

# Objective Probability as a Guide to the World

Michael Strevens

*Philosophical Studies* 95, 243–75, 1999

## ABSTRACT

According to principles of *probability coordination*, such as Miller's Principle or Lewis's Principal Principle, you ought to set your subjective probability for an event equal to what you take to be the physical probability of the event. For example, you should expect events with a very high probability to occur and those with a very low probability not to occur. This paper examines the grounds of such principles. It is argued that any attempt to justify a principle of probability coordination encounters the same difficulties as attempts to justify induction. As a result, no justification can be found.

## ACKNOWLEDGMENTS

Special thanks to Barry Loewer and Bas van Fraassen for helpful discussion.

A dropped egg stays dropped. Or rather, it does so with very high objective probability, for statistical mechanics tells us there is a chance—a microscopic chance—that the energy dispersed by the impact and destruction of the egg will find its way back to the egg in just the right form that the egg will be reconstituted, and will leap back into the air to the point at which it was dropped.

This is one of those events that should not be expected to happen in a billion lifetimes of the universe; certainly, we expect *our* dropped eggs to break and to stay broken. Now, although it may not be obvious, this expectation of ours is based on an inference, an inference from the fact that there is very low *objective probability* of reconstitution, to the conclusion that we ought to have a very low *expectation* of reconstitution. The inference seems to be licensed by a very general rule telling us what to do with objective probabilities: set your levels of expectation, or subjective probabilities, equal to the objective probabilities. (This kind of inference is known as *direct inference*.)

Why is this rule of direct inference reasonable? Why is it rational to set our expectations in accordance with objective probabilities, and irrational to do otherwise? The most satisfying answer to this question would be a demonstration that the rule is correct. There may be other routes to the appreciation of a rule's reasonableness,<sup>1</sup> but proof from first principles is the royal road. In this paper I do my best to show that there is no such proof of the rule. I hope the radical nature of this position is clear: if I am right, then it is impossible to provide any non-circular argument that we should expect dropped eggs to remain dropped. The person who sits around waiting for the

---

1. I have in mind in particular the idea that a rule of inference is reasonable if it is one of the coherent set of rules that is fixed upon when reflection reaches equilibrium (Goodman 1983). It is not my purpose in this paper to launch a critique of the reflective equilibrium notion of justification (for such a critique, see Stich 1990). See also the comments in the conclusion to this paper.

egg to spring back into their hand is not crazy—at least, not demonstrably so—just different.

The implications are not confined to the kitchen. If quantum mechanics is probabilistic then objective probabilities are everywhere. Even if quantum mechanics is deterministic then objective probabilities are almost everywhere. For example, (a) any chemical reaction has its rate and direction determined by the probabilistic generalizations of statistical mechanics; (b) the central concept of evolutionary biology, fitness, is a measure of the probability that an organism will survive and proliferate; and (c) probabilistic generalization is the rule in the medical sciences: heavy smoking greatly increases your chances of dying from lung cancer, but it does not always kill you.

In short, our best theories of physical, chemical and biological phenomena are all probabilistic. They make no definite predictions. Rather, they predict events with, at best, extremely high objective probability. But this information is useless in itself. For the purposes of science and of everyday life, we need something that translates objective probabilities into rational expectations. This rule will take the relevant scientific probabilities and assure us, for example, that we can expect a chemical reaction to occur in a certain way, or that if we want to live long, happy lives, we really ought to consider giving up smoking. If no such rule is demonstrably rational, we are in serious trouble.

## 1. Preliminaries

### 1.1 *Rules For Probability Coordination*

Let us call the practice of setting subjective probabilities equal to objective probabilities the practice of *probability coordination*. Philosophers have not been able to agree on the exact form of the rule for probability coordination. A simple version of the rule was framed by David Miller (1966) (although

the rule can be traced back at least as far as Leibniz):<sup>2</sup>

One's subjective probability for an event  $A$ , on the supposition that the objective probability of  $A$  is  $x$ , ought to be  $x$ .

That is, it is a constraint on one's subjective probability function  $C(\cdot)$  that

$$C(A|T) = P_T(A)$$

(where  $T$  is a probabilistic theory—for example, statistical mechanics—that assigns  $A$  the objective probability  $P_T(A)$ ).<sup>3</sup> This rule is often called Miller's Principle.

As it stands, the rule is inadequate as a rule of direct inference. The initial difficulty is this: when setting one's subjective probabilities for an event, one wants to conditionalize on all one's knowledge, not just on knowledge of some theory  $T$ . But Miller's principle does not explicitly say what to do with background knowledge. This suggests the following amendment of the constraint on  $C(\cdot)$ :

$$C(A|TK) = P_T(A)$$

where  $K$  is any body of propositions (in the case of direct inference, one's background knowledge).

But now at least two new problems are introduced. First, our background knowledge may contain another theory that assigns a different probability to the event  $A$ .<sup>4</sup> Second, there are situations (perhaps rare, or even nomo-

---

2. Leibniz wrote "*quod facile est in re, id probabile est in mente*" (quoted in Hacking 1977, 128). It is hard to translate this sentence without prejudicing the issue, but it seems very likely that Leibniz was doing his best to articulate the principle under discussion (although as Hacking argues in the same work, Leibniz did not make quite the same distinction between objective and subjective probability as we do today).

3. I am being vague about whether probability functions range over events or propositions. In the formulation of the probability coordination rule, ' $A$ ' may be taken to stand for the proposition that the event  $A$  occurs.

4. A well known example: suppose we know that the probability that a randomly selected Texan is rich is very high, and that the probability that a randomly selected philosopher is rich is very low. How should we set our subjective probability that a randomly selected Texan philosopher is rich?

logically impossible) where we ought not to follow the above rule, namely, situations where we know more about an event than probability tells us. (Perhaps we have a crystal ball that tells us for sure whether or not the probabilistically determined event occurs; see Hall 1994.)

Clearly, if the above principle is to serve as a rule of direct inference, we must place some sort of constraint on the body of propositions  $K$  (so as to exclude, for example, information supplied by crystal balls). Thus the correct rule for setting subjective probabilities will have the following form:

$$C(A|TK) = P_T(A)$$

where  $K$  is any body of *admissible* propositions. This is a schema for rules relating subjective and objective probability; particular rules will follow from particular definitions of admissibility. The term ‘admissibility’ was introduced by David Lewis. He fleshes it out in a proprietary way in Lewis 1980 to produce his own rule for probability coordination, which he calls the “Principal Principle”.

There has been considerable work on the question of what is admissible. Isaac Levi argues that information is admissible only if it is stochastically irrelevant to the event  $A$  (that is, only if  $T$  states that the probability of  $A$  conditional on  $K$  is equal to the probability of  $A$ ). Henry Kyburg admits evidence provided it is not known that the evidence is stochastically relevant. (For the views of Kyburg and Lewis, see their summaries of their own work in Bogdan 1984.) David Lewis proposes a working definition of admissibility on which a proposition is admissible if either (a) it contains only information about objective probability, or (b) it contains only information about events that occur before the event  $A$  concerning which a subjective probability is to be set (Lewis 1980; Lewis notes that the definition is not fully general because it assumes that there are no “crystal balls”. For further thoughts see Lewis 1994 and Thau 1994). Lewis’s definition is predicated on his very particular

views about objective probability;<sup>5</sup> his intention seems to be to exclude any stochastically relevant information.<sup>6</sup>

I propose to avoid this difficult issue. All parties to the dispute agree that there are cases in which all evidence is admissible. These are the cases where all evidence is known to be stochastically irrelevant, that is, cases in which the relevant theory of probability  $T$  entails that  $P_T(A|K) = P_T(A)$  (in words, the objective probability of  $A$  conditional on  $K$  is equal to the objective probability of  $A$ ). Thus if *any* of the various rules of probability coordination can be demonstrated, then the following rule can be:

$$C(A|TK) = P_T(A)$$

if  $P_T(A|K) = P_T(A)$ .<sup>7</sup> My purpose, then, can be served by showing that there is no good argument for this narrow rule. It will follow that there is no good

---

5. In particular, the connection between Lewis's definition of admissibility and stochastic irrelevance depends on the claims (a) that objective probability cannot exist where there are underlying deterministic processes (this ensures that events occurring prior to the outcome of interest are stochastically irrelevant) and (b) that a good theory of objective probability assigns a single objective probability to any particular event (this rules out the sorts of cases described in note 4).

