze, do you

⁵ "A Subjectivist's Guide to Objective Chance," in R.C. Jeffrey, ed., *Studies in Inductive Logic and Probability, Volume II* (Berkeley: California UP, 1980).

not ideally want to know the *current* single-case probability of their emerging from each of the various exits? This suggests the following principle:

- (3) The *single-case principle*: the correct subjective probabilities for an agent to attach to outcomes $\{o_i\}$ are equal to the single-case probabilities of the $\{o_i\}$ at the time of the agent's decision.

Note that this single-case principle is equivalent to the perfect principle in all those cases where there is any practical possibility of humans using the perfect principle. For we can only now know whether some future outcome is going to occur if that outcome is already determined; and in such cases, the single-case probability of that outcome will already be 0 or 1.

It may seem obvious that, if there is a right and a wrong about which degrees of belief you ought to act on, then it must be as specified by the single-case principle. For the only way to attach a unique number to the objective probability of an outcome at the time of some decision is to equate it with the single-case probability at that time. So it seems to follow that acting on the single-case probability is the only way to guarantee that you will always choose the action that will maximize your objective expected utility.⁶

Nevertheless, despite this plausible line of reasoning, we think that the single-case principle is not the right way to link degrees of belief with objective probabilities. We think there are many situations where the objectively correct degrees of belief do not match single-case probabilities. We shall formulate a different principle, which we shall call the *relative principle*, which will cover these cases as well. This more general principle will assert that, in general, the correct degrees of belief do not match single-case probabilities, but rather a kind of probability we shall call *relative probability* (because it is relative to a description of the set-up facing the agent).

The relation between our preferred principle, the relative principle, and the single-case principle is analogous to that between the single-case principle and the perfect principle. We earlier rejected the perfect principle in favor of the single-case principle on the

⁶ Of course, degrees of belief other than those which match single-case probabilities can lead to the right choice in *particular* decision situations. The point of the single-case principle is rather that it ensures maximization of single-case expected utility in *every* decision situation. (Analogously, in nonprobabilistic contexts, while a choice of means recommended as effective by false beliefs will *sometimes* lead to desired results, those so recommended by true beliefs will *always* do so. Cf. Papineau, *Philosophical Naturalism* (Cambridge: Blackwell, 1993), ch. 3.)

grounds that the perfect principle does not cover all cases of prudentially correct choices. We reject the single-case principle for the same reason. There are prudentially correct choices beyond those recommended by the single-case principle.

In what follows, we shall first illustrate the relative principle with examples and explain how it works. We hope these examples will help make the relative principle plausible. But our ultimate rationale for favoring this principle over the single-case principle will be that it turns out to yield a simpler overall theory of prudential rationality. By the time we reach the end here, we shall be able to show that the relative principle yields a uniform and integrated account of prudential rationality, while the single-case principle does not.

Although we oppose the single-case principle, we suspect that it captures many people's intuitive understanding of prudentially correct degrees of belief. Even if this suspicion is right, however, it is difficult to find any straightforward expression of the single-case principle in the philosophical literature. Those authors who come closest to arguing that the right degrees of belief must match the single-case probabilities do not generally adopt the simple version of the single-case principle given above. Instead, they defend more complicated formulations, designed in part to block some of the difficulties we are about to raise. Thus, D. H. Mellor, Wesley Salmon,⁷ and Lewis (*op. cit.*) all forge a central connection between degrees of belief and single-case probabilities, but in each case via something more complicated than the simple single-case principle given above. We shall comment on these writers in due course. But for the moment, we shall retain the single-case principle as our target. This is because our disagreement is not on matters of detail, but rather on the whole idea that the objectively correct degrees of belief are those which match the single-case probabilities. It will be easier to make our objections clear if we start with a simple version of this idea.

VI. RELATIVE PROBABILITIES

It is high time we explained what we mean by the relative principle.

- (4) The *relative principle*: the correct subjective probabilities for an agent to attach to outcomes $\{o_i\}$ are the probabilities of the $\{o_i\}$ relative to the agent's knowledge of the set-up.

To illustrate this principle, and the associated notion of relative probability, consider this example. Alma tosses a fair coin and places

⁷ Mellor, "Chance and Degrees of Belief," and *The Matter of Chance* (New York: Cambridge, 1971); and Salmon, "Dynamic Rationality: Propensity, Probability and Credence," in J. Fetzer, ed., *Probability and Causality* (Dordrecht: Reidel, 1988).

her hand over it as it lands. We shall say that the relative probability of finding heads when Alma's hand is removed, *in relation to the description just given of this situation*, is 0.5. Note that this probability of 0.5 is not a single-case probability, since the single-case probability of heads is already either 0 or 1. Yet it is clearly the right probability to act on if you have not seen under Alma's hand, in line with the relative principle. For example, it is the right probability for working out that you should accept the bet we earlier described Alma as offering, since this 0.5 probability implies that the bet offers an expected gain of 50p.

It is not essential to the point we are making that the relevant result already has occurred. Supposed Alma tosses the coin and then offers her bet when it is in the air. Let us assume, as seems plausible, that the result, though still in the future, is by now determined. So the single-case probability of heads is already 1 or 0. Yet the relative probability of 0.5 is still clearly the probability on which you ought to act.

Nor is it essential to the point that the single-case probability of the relevant outcome already be 1 or 0. Suppose that Alma has a machine that makes one of two kinds of devices at random. On each run of the machine, a radioactive source decides with probability 0.5 whether the machine will make device *A* or device *B*. Each device will then print out 'heads' or 'tails', again on a quantum-mechanical basis. Device *A* gives 'heads' a single-case probability of 0.8. Device *B* gives 'heads' a single-case probability of 0.2. You know that the device in front of you is made by the machine, but you do not know whether it is an *A* or a *B*. Again, it seems obvious that you ought to decide whether to accept Alma's bet on the basis of the relative probability of 0.5 for 'heads', even though the single-case probability is already 0.8 or 0.2.