6. Lewis's view is in this respect similar to Levi's. Lewis, unlike Levi, thinks that all propositions can be assigned an objective probability, so that there is always a determinate fact of the matter about stochastic relevance.

7. Lewis (1980, 1994) raises the prospect that on certain philosophical views about probability (including frequency views),  $T$  itself may be inadmissible in some circumstances, that is, that  $P_T(A|TK) \neq P_T(A|K)$ . In such cases Lewis introduces a new probability coordination rule,  $C(A|TK) = P_T(A|TK)$ . (See also Hall 1994.) This "New Principle" appears to differ from the more intuitive rules of probability coordination described at the beginning of this paper. In my paper on the subject (Strevens 1995) I show that this is not so: the New Principle is entailed by an appropriately fleshed out old-fashioned, intuitive rule of probability coordination

Nevertheless, Lewis and I agree that on a frequency view of probability,  $P_T(A|T)$  is not in general equal to  $P_T(A)$ , and that in such circumstances, we should set  $C(A|TK)$  equal to  $P_T(A|T)$ , not  $P_T(A)$ . If such a view of probability is correct, the simple principle I try to justify throughout this paper (see below) ought not to be what I call  $\text{PCP}$ , i.e.,  $C(A) = P_T(A)$ . It ought to be  $C(A) = P_T(A|T)$ . This reformulation would not affect the arguments in this paper, partly because in the cases we are concerned with,  $P_T(A|T)$  and  $P_T(A)$  differ by a very small amount. See also note 12.

argument for a rule that advises probability coordination in a superset of the cases to which the narrow rule applies.

### 1.2 A Simple Probability Coordination Rule

I will focus on a more simple probability coordination rule still. I will call this rule the *simple probability coordination principle*, or PCP for short. It is as follows. We should respect the following constraint on our subjective probabilities:

$$(PCP) C(A) = P_T(A)$$

provided that (a) we know that  $T$  is the correct theory of probability, and (b) all of our background information is stochastically irrelevant. The PCP is a consequence of every probability coordination rule considered so far. (This for the following reason: in the special case in which  $T$  is known to be true and  $K$  is our background knowledge,  $C(TK) = 1$ , and thus  $C(A) = C(A|TK)$ .) The PCP applies to a subset of the situations to which the previous probability coordination rule applied, namely, those in which we know the values of objective probabilities for sure, and in which we wish to use these values to set a value for  $C(A)$  that takes into account all background knowledge.

Because all rules for probability coordination entail PCP, if there is no demonstration that PCP is correct, there can be no demonstration that any probability coordination rule is correct. I will show that there is no way to prove PCP.

There is a second reason for thinking that a failure to demonstrate the rationality of PCP indicates the impossibility of demonstrating the rationality of any rule of probability coordination. It is this: there is good reason to think that the reasonableness of more general probability coordination rules is derived in some way from the reasonableness of probability coordination in the cases covered by PCP. Consider: why is it rational to set subjective probabilities equal to what we *believe* to be the correct objective probabilities

(given the irrelevance of our background knowledge)? Presumably because we should act as though our beliefs are true. (The only obvious alternative explanation is that failing to coordinate probabilities is somehow internally inconsistent, but as I show later in the paper, this is not the case.) That is, we very much want to set our subjective probabilities equal to the actual objective probabilities, so, in the absence of certainty concerning the objective probabilities, we use our best guess as to their values.<sup>8</sup> If this is correct, then any justification of probability coordination (as a strategy for direct inference) must begin with a demonstration that it is reasonable to desire to set one's subjective probabilities equal to the actual objective probabilities.<sup>9</sup> It follows that if probability coordination is justified in any cases at all it must be justified in the cases covered by PCP.

I have given two reasons, then, that if any (known) rule of probability coordination can be shown to be rational, PCP can be shown to be rational. The second argument is illuminating but non-demonstrative; the first argument is decisive. In what follows I show that there is no argument that PCP is rational.

### *1.3 Subjective Probability*

To fully understand the implications of PCP, it is necessary to say something about subjective probability. I make three assumptions (none of which is intended as a definition of subjective probability).

1. To have a subjective probability for an event is to be in some kind of psychological state.

---

8. More precisely, we weight our various guesses according to our subjective probabilities that those guesses are correct.

9. Please note the 'desire' here; I am not suggesting that it is irrational not to set one's subjective probabilities to the actual objective probabilities. Otherwise we would all be irrational, for none of us know the values of all of the objective probabilities.

2. As I have been assuming all along, our subjective probability for some event  $A$  usually reflects our felt level of expectation concerning  $A$ 's occurrence. Thus, the higher our subjective probability for  $A$ , the greater our expectation that  $A$  will occur.
3. Entertaining any given subjective probability normally entails a certain disposition to action, as follows: to have a subjective probability of  $p$  for an event  $A$  is to be willing to accept odds of up to  $p : (1 - p)$  that  $A$  will (or did) in fact occur.

Now the (simple) probability coordination principle can be restated without the technical term 'subjective probability'. The principle tells us that we should expect events to occur to the degree that they are objectively probable, and that we are justified in gambling (i.e., taking any action in uncertain circumstances) at odds proportional to this expectation.

## 2. Attempts to Justify Probability Coordination

There have been a few philosophical attempts to justify probability coordination, such as those of Braithwaite (1966), Mellor (1971) and Howson and Urbach (1993). Van Fraassen (1989) and Earman (1992) comment on the difficulty of grounding probability coordination using the strong law of large numbers (for a discussion, see note 11 below). Charles Peirce's contribution is discussed in the appendix. There is not, however, much sustained discussion in the literature. It is one of the aims of this paper to provide a more or less systematic survey of the ways in which PCP might be vindicated, and of the problems that stand in the way.

Before I begin to examine the options for justification, I will introduce a major theme of this paper: that there is a resemblance between rules for probability coordination and the rule of enumerative induction, and a corresponding resemblance in the difficulties encountered in pursuit of their justification.

The rule of enumerative induction is the rule that gives us warrant to assume that some observed correlation—for example, between being a raven and being black—will continue to obtain in the future. Like all inferential rules, it connects two domains of facts. Premises drawn from one domain lead to conclusions about the other. In the case of induction, the domains are the past and the future. The problem with the rule of induction is that it is the *only* inferential connection between the past and future. So in justifying induction, we face an overwhelming logical problem: if the principle in question is the *only* bridge between two domains, on what foundations could this bridge possibly be built?

Efforts to solve the Humean problem have traditionally run into one of three problems, each arising because the rule of induction seems to be the only available rule connecting the domains:

1. Appeal to the self-evidence of the rule founders on the possibility that the rule is a mere psychological “habit” that has no logical basis,
2. Attempts to derive the rule tacitly invoke the rule itself, or
3. Attempts to derive the rule rely on a global empirical conjecture connecting the two domains that cannot itself be established without using the rule. (In the case of Humean induction, that conjecture is what is often called the principle of the uniformity of nature.)