VII. PROBABILISTIC LAWS

It will be helpful to be more explicit about relative probabilities. For this we need the notion of a *probabilistic* law. We shall take the following as the basic form of such a law: 'The probability of an *A* being a *B* is *p*' ('The probability of a tossed coin concealed under a hand turning out to be heads is 0.5'). We take it that there are many objective truths of this form. Moreover, we take it that many of these are known to be true, on the basis of inferences from statistical data about the proportions of *Bs* observed in classes of *As*.

Given a particular situation that satisfies description *D* and a particular exemplification *o* of type *O*, the probability of *o* relative to *D* is *p* if and only if there is a probabilistic law 'The probability of a *D* being an *O* is *p*'. Note that these relative probabilities are like single-

case probabilities in that they attach to a particular outcome (a particular exemplification of a type) rather than to the type itself. But they do not attach to it simpliciter, but rather to the pair of the outcome and some description of the decision situation.

Single-case probabilities can be thought of as a special case of relative probabilities, namely, as the relative probabilities fixed by *all* the relevant features of the situation. If we think of single-case probabilities in this way, then the single-case probability at t of o is p if and only if, for some D applying to the relevant situation at t , it is a law that the probability of a D being an O is p , and D is maximal in the sense that, for any further E applying at t , the probability of a $D\&E$ being O is p as well. (Note how the maximality of such a D explains why D will exhaust all the *relevant* features of a particular situation; for if D is maximal, then any further E s will not alter the relative probability of o .)

Note that the relative principle concurs with the single-case principle in those cases where agents have information that fixes the single-case probabilities. This follows from the fact that single-case probabilities are special cases of relative probabilities. Consider agents who know all the relevant features of the particular decision situation. The relative probabilities of outcomes fixed by these agents' knowledge of the set-up will coincide with their single-case probabilities, and so these agents will conform to the single-case principle by conforming to the relative principle. The above examples make it clear, however, that this is not the only kind of case where the relative principle specifies the correct degrees of belief on which to act. Often agents' relative probabilities will not be single-case probabilities, yet agents should act on these probabilities in just the same way.

The relative principle accommodates any number of sensible choices that the single-case principle excludes. When informed people make real-life choices (whether to stop smoking, say, or to take some medication, or not to draw to an inside straight), they characteristically act in line with the relative principle but not the single-case principle. However detailed their knowledge of their specific situations, there are always further unknown features (about their metabolism, say, or the next card in the pack) that mean their relative probabilities are not single-case probabilities. But this is no reason to condemn their actions as prudentially incorrect.

There is more to be said about the choice between the relative and single-case principles, and we shall return to this in sections x-xiii below. But first we would like to clarify some other points.

VIII. THE RANGE OF THE RELATIVE PRINCIPLE

Some readers might be tempted to raise the following objection. It seems implausible to suppose that there is a probabilistic law of the form ‘the probability of a D being an O is p ’ for *every* property D that the agent might know the set-up in question to have. But where there is no relevant probabilistic law to be had, the relative principle remains silent about how the agent’s degree of belief in the outcome is objectively constrained.

Suppose, for instance, that you are presented with a black box that you are told makes coins and asked for your degree of belief that the next coin it produces will land heads when tossed. You know nothing further about the box. You just know it is a black box that manufactures coins. We take it that there is no true probabilistic law stating that the probability of heads on a coin from any coin-manufacturing black box is p . The category of coin-manufacturing black boxes is too open-ended and heterogeneous to sustain any serious general patterns. So, given our explanation of relative probabilities in the last section, there is no objective probability of heads relative to your knowledge of the set-up, which means that the relative principle specifies no prudentially correct degree of belief in this case.

Does this not show that the relative principle does not cover all decision situations, and hence that it does not say everything there is to be said about how degrees of belief in outcomes are objectively constrained? We say “no.” We are happy simply to concede that not all degrees of belief are so constrained. If the features of the set-up the agent knows about do not constitute the antecedent of any probabilistic law about the outcome in question, then there is no objective constraint on the agent’s degree of belief in that outcome. In the above example, there is no prudentially correct degree of belief in heads on the next coin from the black box.

IX. PHILOSOPHICAL INTERPRETATIONS OF PROBABILITY

Most relative probabilities are not single-case probabilities. Does this mean that the notion of relative probability presupposes the frequency interpretation of probability? If so, the notion of relative probability is in trouble, for the frequency interpretation of probability is no good.* We do not think this is a good reason for rejecting relative probabilities. The topics discussed here, and, in particular, the notions of relative probability and nonmaximal probabilistic law,

* For a quick account of the central flaw in the frequency interpretation of probability, see Papineau, “Probabilities and the Many-Minds Interpretation of Quantum Mechanics,” *Analysis*, LV (1995): 239-46.

are independent of disputes among alternative philosophical interpretations of probability. It is true that, if you are interested in philosophical interpretations of probability, then the frequency interpretation, which identifies the probability of an *A*'s being a *B* with the limiting frequency of *As* in the sequence of *Bs*, offers the most immediate explanation of relative probabilities and nonmaximal probabilistic laws. It is also true that the currently most popular alternatives to the frequency interpretation, namely, *chance* or *propensity* theories that take the notion of single-case probability as primitive,⁹ have to do some work to account for relative probabilities and nonmaximal probabilistic laws. But we certainly do not think that the notion of relative probability stands or falls (and so falls) with the frequency interpretation.

On the contrary, it seems to us a boundary condition on *any* satisfactory interpretation of probability that it make good sense of relative probabilities and nonmaximal probabilistic laws. Note in this connection that standard statistical tests are indifferent to the maximality or otherwise of the probabilistic laws under investigation. There are well-established empirical methods for ascertaining the probability that a certain kind of machine part, say, will fail in circumstances of some kind. It does not matter to these methods whether or not there are further probabilistically relevant differences between parts and circumstances of those kinds.

Perhaps this is a good reason for preferring those versions of chance or propensity theories which, in effect, take our notion of relative probability as primitive rather than single-case probability.¹⁰ But we shall not pursue this topic further here. We shall simply take the notion of relative probability and probabilistic law as given, and leave the philosophical interpretations of probability to look after themselves.

X. IN DEFENSE OF THE RELATIVE PRINCIPLE

We claim that the relative principle is the fundamental prudential truth about the correct degrees of belief for agents to act on. It may strike some readers as odd to accord the relative principle this authority. After all, many agents will have relative probabilities that dif-

⁹ For example, Lewis; R. Giere, "Objective Single-case Probabilities and the Foundations of Statistics," in P. Suppes et alia, eds., *Logic, Methodology and Philosophy of Science* (Amsterdam: North-Holland, 1973); Fetzer, "Reichenbach, Reference Classes, and Single-case Probabilities," *Synthese*, II (1977): 185-217.

¹⁰ Cf. I. Hacking, *Logic of Statistical Inference* (New York: Cambridge, 1965); Levi, *Gambling with Truth* (New York: Knopf, 1967); D. Gillies, "Popper's Contribution to the Philosophy of Probability," in A. O'Hear, ed., *Karl Popper* (New York: Cambridge, 1995).

fer from the single-case probabilities, and, as a result, the relative principle may specify that they should choose actions other than those which maximize the single-case expected utility.¹¹

For example, imagine that you accept Alma's £2 to your £1 that 'heads' will be printed out by a device from the machine that makes devices *A* and *B*. You do so on the grounds that your probability for 'heads' relative to your knowledge of the situation is 0.5. But, in fact, this is a device *B*, for which the single-case probability of 'heads' is already 0.2. So declining Alma's bet would have a higher *single-case* expected utility than accepting it (0p instead of -40p).

Here, the relative principle seems to advise the choice that is objectively inferior. Why then do we say it is superior to the single-case principle? On the face of things, it looks as if the relative principle will only give the right advice in those special cases where it coincides with the single-case principle.

Despite this, we maintain that the relative principle says all we need to say about correct degrees of belief. Even in cases where your relative probabilities do not coincide with single-case probabilities, it is objectively correct for you to set your degrees of belief equal to the relevant relative probabilities, and to act accordingly.

Our attention can be distracted here by a point we shall later discuss at some length. There is a sense in which it would be prudentially better if you could *find out* about the single-case probabilities before you choose between the alternative actions on offer. But this does not show that there is anything incorrect about setting your degrees of belief equal to merely relative probabilities when you *have not* found out those single-case probabilities.

Indeed, as we shall see, the explanation of *why* it is prudentially better to find out about the single-case probabilities, if you can, is itself a consequence of the correctness of making choices on the basis of relative probabilities. We shall argue (in sections XIX-XX, following Frank Ramsey and I. J. Good) that the best understanding of why it is *better* to act on the single-case probabilities, if you can, is simply that it is correct, in the sense of the relative principle, to find out the single-case probabilities before acting, if you have the choice. It is certainly true that it is prudentially better to find out the single-case probabilities before deciding, if you can. But this truth, far from

¹¹ When we speak of "an agent's relative probability" for some outcome, this should be understood as shorthand for the relative probability of that outcome in relation to *D*, where *D* is the agent's knowledge of the set-up in question.

challenging the relative principle's notion of correctness, depends on it.

XI. A USEFUL COMPARISON

We shall return to these complexities in due course. At this stage, however, some readers may still need persuading that there is *anything* worthwhile about the relative principle when it diverges from the single-case principle. How can it be correct to act on relative probabilities when they advise actions different from those which would be advised by the single-case probabilities, as in the case where you accept Alma's bet not knowing the device in front of you is, in fact, a *B*?

The following comparison may be helpful. First, suppose that, as above, you are offered Alma's bet on a device you know to be from the machine that makes *As* and *Bs*, but where you do not know which sort your device is. Now, compare yourself with Dorinda, say, who is offered the same bet, but in connection with a device that she knows will print out 'heads' with a quantum-mechanical single-case probability of 0.5.

You and Dorinda will both accept Alma's bet, for you and Dorinda will both act on a probability of 0.5 for 'heads'. But where Dorinda is acting on a single-case probability, you are acting on a relative probability instead. Now, would you rather be offered Dorinda's bet than your bet? No. People who make your bet will make money just as fast as people who make Dorinda's bet. Even though Dorinda's bet is in line with the single-case probabilities, and yours is not, yours is objectively just as good a bet. Both of you are quite rightly wagering £1 because your expected repayment is £1.50. (How *can* yours be a good bet, given that it is possible you are betting on a device *B*, which would mean your *single-case* expected repayment is only 60p? But note that we could equally well ask: How *can* Dorinda's bet can be a good bet, given that it is possible her device will in fact print out 'tails', which would mean her repayment is going to be 0p?)

We take this to show that it is objectively worthwhile to conform to the relative principle, even in cases where it could lead you to choose an action that will not, in fact, maximize the single-case expected utility. When you accept Alma's bet, knowing only the relative probabilities but not whether you are betting on a *B* or an *A*, you are making just as good a choice as Dorinda, who knows the single-case probabilities but not what the result will actually be. So, if Dorinda is making a good choice, from a prudential point of view, then so are you.

XII. "A GOOD BET"

It will be helpful to expand briefly at this point on a notion that we assumed in the last section—the notion of an objectively “good bet,” a bet that is worth accepting instead of declining. This notion encapsulates the idea of the objectively correct choice which we are trying to explicate.