Now observe that the correct rule for probability coordination (whatever it may be) also seems to be the only connection between two domains, in this case the domain of objective probabilities and the domain of future events. In what follows, we will see that attempts to justify PCP run into the same three problems listed above, for exactly the same reasons.

There are two approaches to the justification of PCP. We may try to establish the principle either by appealing to the fact that it is self-evident, or by deriving it from other facts that are themselves self-evident or well-confirmed. I will consider these possibilities in turn.

### 3. The Simple Probability Coordination Principle as Self-Evident Truth

A self-evident truth—a specimen of what might be called *immediate a priori* knowledge—should be obvious, and obvious because it is true. The simple probability coordination principle PCP certainly seems to satisfy the first of these criteria, obviousness. That is, it looks true to us, so much so that we are inclined to question the sincerity of anyone who denies it.

“Obviousness”, however, is in the first instance a psychological property, and since Hume, philosophers have had to take seriously the possibility that it might sometimes have a purely psychological (rather than a logical) explanation. That is, obviousness might be a product not of logical self-evidence—whatever that is—but of psychological (or perhaps evolutionary) “habit”. What is “obvious” may seem that way only because we rely on it so instinctively, so often, in our daily lives. This kind of obviousness is clearly no guarantee of truth.

The philosophical suspicion of the *a priori* nurtured by Hume has been, to some extent, vindicated by science. Even a proposition as apparently “self-evident” as that space is Euclidean has turned out to be suspect. It would be overzealous to call into question every *a priori* principle on these grounds, but I will make the following cautious observation: when it comes to the fundamental physical attributes of the world, what is “obvious” to us can be quite wrong. Unfortunately, the subject matter of PCP—objective probability—falls squarely into this class of fundamental physical attributes of the world. (Although PCP is a normative principle, its normative force must depend at least in part on the nature of objective probability.) We must, therefore, be very careful about the implications of the obviousness of PCP. The reliability of our intuitions concerning something as mysterious as probability is not obvious at all.

How can we restore faith in the *immediate a priori*? Ironically, what is needed is a *theory* of the *a priori*. Such a theory will enable us to see how, in some cases at least, the obviousness of a proposition stems from its truth, not

from our psychology. Equipped with this theory, we may hope to distinguish in advance the cases in which obviousness guarantees truth.

The history of philosophy is full of attempts to use such theories to distinguish the true a priori from what is merely psychologically salient. No theory of this kind is now universally accepted, but there is one account that retains some respectability: the view that a priori truths are analytic truths, that is, propositions that are true in virtue of their meaning.

Could PCP be an analytic truth? If it were, there would have to be some semantic connection between the notion of ‘rational subjective probability for *A*’ and that of ‘objective probability of *A*’. Such a connection could arise in a number of ways.

First, there could be a semantic connection between the notion of objective probability and that of subjective probability. This, however, is implausible. Objective probability is a mysterious aspect of the world, whereas subjective probability is something psychological. The only thing the two kinds of probability have in common is that they both satisfy the axioms of the probability calculus, a fact that cannot be used to establish the truth of any rule for probability coordination.

Second, it might be claimed (or stipulated) that it is simply constitutive of rationality that we ought to set our subjective probabilities equal to known objective probabilities. Given such a definition of rationality, though, it makes sense to ask the following question: if ‘rationality’ is just the name of some particular cognitive practice, defined in part by PCP, why be rational? How is the practice deemed ‘rationality’ better than other practices? Of course, this is just the original question rephrased: why follow PCP?

Finally, it might be thought (or stipulated) that it is constitutive of objective probability that PCP is true, on the grounds that the term ‘objective probability’ can refer only to something of which the correct rule for probability coordination is true. (For example, it might be claimed that PCP is part of our intuitive theory of probability, and that objective probability is that

thing that best satisfies the theory. David Lewis makes this suggestion (Lewis 1986 xv–xvi.)

I will avoid contemplating the nature of objective probability at this time. Let me just point out some drawbacks to the above proposal. First, in order to justify PCP, it is not enough simply to define objective probability as whatever makes probability coordination rational. In addition, it must be shown that there *is* something in this world with which it is rational to coordinate subjective probabilities. Second, having found such a thing, you must have the courage (or the recklessness) of your convictions. If it turns out that it is rational to coordinate subjective probabilities with certain short term frequencies (which in fact it is, as I will show later), then you have to accept—and happily!—that those short term frequencies are in fact objective probabilities. This position will make a probabilistic science pointless. No sooner will you have discovered the value of a probability than the value will change.

In conclusion, it is just not possible to tweak definitions so that PCP comes out as true. But if PCP is not true by definition, then its obviousness is suspicious. Certainly, it looms large in our psychology, but why should we trust what we find in our minds?

#### 4. Arguments for Probability Coordination

If PCP is not an immediate a priori truth, it must be justified by way of an argument. The arguments philosophers have suggested (e.g. Mellor, Howson & Urbach) tend to be consequentialist arguments.<sup>10</sup> Before I move on to these arguments, I will consider the following non-consequentialist argument for probability coordination, due to Richard Braithwaite (1966).

Suppose that a probabilistic experiment has  $n$  mutually exclusive and exhaustive outcomes, each of which has the same objective probability of

---

10. Technically, Howson and Urbach's argument is not consequentialist, but in spirit it is (see the discussion later in the main text).

occurring, namely  $1/n$ . Braithwaite notes that this assignment of probabilities is invariant under permutation (i.e., if you swap the chances of two outcomes, you do not change the assignment). It follows, Braithwaite argues, that one's subjective probability over these outcomes ought also to be invariant under permutation (and thus that the subjective probability for any outcome must be  $1/n$ , equal to the objective probability). Why must objective invariance under permutation be reflected in subjective invariance under permutation? Because, Braithwaite answers, to set one's subjective probabilities otherwise would be "irrational to the point of schizophrenia" (p. 273).

I suggest that this compulsion Braithwaite feels is generated by none other than PCP itself, which was to be proved, not presupposed as the only sane option. Why? The PCP requires that one's subjective probabilities reflect the objective probabilities, so if the objective probabilities are all the same (or equivalently, are "invariant under permutation"), then the subjective probabilities ought to be, too. Braithwaite's argument has much in common with that of Howson and Urbach, discussed below.

## 5. Consequentialist Justification of Probability Coordination

A consequentialist justification of PCP defends the principle on the grounds that it is good for us: if we adopt PCP, the story goes, we will get what we want. More exactly, if we follow PCP then we will adopt subjective probabilities that motivate actions that further our interests.

### 5.1 *The Game of Life*

I will now attempt to say more exactly what is meant by the proposition that using PCP as a rule of direct inference will further our interests. The furthering of our interests normally depends on many strategies in addition to direct inference. To consider the advantage of PCP in isolation, we need to consider an aspect of life, perhaps artificial, to which these other strategies are

not relevant. The obvious scenario is a gambling game in which the objective probabilities are already known. I will develop a particular version of this scenario, which I call simply the Game.

In the Game, there are two participants, someone like us who I will call the Player, and an entity we might call Mother Nature. Each instance of the Game (each game with a small 'g') is centered around a chance setup, such as a coin toss, which I call the *game setup*, and a designated outcome, say the coin's landing heads. The probability of the designated outcome's occurring is known to have a fixed value  $p$  for any game.

A game consists of a number of rounds, each organized around a trial on the game setup (e.g., a toss of the coin). Any particular round proceeds in this way. Mother Nature names the stake and the odds for the round (say, a stake of ten utility points and odds of three to two in favor of heads). The Player chooses whether to bet for or against the designated outcome. Mother Nature takes the other side. The two participants contribute to the stake according to the odds. (In the above example, the participant who bets on heads puts up three fifths of the stake (six utility points), the other participant two fifths (four points).) The coin is tossed, and the winner takes all.