What makes a bet good? An initial temptation is to say that what makes a bet good is that it *will* make you money (a temptation to which we succumbed in the last paragraph, in the interests of dramatic emphasis, when we said that “People who make your bet will make money just as fast as people who make Dorinda’s bet”). But, of course, we cannot strictly say this, since even a good bet may turn out unluckily (your device may print out ‘tails’). Another temptation is to say that a good bet will make money *in the long run*. But this is little better, since any finitely long run of bets can still turn out unluckily.¹²

Once we resist these temptations, we realize that the only thing that can be said about an objectively good bet is that it is objectively *probable* it will make you money; more generally, that the objectively right choice is that which offers the maximum *objective expected utility* among the alternatives open to you. It may be surprising that there is nothing more to be said, but it is a conclusion that a long tradition of philosophers addressing this question have been unable to avoid.¹³

What is even more surprising, perhaps, is that it does not matter to the objective correctness of a choice whether the objective probabilities that make it correct are single-case or not. All that matters is that they are the objective probabilities relative to your knowledge of the decision situation. The comparison in the last section makes this point clear. It does not matter to the goodness of the bet offered by Alma whether the single-case probability of heads is currently different from 0.5. As long as you know the decision situation as one in which the relative probability of heads is 0.5, it will be right for you to accept the bet; and you will be just as likely to make money, in the objective sense of probability that matters to choice, as someone who

¹² If you believe in the frequency theory of probability, you can argue that a good bet will definitely make money in the *infinite* long run. But, even for frequentists, it is not clear why this should *make* a bet good, given that none of us is going to bet forever.

¹³ Cf. C.S. Peirce, “The Doctrine of Chance,” in M.R. Cohen, ed., *Chance, Love, and Logic* (New York: Harcourt, 1924); and Hilary Putnam, *The Many Faces of Realism* (La Salle: Open Court, 1987).

accepts a similarly structured bet on the basis of knowledge that the single-case probability is 0.5.¹⁴

XIII. DECISION SITUATIONS

One way of conveying our picture of prudential rationality would be to say that the practical predicament facing an agent cannot be separated from the agent's knowledge of his surroundings. Let us use the admittedly inelegant term *decision situation* for the context of a choice. It is tempting to think of such a decision situation as consisting solely of the objective set-up that will generate the outcomes of interest—such as, for example, some physical device that will print 'heads' or 'tails' with certain probabilities.

If we think of decision situations in this way, then it is natural to conclude that, if there are correct degrees of belief, they must be those which correspond to the current single-case probabilities. For there is no way of identifying a unique probability for an outcome, given only the set-up in all its particularity, except as the probability fixed by all current features of the set-up—that is, the current single-case probabilities.

But suppose we think of the decision situation facing an agent as consisting of a set-up *plus* the properties the agent knows the set-up to have. Now, there is another way of identifying a unique probability for an outcome—namely, the outcome's probability relative to all those properties which constitute the decision situation.

Note that on this view it is still an objective feature of the world which makes an agent's degree of belief correct. For it is a perfectly

¹⁴ Is there not a danger, when betting on merely relative probabilities, that your opponent may know the single-case probabilities, and so have the advantage of you? (Alma may have a secret way of telling if the device is an *A*.) But note that in such betting situations you will have extra information apart from knowing about the structure of the betting device—you will also know, for example, that Alma is offering 2-1 against heads. So the recommendation of the relative principle in such a case depends crucially on the objective probability of heads *relative to the fact that someone is offering these odds*. If this is still 0.5, then your bet is fine. But if it is significantly less than 0.5, as would be the case if people like Alma have a way of fleecing suckers, then you should avoid such bets.

We would all do well to remember the advice given by Sky Masterson's father in Damon Runyon's "The Idyll of Miss Sarah Brown" (in *Runyon a la Carte* (London: Constable, 1946)): "Son, no matter how far you travel, or how smart you get, always remember this: Some day, somewhere, a guy is going to come to you and show you a nice brand-new deck of cards on which the seal is never broken, and this guy is going to offer to bet you that the jack of spades will jump out of this deck and squirt cider in your ear. But, son, do not bet him, for as sure as you do you are going to get an ear full of cider." In what follows, however, we shall simplify the exposition by continuing to assume, contra Mr. Masterson, Sr., and common sense, that the existence of attractive bets is in general probabilistically irrelevant to the relevant outcomes.

objective fact that the probability of outcome o relative to some property D of the set-up is p . Many such facts obtain quite independently of what agents know, or even of whether there are agents at all. Of course, *which* such fact makes a particular agent's degree of belief correct will depend on that agent's decision situation. But it is scarcely surprising that which objective probabilities matter to decisions should be different for different decision situations.

XIV. THE PRINCIPAL PRINCIPLE

It is an implication of Lewis's "principal principle" (*op. cit.*) that, if you know the single-case probability of some outcome and nothing else relevant (nothing "inadmissible" in Lewis's terminology), then you ought to set your degree of belief equal to that single-case probability. It will be illuminating to compare the position we are defending here with Lewis's.

One immediate point is that Lewis seems as much concerned with issues of evidential rationality as with prudential rationality. His principal principle is part of a more general Bayesian account of the degrees of belief required of rational thinkers who have such-and-such evidence, and bears on questions of the prudential correctness of degrees of belief, if at all, only derivatively.

Even so, Lewis holds that connection forged by the principal principle between degrees of belief and single-case probabilities exhausts the ways in which degrees of belief are rationally constrained by knowledge of objective probabilities. So it is natural to ask what Lewis would say about agents who know relative probabilities but not single-case probabilities. Take the case of the machine that makes devices A and B . As before, you know that the device in front of you comes from this machine but not whether it is A or B . So you know that the probability of 'heads' relative to your information is 0.5, but you are ignorant of the single-case probability of this outcome. Lewis's principal principle would seem to have no grip here. Yet, surely, evidential rationality, just as much as prudential rationality, requires you to have a degree of belief of 0.5 in heads.¹⁵

Lewis's response to this challenge¹⁶ hinges on the fact that the principal principle, unlike the single-case principle, does not refer specifically to *current* single-case probabilities. Rather, it says that, if you know the value of the single-case probability for any time t , and

¹⁵ Cf. Levi, Review of R.C. Jeffrey, ed., *Studies in Inductive Logic and Probability, Volume II* [*op. cit.*], in *Philosophical Review*, xcii (1983): 120-21; and Kyburg, "Principle Investigation," this JOURNAL, lxxviii, 12 (December 1981): 772-78.