There are some different versions of the Game that depend on whether or not the odds and/or the stake are varied from round to round of a particular game. I will focus on the kind of Game in which both odds and stake are constant. It turns out that variable odds games can be dealt with as easily as constant odds games, but that varying the stake introduces an entirely new kind of difficulty. In the main text of this paper I will describe difficulties enough; the complications of variable stakes games are discussed in the appendix.

The strategy advised by PCP in the Game is as follows. If the participant who bets on the designated outcome is required to put up a fraction  $x$  of the stake, then according to PCP, the Player should take that side of the bet if  $x$  is less than the probability of the designated outcome, and the other side of the

bet if  $x$  is greater than the probability.

## 5.2 *Success in Life*

The consequentialist wants to show that the Game strategy advised by PCP furthers our interests, and thus that the principle is a good strategy in any of our probabilistic contests with Nature, all of which, I assert, bear some abstract similarity to the Game. What does it mean for a strategy to further our interests?

One way of showing that PCP is a good strategy is show that, for most games we play, following PCP will result in a net gain. (Say that such a strategy is a *winning strategy*.) A weaker result would be to show that we have good reason to be confident that PCP will be a winning strategy.

A further distinction: a strategy may be winning either in the *long run* or in the *short run*. To be winning in the long run, a strategy must (usually) result in a net gain if it is played *long enough*. That is, for most games, there must exist some number of rounds such that, if the game has that many rounds or more, the strategy yields a net gain. (Such a strategy may lose for a while, but there must be some point in the future, however distant, after which the player is ahead and remains so. I call this the *profit point*.) To say that a strategy is winning in the short run is to say that the strategy provides a net gain in most games, however short.

I will begin by considering PCP as a long run strategy.

## 6. Probability Coordination as a Winning Long Run Strategy

### 6.1 *Failed Attempts At A Proof From First Principles*

How might it be shown that PCP advises a winning strategy in the long run? It cannot be shown. In any round of a game, either outcome is possible, so in any sequence of rounds any sequence of outcomes is possible. That is the nature of chance phenomena. With exceptionally bad luck, a player can lose

every round of every game they play, regardless of the strategy they use.<sup>11</sup> We can conclude, then, that there simply could not be a proof that PCP advises a winning strategy.

Note, however, that to lose so often *is* exceptionally unlucky. We do not expect this sort of thing to happen very often. So are we not justified in being highly confident, if not exactly certain, that PCP will advise a winning strategy?

There is something intuitively compelling about this reasoning, but it is compelling for the wrong reasons. This becomes obvious when the argument is restated more formally, as follows.

It is easy to show that for any game, PCP advises a strategy that has a very high *objective probability* of delivering a net gain in the long run. This is a consequence of two facts:

1. The PCP advises a winning strategy in games where the frequency of the occurrence of the designated outcome is close to the probability of that outcome. (See the appendix for a proof of this result, and a quantification of ‘close’.)
2. The probability of the frequency approximately matching the probability is very high. (This is a consequence of the weak law of large numbers.)

---

11. Some people—including Mellor (1971), on one reading (Salmon 1979)—have thought that the strong law of large numbers offers a way around this problem.

This is not correct. Although the strong law says that certain outcomes will occur with probability zero, this probability is compatible with those outcomes occurring. Consider an infinite number of coin tosses. There is a probability of zero that any particular sequence of outcomes will result, yet some such sequence must be produced.

Van Fraassen (1989) and Earman (1992) present a slightly different argument against justifying probability coordination using the strong law. They argue that, regardless of the above consideration about coin tosses, to take a probability of zero or one as entailing facts about frequencies must always involve an instance of probability coordination. Mellor’s later defense (1982) of his 1970 argument is vulnerable to this objection.

Then, the reasoning goes, we are entitled to convert this high objective probability of winning to a high level of confidence, that is, a high *subjective* probability.

The problem: to do so, we need to use some rule for probability coordination. And it is the status of probability coordination that is in doubt.

Let us take a step back. There is a reason that this attempt to justify probability coordination involves a tacit appeal to the very practice in question. The PCP, as I have already argued, establishes a connection between objective probabilities and future events. To justify such a connection, we would need to have some independent grasp of the relation between probabilities and events. But—and this is what made probability coordination so important to begin with—the only grasp we seem to have of this relation is encapsulated in PCP itself. There is simply no solid ground from which we can assess PCP's truth. As with Humean induction, so with probability coordination: we cannot conjure a connection between the past and the future, or between the probabilistic and the non-probabilistic, from mathematics and deductive logic alone.

## 6.2 *Foundations For A Long Run Match Between Frequency And Probability*

The conclusion of the last section suggests that the problem of justifying probability coordination is simply intractable. Not only are there no known ways to justify PCP, it seems that there could be no such ways. Despair, however, is inevitable only if we concede that the rule for probability coordination *is* the sole connection between probabilities and events. There might be another connection, more fundamental than PCP, that is able to act as a foundation for probability coordination.

An obvious possibility is what I will call the *long run frequency postulate*:

(LRFP) If the objective probability of a type  $E$  outcome is  $p$ , the long run frequency of type  $E$  outcomes will be equal to  $p$ .

If this postulate is true, then PCP will advise a winning strategy for one game at least: the game that comprises all actual trials on the chance setup to which the probability  $p$  is attached, for example, a game in which you bet on every coin toss there ever was or ever will be. Of course no one ever plays these games, so none of what follows bears on any actual, human application of PCP, but it is worth going ahead because the conclusions drawn will be useful when more realistic cases are considered below.

Assume, then, that LRFP provides grounds for adopting subjective probabilities as advised by PCP. The next question is whether LRFP has any foundation itself. There are two possible positive answers.

The first way to found LRFP is to make it true in virtue of the definition of objective probability; that is, to define probability as long run (limiting) frequency.<sup>12</sup> Howson and Urbach's (1993) argument, treated below, is along these lines.

The second possible foundation for LRFP is even simpler: it might just be that LRFP is true for most or all probabilities. If so, we would have a kind of externalist justification for PCP. Though we could not know it, or have any rational degree of confidence that it is true, PCP would in fact be a long run winning strategy. (This position may be compared with the "postulationist" approach to the problem of induction described by Bertrand Russell (1948) or to Hans Reichenbach's (1949) vindication of induction. There are many

---

12. A frequentist definition of probability is, *prima facie*, inconsistent with PCP, for reasons that surfaced in the last section. Given probabilistic independence, the law of large numbers implies that there is a non-zero objective probability that the relevant long run frequency fails to match the probability (at least when there are only a finite number of outcomes). The PCP requires us to convert this probability into a non-zero rational degree of belief, which is arguably an acknowledgment of the possibility that the frequency will fail to match the probability, a conclusion clearly at odds with frequentism. The frequentist can avoid this problem in two ways. The first is to deny probabilistic independence (at least when large numbers of trials are involved). The second way out is provided by the "New Principle" proposed by Ned Hall and David Lewis (see note 7). The New Principle directs the frequentist to adopt a subjective probability of one that the frequency will equal the probability.

variants on this kind of defense of induction; all are agreed that there is no incontrovertible proof that induction is warranted.)

The cost of this philosophical tactic is a dilution of our notion of justification. Up until this point I have been looking for an *argument* for probability coordination, something that would give us confidence that our coordination of probabilities is warranted. But perhaps this world is just so wild at heart and weird on top that we can never have such a justification for anything, like PCP, that truly matters.

Once again we are reminded of Hume's problem of induction. The resolution of that problem, too, seems to turn on the truth (or otherwise) of a spatiotemporally global postulate, the so-called principle of the uniformity of nature. In both cases, we could know that the postulate held or failed to hold only if we had surveyed the length of time and the breadth of the universe, that is, only when it would be too late anyway. We can speculate about these postulates; we can act, full of hope, as though they were true; but we cannot adopt any reasoned epistemic attitude towards them.