¹⁶ See Lewis, "Postscripts to 'A Subjectivist's Guide to Objective Chance,'" in his *Philosophical Papers, Volume II* (New York: Oxford, 1986), pp. 114-32.

nothing else relevant, then you ought to have a matching degree of belief, even at times other than t . In the case at hand, Lewis would argue that you know the single-case probability of heads was once 0.5 (before the machine opted between A and B), and you know nothing else relevant (since your later information tells you nothing about which device the machine made). So the principal principle requires you to set your degree of belief equal to the relative probability of 0.5 after all.

XV. A MODIFIED SINGLE-CASE PRINCIPLE?

The apparent success of Lewis's response might suggest that we were perhaps too quick to replace the single-case principle by the relative principle. Perhaps we should have tried:

- (5) The *modified single-case principle*: the correct subjective probabilities for an agent to attach to outcomes $\{o_i\}$ are equal to the single-case probabilities of the $\{o_i\}$ at that earlier time when nothing probabilistically relevant to o was yet determined, apart from those features which the agent now knows about.

This would not require that the correct degrees of belief are the *current* single-case probabilities. Indeed, this principle would agree with the relative principle, to the extent that it specifies that the correct degrees of belief are a function, inter alia, of which facts the agent now knows about. But it would at least retain the original idea of equating prudentially correct degrees of belief with single-case rather than relative probabilities.

This modified single-case principle will not do, however. Consider once more the case where you do not know whether the machine has given you device A or B . But now suppose that you have observed your device print out 'heads' as its first two results. What is the correct degree of belief in 'heads' next time? A simple calculation using Bayes's theorem will show that the probability of 'heads' next time on a device that has already produced two 'heads' is $13/17$. This probabilistic law is a consequence of the other probabilistic laws already specified for this example. So the relative principle specifies, quite rightly, that the correct degree of belief for 'heads' on the next toss is around $13/17$. Somebody who uses this figure to decide which bets to accept will make the correct choices.

But note that $13/17$ is not, *nor ever has been*, the single-case probability of 'heads' next time. The single-case probability of heads on any toss was 0.5 before the device was made, and is now either 0.2 or 0.8. So even the modified single-case principle cannot cover this case. The only fully general principle governing prudentially correct degrees of belief is that they should equal the relative probabilities.

We should stress that this example does not refute Lewis's view. For Lewis is not committed to the modified single-case principle as the only principle governing reasonable degrees of belief. Rather, Lewis's full theory of reasonable degrees of belief contains two principles that together accommodate cases that escape the modified single-case principle: these are (a) the principal principle by which he links reasonable degrees of belief to single-case probabilities, and (b) a principle of Bayesian conditionalization for updating reasonable degrees of belief. So he can deal with the present example by saying that rational agents will first set their degrees of belief for *A* and *B* equal to 0.5, in line with the principal principle; and then they will update these degrees of belief by Bayesian conditionalization to get the right answer after they observe that their device has produced two 'heads' already.

Still, apart from demonstrating that prudentially correct degrees of belief do not always match single-case probabilities (and thus that the modified single-case principle on its own is inadequate), our example also shows that there is a measure of theoretical economy to be gained by upholding the relative principle rather than Lewis's principal-principle-plus-Bayesian-conditionalization. The advantage is that we do not need to adopt conditionalization as a separate principle, at least in cases structured by objective probabilistic laws. For it can be shown that, in such cases, conditionalization falls out of the relative principle, in the way illustrated by our example: if you set your new degree of belief equal to the relative probability of getting the outcome in situations satisfying your initial-knowledge-plus-new-information, then you will automatically be conditionalizing. By contrast, note that conformity to conditionalization does not follow from the principal principle, but must be added as an independent requirement on rationality, since updated degrees of belief often will not match single-case probabilities, as in our example. It seems to us preferable to avoid postulating Bayesian conditionalization as an extra requirement, given the well-known difficulties of providing this principle with a satisfactory rationale.

XVI. LEWIS AND DETERMINISM

We would like to offer a different and more direct reason for preferring our relative principle to Lewis's package of principal-principle-plus-Bayesian-conditionalization. Consider Elspeth, a convinced determinist, who is asked to bet on a certain coin coming down heads. Elspeth has performed a large number of trials on coins of this manufacture and has overwhelming statistical evidence for a probabilistic law that says that the probability of heads with coins just

like this one is 0.9. Suppose also that determinism really is true. What ought Elspeth's degree of belief in the coin's landing heads on the next throw to be?¹⁷

Since Elspeth thinks that heads has a single-case chance of either 0 or 1, the full version of Lewis's principal principle implies that her degree of belief in heads should be the average of these two chances weighted by her degree of belief in each, which will, of course, equal her degree of belief in the latter chance. But what ought this degree of belief to be? To us it seems clear that it should equal 0.9. But there is nothing in Lewis's principal principle, nor in his principle of Bayesian conditionalization, to say why. The relative principle, on the other hand, can deal with the case: the objective probability of heads relative to Elspeth's knowledge of the situation is 0.9, so 0.9 is the correct degree of belief for Elspeth to have.