In short, either by going frequentist or by going externalist we can provide a foundation for probability coordination. But these justifications, such as they are, apply only to very special instances of the Game. This is a crippling qualification.

## 7. Probability Coordination as a Winning Short Run Strategy

### 7.1 *The Short Run Game*

We are mortal, and so we only ever play games of a limited number of trials, which I will call *short run* games. Even if we are convinced that PCP supplies a winning long run strategy, then, we have to ask whether we are justified in applying the same strategy to short run games.

Successful use of a winning long run strategy requires that we be willing and able to play as many games as it takes to reach the profit point, the point

at which we have made a net gain that, as long as we continue to play, can never entirely disappear. But no one who plays short run games can conform to this requirement, because a short run game may end before the profit point is attained. We may often *in fact* pass the profit point, but it cannot be guaranteed—nor can we have even a high rational expectation—that we will ever do so. It is this guarantee or expectation that we would need to justify the application of a long run strategy to a short run game.

The consequentialist must start over. This time around, the problem concerns the performance of PCP as a strategy in games of limited length.

## 7.2 *More Failed Attempts*

Some possibilities can be eliminated at once. We cannot demonstrate that PCP usually returns a net gain, because probabilities guarantee nothing. Neither, it seems, will we be able to show that we can be fairly confident of a net gain. Although there is a high probability of PCP advising a winning short run strategy—the longer the game, the higher the probability—we cannot convert this into a rational expectation of gain without PCP itself.<sup>13</sup>

Some philosophers, however, have thought that the short run success of PCP can be derived from the long run frequency postulate. (This possibility will obviously appeal to frequentists, who get LRFP by definition. The philosophers I am referring to, Colin Howson and Peter Urbach, are indeed frequentists.)

As before, there are two ways of proceeding. One is to show that if LRFP is true, PCP usually returns a net gain in short run games. The other is to show that we can rationally expect a net gain in the short run, though it may not be guaranteed.

It is not guaranteed. Even if long run frequencies match probabilities,

---

13. For very short games, the objective probability of winning is not very high. In what follows I consider only reasonably long games; if the application of PCP to these cannot be justified, it certainly cannot be justified for very short games.

short run frequencies can be anything you like, because the outcome types in question might not be evenly spread across space and time. Suppose, for example, that one half of all coin tosses land heads. This is consistent with the possibility that, at this end of the universe, all coins land heads. Why? Because at the other end of the universe, they might all land tails, making for a 50 : 50 distribution over the whole universe. In such a world, the adherent of PCP could well lose badly, if the odds suggest betting on tails in each round.

That is an extreme case, but nature need not be so radically out of balance for probability coordination to fall through. It might be that over all the games you play, you get heads half the time and tails half the time. But this is consistent with your getting all heads in one game, all tails in the next, and so on. Again, if this is the case, PCP might well consistently advise a losing strategy. Of course, it is objectively unlikely that that the distribution of outcomes will be seriously skewed in any of these ways, but for now familiar reasons, this is no help to us.

### *7.3 Indifference*

There is a further possibility to consider. Given LRFP, we can derive a high subjective probability for winning with PCP by appealing to a principle of indifference such as the following:

If we have exactly the same information about a number of different possible events, then our subjective probabilities for the events ought to be the same.

If the principle is granted, we may reason as follows. Suppose that our short run game consists of ninety trials on a setup  $X$ . Consider the set of all actual trials on  $X$ . Our particular ninety trials could be any ninety of these actual trials; we do not know which. The principle of indifference tells us that we should assign equal subjective probabilities that our ninety-trial sequence is any one of the possible selections of ninety actual trials. It is easy to show

that, if LRF is true for  $X$ , the overwhelming majority of these possible sets are such that the frequency of designated outcomes approximately matches the probability.<sup>14</sup> It follows that we should have a high subjective probability that our particular ninety-trial sequence will be one in which the frequency is close to the probability. Should this be the case, PCP advises a winning strategy, so we have an argument justifying our confidence that PCP gives a winning short run strategy.

The reasoning is flawless, but the key premise—the principle of indifference—is highly suspect. It has been recognized for some time that strong principles of indifference such as the one under consideration do not determine unique assignments of subjective probability. (For a recent survey of the relevant arguments, see van Fraassen 1989, chap. 12.)

The difficulty is easily illustrated by the case at hand. In the argument above, our obtaining each possible ninety-trial combination is assigned an equal subjective probability. But the principle may just as well be used to assign an equal subjective probability to each of the following events: that a ninety-trial game contains no instances of the designated outcome, that it contains one instance, that it contains two instances, and so on. Now our subjective probability for (rough) equality of frequency and probability will be of the order of one ninetieth rather than one, entirely at odds with both the previous assignment of subjective probabilities and the goal of justifying PCP.

A possible response to this problem is to impose some plausible restriction on the use of the principle of indifference, by limiting the cases in which sameness of information mandates sameness of subjective probability. In this way Bertrand Russell (1948) argues that subjective probabilities can be set in accordance with frequencies; such a restricted principle is also the basis of

---

14. The result is not so straightforward if there are an infinite number of actual trials on  $X$ . But as always, I will proceed by examining the most favorable case for the would-be justifier of PCP.

Howson and Urbach's (1993) attempt to justify probability coordination. In the next section I will examine Howson and Urbach's justification, arguing first that it invokes a principle of indifference, and second that no plausible restricted version of such a principle saves the argument.

#### 7.4 *Howson and Urbach on Probability Coordination*

As already noted, Howson and Urbach define probability as long run limiting frequency (after von Mises), and so establish LRF by stipulation. Their justification of probability coordination proceeds thus.

Consider the following problem. Suppose that all you know of some particular event (e.g., the outcome of a certain probabilistic experiment) is that its objective probability of being type  $E$  is  $p$ . Then you have no rational choice but to reason in this way:

1. To say that the event has an objective probability  $p$  of being a type  $E$  outcome is (on the frequentist account) to say that the event belongs to a sequence of outcomes in which the limiting frequency of type  $E$  events is  $p$ .
2. If we assign a certain subjective probability to the event's being of type  $E$ , then we must assign the same subjective probability to any other event in the sequence's being of type  $E$ . In other words, the same probability of being  $E$  must be assigned to every event in the sequence.
3. We can assign every event in the sequence either a probability of  $p$  (the long run frequency of  $E$  outcomes) or some probability that is not equal to  $p$ . As we have already seen, to assign  $p$  to all those events is a winning long run strategy. To assign something other  $p$ , however, is in many cases a losing long run strategy (i.e., it is a strategy that, given certain odds, is guaranteed to produce a net loss in the long run).
4. We must assign the event's being  $E$  a subjective probability of  $p$ .

(There are two ways of proceeding from (3) to (4). The consequentialist must depend on the losses and gains that would result from the long run game if it were played. Howson and Urbach proceed slightly differently: they think that to knowingly adopt subjective probabilities that will lead to a long run loss is *inconsistent*, because to adopt a subjective probability  $p$  for an outcome is to “implicitly commit oneself to the assertion that  $p$  gives fair odds” on that outcome’s occurring. This difference is not important, but it does help in understanding the passage in Howson and Urbach from which it is excerpted, quoted below.)

The questionable premise in this argument is (2), which can only be motivated by a principle of indifference. Here is the relevant passage in Howson and Urbach (their example is a coin toss):

But that toss was specified only as a member of [a sequence] characterised by its limit value  $p$ . Hence [by assigning a subjective probability to a particular coin toss’s yielding heads] you have *implicitly committed* yourself to the assertion that the fair odds on heads occurring at *any such bet*, conditional on just the same information that they are members of [a sequence] with [limiting frequency]  $p$ , are [the same] (p. 345, my italics).