There are various responses open to Lewis. An initial possibility would simply be to deny that there would be any probabilistic laws if determinism were true; but this seems to us a thesis that needs some independent argument in its favor. An alternative is to appeal to the notion of "counterfeit chance"—a degree of credence that is resilient in the sense that it would not be undermined by further "feasible" investigation—as a surrogate for the notion of a probabilistic law under determinism.¹⁸

Counterfeit chances have an indeterminacy, however, that makes Lewis suspicious of their credentials as constraints on degrees of belief. He says:

Counterfeit chances will be relative to partitions; and relative, therefore, to standards of feasibility and naturalness; and therefore indeterminate unless the standards are somehow settled, or at least settled well enough that all remaining candidates for the partition will yield the same answers. Counterfeit chances are therefore not the sort of thing we would want to find in our fundamental physical theories, or even in our theories of radioactive decay, and the like. But they will do to serve the conversational needs of determinist gamblers.¹⁹

Put in our terms, Lewis is here appealing to the point that probabilistic laws will often imply different probabilities for the same outcome, depending on what nonmaximal descriptions of the set-up we

¹⁷ Cf. Levi, Review of Jeffrey, and "Chance," *Philosophical Topics*, xviii (1990): 117-49.

¹⁸ Cf. Bryan Skyrms, *Causal Necessity* (New Haven: Yale, 1980); and Lewis, "Postscripts to 'A Subjectivist's Guide to Objective Chance'."

¹⁹ Lewis, "Postscripts to 'A Subjectivist's Guide to Objective Chance'," p. 121.

employ. Lewis seems to be suggesting that such probabilities cannot therefore place objective constraints on agents' degrees of belief. But there is no good basis for this conclusion. It is only if you assume that such constraints must always be supplied by single-case probabilities that you will find their relativity to descriptions worrying. By contrast, this relativity is unproblematic from our point of view, since our tie between degrees of belief and objective probabilities is forged by the relative principle, not the principal principle, and this explicitly specifies that correct degrees of belief are relative to agents' knowledge of the set-up.

Lewis gives no good reason to regard nonmaximal probabilistic laws as second-class citizens. We may not want to find such laws in our fundamental physical theories, but that does not mean that they are *only* capable of serving the conversational needs of determinist gamblers. On the contrary, the part that nonmaximal probabilistic laws play in the relative principle makes them fundamentally important in understanding how degrees of belief in outcomes are objectively constrained. To deny this is to deny that there is any objective feature of the world which makes it correct for Elspeth to attach a 0.9 degree of belief to heads.

XVII. SALMON'S VIEW

Salmon argues forcefully that subjective degrees of belief ought to match objective probabilities, in addition to satisfying internal constraints of coherence and conditionalization (*op. cit.*). Although Salmon thinks of objective probabilities as long-run frequencies, he is sensitive to the distinction between single-case and relative objective probabilities: for him, this is the distinction between frequencies in objectively homogeneous reference classes (genuine *chances*) and frequencies in merely epistemically homogeneous reference classes.

When Salmon addresses this distinction (*op. cit.*, pp. 23-24), he argues that the fundamental external constraint on degrees of belief is that they should match the chances (the single-case probabilities); that is, his basic commitment is to the single-case principle. He does recognize, however, that we need some account of the worth of choices informed by merely relative probabilities. His response is to view degrees of belief that match merely relative probabilities as the *best available estimates* of single-case probabilities. In effect, he views the relative principle as a matter of trying to conform to the single-case principle under conditions of limited information.

This does not work. Consider again the machine that makes *As* and *Bs*. You have a device from this machine, but you do not know which kind it is. Your relative probability for 'heads' is 0.5. This does

not seem a good estimate of the single-case probability—after all, the single-case probability is certainly either 0.2 or 0.8. At best, Salmon could argue that the relative probability is a weighted average of the possible single-case probabilities. But then the weighting factors (0.5 for the 0.2 single-case probability, and 0.5 for the 0.8) are relative probabilities once more, and we are left once more with an unexplained appeal to the worth of matching degrees of belief to relative probabilities. Whichever way we turn it, Salmon does not give us a way of explaining the relative principle in terms of the single-case principle.

XVIII. MELLOR'S VIEW

In *The Matter of Chance*, Mellor argues that the correct degrees of belief are those which match the chances, and to that extent is in prima facie agreement with our single-case principle (*op. cit.*). He then considers an example that is effectively equivalent to our machine that makes *As* and *Bs*. In connection with this example, he appeals to another thesis defended in his book, namely, that chances are functions not just of specific outcomes ('heads' next time) but also of the *kind of trial* generating the outcome. So, if the trial is (a) *the machine makes a device that displays a result*, then the chance of 'heads' is 0.5; but if it is (b) *this particular device (that came from the machine) displays a result*, then the chance of 'heads' is either 0.2 or 0.8.

So Mellor makes chances relative to the kind of trial. What determines the kind of trial? If this depends on the agent's *knowledge* of the set-up, in such a way that the trial is of kind (a) whenever the agent does not know whether the device is an *A* or a *B*, and is of kind (b) whenever the agent does know which it is, then Mellor has effectively switched to the relative principle, for his "chances" are equal to probabilities relative to the features of the set-up the agent knows about.

Mellor specifies, however, that his "chances" are not relative to knowledge, but rather to what "counts as repeating the bet" at hand: if the bet is on the *device-generating machine*, so to speak, the "chance" is 0.5, while if it is on *this device*, the "chance" is 0.2 or 0.8. Mellor's resulting position is certainly different from the relative principle, since repetitions of a bet can presumably be on *this device* rather than the machine, even though the agent does not know whether this device is an *A* or *B*. But for just this reason, Mellor's position is unsatisfactory, since he is left without any explanation of why the prudentially correct degree of belief for this kind of agent is 0.5.

In his more recent "Chances and Degrees of Belief," Mellor returns to this issue, and argues that such an agent should act on a 0.5

degree of belief after all, since of the two “chances” that attach to the outcome ‘heads’ (0.5 relative to trial (a); 0.2 or 0.8 relative to trial (b)), this is the only one that the agent knows (*op. cit.*). We agree that this is the right answer, but would observe that Mellor is only able to give it because he has switched to the relative principle. In our terms, he now holds that agents should act on the probabilities relative to the features of set-up that they know about, that is, that their degrees of belief should match their relative probabilities.

XIX. MORE INFORMATION IS BETTER

We have argued that, if your degrees of belief match your relative probabilities, it is irrelevant to the goodness of your consequent choice whether or not these are also single-case probabilities. Degrees of belief that correspond to the relative probabilities on one gambling device can be just as correct, from a prudential point of view, as numerically identical degrees of belief that match single-case probabilities on a different device.

As we admitted earlier, however, there is a different sense in which it *is* better to act on single-case probabilities, if you can find out about them. To pin down this sense, it will be helpful to consider, not two different gambling devices, but a single device, about which you might have more or less information. Imagine, for example, that you are facing a device from the machine that makes *As* and *Bs*. You are able *either* to act on your present 0.5 relative probability for heads on a device from this machine, *or* to find out whether your device is an *A* or a *B*, and so act on the single-case probability of heads. In this kind of comparison, where the single-case probability involves *more* information about the *same* device, we agree that it is prudentially better to find out the single-case probability before deciding.

We deny, however, that this point calls for any modification of, or addition to, the relative principle. That it is prudentially better to find out the single-case probabilities if you can, before committing yourself to action, is itself a consequence of the relative principle.

What is more, this corollary of the relative principle is simply a special case of a more general point, namely, that it is always prudentially better to act on probabilities relative to more rather than less information about a given set-up, if you can. Not only are single-case probabilities better than relative probabilities, but relative probabilities that are relative to greater amounts of information are better than relative probabilities that are relative to less. Indeed, the same point applies within the category of single-case probabilities. As long as you can keep your options open, it is prudentially better to wait and see how the single-case probabilities evolve, instead of deciding

straight off knowing only their initial values. Best of all, of course, is to wait until you know the outcome itself. If you can delay until Alma has removed her hand, then you can accept her bet whenever the coin shows heads, and decline whenever it shows tails.

XX. THE RAMSEY-GOOD RESULT

Let us first offer an example to illustrate why it is always prudentially better to act on probabilities relative to more rather than less information about a given set-up, if you can. We shall then sketch a general demonstration.

Imagine you are presented with three sealed boxes. You know that a ball has been placed in one, and that the other two are empty, and that some chance mechanism has given each box the same $1/3$ single-case probability of getting the ball. Felicity offers you her £3 to your £2, with you winning the total stake if the ball is in box 1. You can either accept or decline the bet now (option *G*, for “go”), or you can look in box 3 and then accept or decline the same bet (option *W*, for “wait-and-see”).

If you take option *G* and choose now, you will decline Felicity’s bet (since you would be putting up £2 for an expected return of £1.67). So, since you will decline, the expected gain of option *G* is £0.

If you take option *W*, then $1/3$ of the time you will find the ball in box 3, and so decline the bet, since you would then be sure to lose it. But $2/3$ of the time you will not find the ball in box 3, and in these cases you will accept the bet, since the expected return on your £2 will now be £2.50. Averaging over these two possibilities, weighted by their respective probabilities, you can expect to gain £0.33 if you take option *W*.

So option *W* is better than option *G*, since its expected value is £0.33 rather than £0. The reason is that option *W* gives you the opportunity of switching to accepting the bet in just the case where your extra information (no ball in box 3) indicates that this is worthwhile.

This example illustrates the general result. Suppose you are faced with alternative actions $\{A_i\}$, and have a metachoice between deciding now (option *G*), or finding out whether *C* before acting (option *W*). If this extra information can affect which A_i you choose, and if the cost of acquiring this information is negligible, then the expected gain from option *W* is always greater than the expected gain from option *G*.

We shall sketch a proof (following Ramsey and Good²⁰) for the simple case of two alternative actions A_1 and A_2 (accept or reject

²⁰ Ramsey, “Weight or the Value of Knowledge,” *British Journal for the Philosophy of Science*, xli (1990): 1-4; Good, “On the Principle of Total Evidence,” *British Journal for the Philosophy of Science*, xvii (1967): 319-21.

some bet, say). If you take option G , then you will now choose one of these— A_1 , say. Let $EU(A_1/C)$ be the expected utility of this action assuming C . Then we can decompose the expected gain of taking option G and so choosing A_1 as

$$(I) \quad EU(G) = EU(A_1/C)Pr(C) + EU(A_1/\text{not-}C)Pr(\text{not-}C)$$

Now, suppose you take option W instead and find out whether C . The discovery that C might make you choose A_2 rather than A_1 . Alternatively, the discovery that $\text{not-}C$ might make you do this. (Both discoveries cannot make you switch; otherwise, you would not have preferred A_1 to start with.)

Suppose, without loss of generality, that discovering C makes you switch to A_2 . Then

$$(II) \quad EU(W) = EU(A_2/C)Pr(C) + EU(A_1/\text{not-}C)Pr(\text{not-}C).$$

But this must be larger than $EU(G)$, because if $EU(A_2/C)$ was not bigger than $EU(A_1/C)$, you would not have chosen A_2 when you discovered C .

XXI. MORE INFORMATION IS RELATIVELY BETTER

We take this result to explain why it is better to find out the single-case probabilities before acting, when you can. More generally, we take it to explain why it is better to delay acting until you acquire more probabilistically relevant information, whenever you can.