Clearly, the “implicit commitment” is taken to be a consequence of our having “just the same information” about each and every toss in the sequence, namely, the limiting frequency of heads in that sequence. Thus though Howson and Urbach do not say so, they tacitly appeal to a principle of indifference along the lines of the one stated above.

Howson and Urbach’s appeal to a principle of indifference might be defended as follows. Although they make no explicit mention of the fact in their argument above, it is a tenet of their theory of probability that if a long run sequence is to determine an objective probability, every outcome in that sequence must be produced by the same physical setup. Thus Howson and

Urbach might be thought to be appealing to a more constrained, thus weaker, yet far more plausible principle of indifference:

If all we know about a set of events is that they have been generated by the same physical setup, then we ought to assign the same subjective probability to each of these events being of some type *E*.

This principle is quite appealing, and for a good reason: it is a disguised version of PCP. To see this, consider the criteria we consider relevant in deciding whether two physical setups are in fact the same. One possibility is to count two setups as the same just in case they are physically identical. This requirement, however, is far too stringent for our purposes. No two setups are physically identical, not even time slices of what we would intuitively call the same setup. As we throw a die, we wear off a few molecules here and a few molecules there, and the mental state of the thrower changes, so two die throws are never exactly the same. For that matter, two atomic decays cannot be identical, as they take place at different points in space-time.

The criterion we want for sameness, then, must be something like this: the setups must be physically identical *in all relevant respects*. But what makes a respect relevant? The molecules that we wear off the die are irrelevant—provided that they are not all worn off the same corner. The location of a radioactive sample is irrelevant—provided that it is the same size, and of the same isotope of the same element as before, and thus has the same half-life. In general, our intuitions tell us, the changes that are irrelevant are those that make no difference to the objective probability of the outcome's being produced. In other words, the above principle's 'same physical setup' means same *chance* setup. Thus the principle really says: assign equal subjective probabilities to outcomes that are produced with the same objective probability. This, of course, is just a special case of PCP. It is now clear why the weaker principle of indifference seems right. It is equally clear that it cannot provide any grounds for *justifying* PCP.

Let me try to tie all these various failures together by asking why, in general, it is so difficult to use LRFP to justify an expectation that short run frequencies will match probabilities.

First of all, there is no direct relation between long run frequencies and short run frequencies. No short run (finite) frequency in itself makes a difference to the value of the corresponding long run (limiting) frequency, nor need the long run frequency be at all reflected in a given short run frequency.

Second, although there is an *indirect* relation between long run frequencies and short run frequencies—that is, a way of inferring short run frequencies from a given long run frequency—the inference goes by way of objective probabilities, and thus depends on probability coordination. This requires some explanation. To get from a long run frequency to a prediction about a short run frequency, we use inverse inference (the process by which we infer the values of probabilities from statistics) or a frequency definition to go from the long run frequency to the relevant objective probability, and from this derive a high objective probability that the short run frequency will match the long run frequency. But to go from this high objective probability to a high level of expectation, we need PCP.

Once again we encounter a difficulty reminiscent of the problem of induction. It seems we could justify the inferential rule in question by appealing to a global empirical posit about the uniformity of things. (In the case of PCP, the posit asserts that long run frequencies are usually reflected in the short run.) But for reasons with which we are now familiar, any argument for the uniformity postulate would have to be based in part on the inferential rule that was itself in need of justification.

### 7.5 *Foundations For A Short Run Match Between Frequency And Probability*

This is what we have come to: Just as in the case of the long run, to justify adherence to PCP in the short run we have to find some connection between

short run frequencies and probabilities that is independent of PCP. The obvious candidate is what I will call the *short run frequency postulate*:

(SRFP) In most of our short run games, the frequency of occurrence of the designated outcome is close to the probability of that outcome.

To vindicate probability coordination we have to provide a foundation for the short run postulate. The options are familiar: a kind of frequentism, or externalism.

I will begin with frequentism. In order to make SRFP true by definition, we would have to define probability so that, in most short run games (of substantial length), the frequency would more or less match the probability. This could be done by defining probability as short run frequency. As I have already commented, such a definition would entail a hopeless instability in the values of probabilities. For example, let  $p$  be the probability of a certain coin's landing heads. Then when that coin happened to be in the middle of a run of heads,  $p$  would be near one; a few minutes later, after a cluster of tails,  $p$  might be down near one quarter. Under such circumstances, probabilistic reasoning would be pointless: as soon as we learned the values of probabilities they would change to something completely, unpredictably different.

A more sophisticated definition of probability is as follows. Probability is set equal to limiting frequency (and so LRFP is made true by stipulation), except that an outcome type is not counted as having a probability at all unless the outcome frequencies are sufficiently homogeneous. That is, to be deemed a probability, a long run frequency must be mirrored in the short run frequencies throughout the length (in time) and breadth (in space) of the universe. The stability of probabilities is thus assured.

This definition is, however, not quite sufficient to found SRFP. We need a further postulate to the effect that *our* games are not among those few sets of (spatiotemporally contiguous) trials that suffer from remaining inhomogeneous.

geneities.<sup>15</sup> To derive a high rational subjective probability that our games are among the majority of sets of trials, we would need either a principle of indifference or a rule for probability coordination, options which are not available. An alternative is to brazen our way out of the situation by adding the new postulate to the definition of probability, in the process relativizing the facts about probabilities to a community of inquirers. A braver thing to do is to surrender to externalism, accepting that there is no way to show that we ought to expect to be unaffected by inhomogeneities.

Before I move on to externalism, I will comment on this tactic of justifying PCP as a principle of direct inference by tweaking the definition of probability. Trying to solve philosophical problems with definitions is much like trying to alter the landscape by redrawing the map. Distinctive topographical features (philosophical problems) do not disappear; they simply acquire new names. In the case at hand, though the new definition of probability makes the justification of direct inference easier, the justification of the rules for inferring the values of objective probabilities (the rules of inverse inference) becomes that much harder, for the following reason.

To justify inverse inference, one must not only provide a method for getting from evidence to probability, one must also give some reason for believing that there are probabilities to be got at. If there is no such reason, then there is no warrant for relying on the numbers that will be churned out, no matter what, by the usual rules of inverse inference.

Now, the more stringent our definition of probability, the bolder the conjecture that there are any probabilities at all. On a finite frequency account of probability, for example, probabilities must always exist (if there are any outcomes). On a limiting frequency account, however, this is no longer the case. The frequency with which a designated outcome occurs may have no

---

15. We could define probability so that there were no inhomogeneities at all, but (technical difficulties aside) the new definition of probability would not be instantiated; one thing we know for sure is that frequencies *sometimes* depart from probabilities.

limiting value (especially if nature is not, as they say, uniform). Thus the proponent of a limiting frequency account must not only justify rules for inferring limiting frequencies from short run frequencies, but also offer some reason for thinking that the relevant limiting frequencies exist in the first place. And so it goes: if it is true by the definition of probability that we do not occupy a local inhomogeneity (a place where the local frequency of the relevant event is not equal to its global frequency), then the justifier of inverse inference must, in order to demonstrate the existence of probabilities, show that there is reason to believe that we do not occupy a local inhomogeneity. As we add further clauses to the definition of probability, we make the existence of probabilities, and thus a justification of inverse inference, more and more difficult to establish.