It is noteworthy, if not immediately obvious, that this explanation makes essential appeal to the relative principle rather than to any single-case principle. Waiting-and-seeing is better than going because it has a higher objective expected utility from the point of view of your *relative probabilities prior to this choice*. In equations (I) and (II), $Pr(C)$ and $Pr(\text{not-}C)$ are not single-case probabilities, but probabilities relative to your initial information; and the same goes for the conditional probabilities $Pr(o_i/C)$ and $Pr(o_i/\text{not-}C)$ implicitly involved in the quantities $EU(A_i/C)$ and $EU(A_i/\text{not-}C)$.

We can illustrate the point with the example of the boxes. Here, $Pr(C) = 1/3$ is the probability of finding the ball in box 3. This is not the single-case probability of that outcome, for that is already either 0 or 1.²¹ Rather, it is the probability relative to your initial knowledge of the set-up. Similarly, the conditional probability = $1/2$ of the ball being in box 1 (call this outcome o), given it is not in box 3,

²¹ Was not $1/3$ *once* the single case probability of C ? Well, then consider an example where your relative probabilities of $1/3$ for each box is the result of conditionalizing on previous results, as in the kind of case discussed in section xv above.

which is central to the expected utility calculations, is a relative conditional probability. It is calculated by dividing $\Pr(o \ \& \ \text{not-}C) = 1/3$ by $\Pr(\text{not-}C) = 2/3$, where these are again all probabilities relative to your initial information. (The *single-case* conditional probability of o given $\text{not-}C$ —that is, $\text{Sing-CPr}(o \ \& \ \text{not-}C)/\text{Sing-CPr}(\text{not-}C)$ —will already either be 0 or 1 or undefined—depending on whether the ball is in fact in box 2, box 1, or box 3, respectively.)

In effect, the Ramsey-Good result reduces the *superiority* of decisions based on single-case probabilities, and better informed probabilities generally, to the *correctness* of waiting-and-seeing instead of going, whenever you have this choice. But, for this result to work, correctness must be understood in the sense of the relative principle, not the single-case principle.²²

XXII. FAILING TO MAXIMIZE SINGLE-CASE EXPECTED UTILITY

We can highlight the dependence of the Ramsey-Good result on the relative principle by noting that the sense in which it is always better for you to acquire additional probabilistically relevant information, if you can, does not guarantee that you will end up choosing the action A_i which maximizes the *single-case* expected utility. Consider the above example again, and imagine the ball is in fact in box 2. You look in box 3, see nothing, and so accept Alma's bet, since it is now a good bet. But it was determined from the start that this bet would lose. So, despite the fact that it is better, from the point of view of the probabilities on which you will quite correctly act, to wait-and-see rather than go, the single-case expected utility of waiting-and-seeing in this particular case is -£1, by comparison with the single-case expected utility of £0 for going.

This is simply a reflection at the metalevel of a point made in section X. The action that is specified as correct by the relative principle need not be the action that maximizes the single-case expected utility. Since the Ramsey-Good result explains the superiority of waiting-and-seeing in terms of the relative principle rather than any single-case principle, it is unsurprising that there should be cases where waiting-and-seeing does not maximize single-case expected utility.

²² David Miller complains that the Ramsey-Good result fails to explain the *principle of total evidence* (*Critical Rationalism* (Chicago: Open Court, 1994), chapter 7). Our view is this. We have no explanation of why you should act on the probabilities relative to *all* your information. That is the relative principle, and we take this to be a primitive truth about prudential rationality. But this principle then explains, via the Ramsey-Good result, why you should arrange to act on probabilities relative to *more* information, if you can.

XXIII. FINDING OUT THE SINGLE-CASE PROBABILITIES

Of course, in the special case where acquiring more information tells you about the current single-case probabilities, instead of just giving you better informed relative probabilities, then the A_i chosen after waiting-and-seeing will indeed maximize the single-case expected utility, since you will then be acting on the single-case probabilities themselves. (Imagine that you can look in *two* boxes, or find out whether your device is A or B , before deciding on Alma's bet.) This may make some readers feel that in cases where you find out the single-case probabilities, at least, the rationale for waiting-and-seeing is independent of the relative principle. Surely, here the justification for waiting-and-seeing is simply that this will enable you to act on the single-case probabilities.

This thought simply takes it for granted, however, that it is better to act on the single-case probabilities if you can, without offering any explanation for this. We think it better to explain things, if you have the materials to do so, rather than take them as primitive. The Ramsey-Good result explains why it is better to act on the single-case probabilities, rather than relative probabilities, when you have the choice. But it does so by assuming the correctness of acting on your initial relative probabilities rather than on your final single-case probabilities.

XXIV. CONCLUSION

It is natural to think that the prudentially correct action is always the action that maximizes single-case expected utility. This, in turn, makes it natural to suppose that the single-case principle specifies the basic notion of prudential correctness, and that the relative principle is at best some kind of derivative poor relation.

Our analysis indicates, however, that this has things the wrong way round. To see this, suppose you postulate the single-case principle as a basic principle governing prudential correctness. You will then need to say something about agents who act on relative probabilities. Since relative probabilities need not be (nor need ever have been) equal to single-case probabilities, you will need separately to introduce another notion of prudential correctness, as in the relative principle. The question will then arise of why it is better to be correct in the first single-case sense than in the second relative sense. You could take this to be a further primitive fact. But it would make more sense to explain it, via the Ramsey-Good result, as itself a consequence of the relative notion of correctness, since you will in any case need this explanation for the more general point that it is better to act on more probabilistically relevant information than less.

By this stage, however, the single-case principle will have ceased to do any work. For if we simply begin with the relative principle, which we have to assume anyway, we do not need separately to postulate a notion of single-case correctness, since agents who act on single-case probabilities are *per se* acting on relative probabilities. The relative principle itself will explain why it is preferable to act on single-case probabilities and, in general, on probabilities relative to more information, when you can.

It may seem unnatural to regard the relative principle as more basic than the single-case principle. But this position has the support of theoretical simplicity.

University College/London

HELEN BEEBEE

King's College/London

DAVID PAPINEAU