In fact, we make it more difficult by *precisely* the amount that the same definitions made the justification of direct inference easier. Those definitions excused the justifier of direct inference from showing that limiting frequencies exist, that designated outcomes are homogeneously distributed, and that we do not occupy one of the remaining inhomogeneities. Now the justifier of inverse inference is faced with the very tasks that were to have been avoided. As I remarked above, redefining the linguistic borders only shifts problems from region to region; in this case, from the domain of direct inference to that of inverse inference. Upon encountering the same problems for the second time, the only option is externalism.

The strategy of the externalist is to hope fervently that the short run frequency postulate is true. If this hope is vindicated, then PCP will usually advise winning short run strategies, though—except in retrospect—we cannot know it. And that is about all that can be said.

## 8. The Psychological Origins of Probability Coordination

We have now seen that no good argument can be given in support of probability coordination. That is the *logical* component of my skeptical argument.

This conclusion notwithstanding, some will argue that the sheer force of our conviction that probability coordination is rational must count for something in itself. (This is the root of the well-known argument that any really powerful skeptical argument must be interpreted as a *reductio*, since the negation of the conclusion is better supported than the premises.) This position will be undermined, however, if it can be shown that the relevant “force of conviction” has arisen for reasons that are unconnected with the merit of the principle. Such a demonstration is the other half, the genealogical component, of the Humean skeptical assault. It is the aim of this section to provide the necessary genealogical argument.

Why does probability coordination seem intuitively so right? This is a psychological problem, the problem of identifying the cause of our firm belief in probability coordination. From an epistemic point of view, causal explanations of this belief can be divided into two classes:

1. *Logical* explanations claim that the cause of our belief in PCP is also a *reason* to believe PCP. That is, we believe the principle because we correctly apprehend its truth, or the reasons for its truth.
2. According a *Humean* explanation, PCP is a mere “habit”, a psychological disposition that has been acquired not through good reasoning but in some other way. If our belief in PCP has a Humean explanation, the fervor with which PCP is maintained bears no relation to its epistemological status. The PCP could be quite unjustified, yet “obvious” to us.

The possibility of a logical explanation of our commitment to probability coordination is destroyed by our failure, for apparently principled reasons, to find any epistemically accessible justification of probability coordination whatsoever. Even if some genius were to come along tomorrow with an unforeseen but brilliant vindication of PCP, this vindication could not be the reason that *we* believe PCP. (Unless, that is, we have temporarily “forgotten”

the argument, like Meno's slave. Some aspects of our mental life—for example, syntactic and visual processing—are indeed buried very deep in the mind, but a priori philosophy does not seem to be one of them.) The cause of our belief in PCP cannot be a logical “because”.

Since there is no logical explanation for our adherence to PCP, the explanation *must* be Humean. To discover the correct explanation, we should consider stories as to how we could possibly have picked up the interesting habit of setting our subjective probabilities equal to the objective probabilities. There are four possibilities:

1. *Cradle or book learning*: The PCP is a cultural artifact passed down from generation to generation.
2. *Faulty (deductive) reasoning*: We have been led to PCP by a bad argument.
3. *Learning by induction*: We have learned to follow PCP by way of some general inductive rule.
4. *Evolution*: The process of evolution has resulted in our being equipped with PCP by our very nature; i.e., more or less as a matter of genetics.

The first suggestion is obviously false. Nor does the second suggestion seem likely: if our belief in the rationality of probability coordination were to be explained by faulty reasoning then (a) we would remember that reasoning, and (b) our faith in the principle would be diminished by considerations such as those offered in the first part of this paper.

The remaining two possibilities, then, are that adherence to PCP is learned or has evolved. In either case, the story of our acquisition of the principle is similar. The PCP has been found to be successful in the past, because past short run frequencies have usually been close to objective probabilities. As a result, the rule is installed in the brain either by way of an inductive learning

rule already in place, or by Mother Nature (who is herself something of an induction artist).

There is an obvious objection to this kind of story. Why should an organism unequipped with PCP ever be sensitive to objective probabilities in the first place? Clearly, the concept of objective probability must have appeared at the same time as the practice of probability coordination. Either the concept and the practice could have sprung up more or less fully formed, or they could have gradually developed from less sophisticated analogues, such as the notion of a frequency and a disposition to project frequencies. (The considerable sensitivity of various animals—pigeons, rats and so on—to frequencies is documented in Gallistel 1990.) I will not concern myself with the details of these stories. The unavoidable conclusion is that some form of either the learning story or the evolutionary story must be true, because these are the only possible explanations of our passion for probability coordination.

In summary, our belief in the rationality of probability coordination seems to be explicable, but not demonstrably justifiable. There is compelling reason to think that the “obviousness” of PCP is a result of the principle’s past success. Thus probability coordination and Humean induction are seen to be related in one last way: our belief in PCP seems to result from the application of a more general inductive rule, either a learning rule implemented in the human brain, or the implicitly inductive rule of natural selection for past success. Will PCP continue to be successful? Will short run frequencies continue to reflect objective probabilities? Will the future resemble the past in the appropriate way? We reluctant disciples of induction can only hope.

## 9. Concluding Remarks

Not all the consequences of my argument are negative. It might have been thought, for example, that any reasonable account of objective probability

must provide a foundation for probability coordination.<sup>16</sup> This paper shows that to adopt such a constraint would be a mistake. However, the news is mostly bad news.

A proponent of probability coordination may claim that I have not looked hard enough, that I have not been clever enough, that there *must* be some kind of philosophical justification for the practice because it is just so obviously rational. To make such claim, though, is to mistake the strength of one's convictions for the quality of one's argument. Our feelings concerning PCP, I have shown, offer no objective support for the principle, but are rather a psychological footprint left by the principle's past success.

There have been attempts to justify induction and the like without providing an argument from solid premises. These attempts may involve "non-vicious circularity" or "reflective equilibrium" (though our "decision" to coordinate probabilities involves very little reflection, if much equilibrium). I will not discuss the virtues of these proposals here (the issue has little to do with what is distinctive about probability coordination or objective probability). Suffice to say that, should you encounter a perfectly rational person patiently waiting for a broken egg to pull itself together, there is nothing you can say that will persuade that person to abandon their lonely post.

---

16. Wesley Salmon (1967) appears to impose such a requirement. But he is well aware of the difficulties of demonstrating the rationality of probability coordination, and he later appears to fall back on a different requirement, that there be some demonstrable relationship between probabilities and frequencies (a requirement easily satisfied by his favored account of probability, the frequency account).

## Appendix A: The Variable Stake Game

If Mother Nature is allowed to name any stake she likes in each round of a game, then neither probability coordination nor any other strategy can guarantee a net gain, even on the assumption that SRFP is true. This is because one huge loss may be so crippling as never to be recovered. (Even in the long run game Mother Nature can employ some devious strategy such as forever halving the stake after a big loss, so that we cannot break even no matter how often we win thereafter.)

In our real life, short run games this is particularly evident: at Mother Nature's table, the stake is frequently life itself, not the sort of loss that can be recouped later. Even the most unheroic of us chooses constantly to put up such a stake. For example, any time that we drive somewhere, we are playing a round in which there is on one hand a very small chance of death, and on the other hand a very large chance of a small to medium amount of utility.

Even given SRFP, the consequentialist has no reason to think that probability coordination is the correct strategy in such games. This point can be made more vividly with the help of a Peircean example (after Charles Peirce, who gave this problem considerable thought (Peirce 1878)). Suppose that you must choose between two alternatives, both of which may lead to death, one with an objective probability of one tenth, one with a probability of nine tenths. How should you choose? The PCP strongly favors the first option, but there can be no justification of this choice by appeal to a guarantee of a net gain over a number of games, for a loss in the first round terminates the game.

A consequentialist justification of probability coordination in the face of death must deal with exactly this kind of case. It is fairly clear that there is only one kind of solution to the problem: to take a wider perspective. Although the consequentialist cannot justify death-defying invocations of PCP at the level of the individual, consequentialism becomes possible again at a higher level, in which we consider the good of a group of individuals as a

whole. (This, incidentally, was Peirce's solution also.) In the Peircean game above, a community will minimize its net loss if all individuals follow the PCP strategy, assuming, as usual, that the relevant short run frequencies are close to the probability. (Thus, it may be noted, coordinating probabilities in such situations will be favored by natural selection, all other things being equal.)

Of course, games can be—and have been—constructed in which groups face death as a whole, rather than individually. To justify probability coordination in these situations, one must move to a higher level still. The individual adherents of the principle, meanwhile, must work up a philosophical concern for their similarly inspired cohorts. Apart from this, there is very little that can be said to ameliorate the problem of varying stakes.

## Appendix B: Proofs

In this section I show that PCP is a winning strategy in constant and variable odds games if the frequency with which the designated outcome occurs is close enough to its probability.

*Lemma:* The PCP advises a winning strategy for any short run game if the relative frequency with which the designated outcome occurs is exactly equal to the objective probability of that outcome.

*Proof:* Consider an  $n$ -trial game in which the probability of the designated outcome is  $p$ . Suppose that the odds offered by Mother Nature are  $j : k$  in favor of the designated outcome (where the entire stake in each game is  $j + k$ ). Without loss of generality, suppose also that  $p > \frac{j}{j+k}$ , so that PCP advises us to bet on the designated outcome in every round (since the chance of winning is greater than the proportion of the stake that must be put up by bettor). Then for every round in which the designated outcome turns up, we win  $k$  utility points, and for every round in which it does not, we lose  $j$  points. At the end of  $n$  games, our net gain is  $ak - (n - a)j$ , where  $a$  is the number of times the designated outcome was realized. By assumption, the relative frequency with

which the designated outcome occurs is equal to the objective probability  $p$ . Thus  $a = np$ , and our net gain is  $npk - n(1 - p)j$ . We can now show that if  $p > \frac{j}{j+k}$ , as assumed, then  $npk > n(1 - p)j$ ; i.e., PCP results in a net gain:

$$\begin{aligned} p > \frac{j}{j+k} &\Rightarrow np(j+k) > nj \\ &\Rightarrow npj + npk > nj \\ &\Rightarrow npk - n(1-p)j > 0 \end{aligned}$$

*Theorem:* The PCP advises a winning strategy in the short run if the relative frequency  $f$  with which the designated outcome occurs is close to the objective probability of that outcome  $p$ , where  $f$  is defined to be close to  $p$  if  $f$  is closer to  $p$  than are the odds offered by Mother Nature, that is, if  $|f - p|$  is less than  $\left| \frac{j}{j+k} - p \right|$ .

*Proof:* We prove a slightly stronger claim, that PCP will be a winning strategy if either (a)  $f$  lies on the other side of  $p$  from the odds, or (b)  $|f - p|$  is less than  $\left| \frac{j}{j+k} - p \right|$ . A game that fits this description has the same properties as a game in which the objective probability is equal to  $f$ , not  $p$ , and the frequency matches the probability exactly. The PCP will advise the exactly the same strategy in both games, and (as a result of the lemma) this strategy will yield the same net gain.

*Theorem:* The PCP advises a winning strategy in a long run game if the limiting frequency  $f$  with which the designated outcome occurs is close enough to the objective probability of that outcome  $p$  (in the sense of ‘close’ just defined).

*Proof:* As before, we prove the slightly stronger claim that PCP will be a winning strategy if either (a)  $f$  lies on the other side of  $p$  from the odds, or (b) is less than . Given a long run game, let  $f_m$  designate the relative frequency with which the designated outcome occurs over the first  $m$  trials of the game. By assumption, the limiting frequency of designated outcomes is equal to  $f$ , which means that for all  $\epsilon$ , there exists  $N_\epsilon$  such that for all  $m$  greater than

$N_\varepsilon$ ,  $|f_m - f|$  is less than  $\varepsilon$ . Choose  $\varepsilon < \left| \frac{j}{j+k} - f \right|$ . Then for all  $m$  greater than  $N_\varepsilon$ , either (a)  $f_m$  lies on the other side of  $p$  from the odds, or (b)  $|f_m - p|$  is less than  $\left| \frac{j}{j+k} - p \right|$ . It follows from the proof of the previous theorem that  $N_\varepsilon$  is the profit point for the game; that is, that after  $N_\varepsilon$  trials, a player following PCP will have made a net gain that can never disappear.

*Proposition (Variable Odds Games):* Let  $J:K$  be the event of Mother Nature's naming the odds  $j : k$  in a given round. Let  $A$  be the event of the designated outcome turning up in a given round. Then if (a) LRF holds for the probabilities  $\Pr(A|J:K)$ , and (b) the events  $J:K$  and  $A$  are stochastically independent, the above results hold for a variable odds game. The proof is obtained by considering a variable odds game to be a mix of a number of constant odds games.

## References

- Bogdan, Radu, ed. (1984). *Henry E. Kyburg, Jr. and Isaac Levi*. Dordrecht: Reidel.
- Braithwaite, R. B. (1966). "Why is it reasonable to base a betting rate upon an estimate of chance?" *Proceedings of the 1964 International Congress for Logic, Methodology and Philosophy of Science*. Edited by Y. Bar-Hillel. Amsterdam: New Holland. 263-273.
- Earman, John. (1992). *Bayes or Bust?* Cambridge, MA: MIT Press.
- Gallistel, Charles. (1991). *The Organization of Learning*. Cambridge, MA: MIT Press.
- Goodman, Nelson. (1983). *Fact, Fiction, and Forecast*. Fourth edition. Cambridge, MA: Harvard University Press.
- Hacking, Ian. (1977). *The Emergence Of Probability*. Cambridge: Cambridge University Press.
- Hall, Ned. (1994). Correcting the guide to objective chance. *Mind* 103:504–517.
- Howson, C. and P. Urbach. (1993). *Scientific Reasoning* 2nd ed. La Salle, Illinois: Open Court Press. (First edition 1989).
- Lewis, David. (1980). "A subjectivist's guide to objective chance". Reprinted in Lewis (1986).
- . (1986). *Philosophical Papers Volume Two*. Oxford: Oxford University Press.
- . (1994). "Humean supervenience debugged". *Mind* 103:473–490.
- Mellor, D.H. (1971). *The Matter of Chance*. Cambridge: Cambridge University Press.
- . (1982). Chance and degrees of belief. In R. McLaughlin, ed., *What*

*Where? When? Why?*. Dordrecht: Reidel.

Miller, David. (1966). "A paradox of information". *British Journal for the Philosophy of Science* 17:59–61.

Peirce, Charles. (1878). "The doctrine of chances." Reprinted in *The Philosophy of Peirce: Selected Writings*. London: Routledge, 1940.

Reichenbach, Hans. (1949). *The Theory of Probability*. Berkeley: University of California Press.

Russell, Bertrand. (1948). *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.

Salmon, Wesley. (1967). *The Foundations of Scientific Inference*. Pittsburgh: The University of Pittsburgh Press.

———. (1979). "Propensities: a discussion review". *Erkenntnis* 14:183–216.

Stich, Stephen. (1990). *The Fragmentation of Reason*. Cambridge, MA: MIT Press.

Strevens, Michael. (1995). "A closer look at the 'New' Principle". *British Journal for the Philosophy of Science* 46:545–561.

Thau, Michael. (1994). Undermining and admissibility. *Mind* 103:491–503.

van Fraassen, Bas. (1989). *Laws and Symmetries*. Oxford: Oxford University Press.

*Department of Philosophy*  
*Stanford University*  
*Stanford, CA 94305*  
*USA*