A Closer Look at the 'New' Principle

Michael Strevens

British Journal for the Philosophy of Science, 46:545–561. 1995.

Introduction

Revolution?

In his 1980 paper 'A Subjectivist's Guide to Objective Chance', David Lewis argued that much of our reasoning about objective probabilities rests on an intuitively compelling rule of inference that Lewis dubbed the 'Principal Principle'. In its simplest version, the Principal Principle states that one ought to set one's subjective probabilities equal to the corresponding objective probabilities, provided that one has no 'inadmissible' evidence. (The principle will be explained at greater length below.)

At the time, Lewis announced that he had one reservation about the Principal Principle: it seemed to be inconsistent with the philosophical doctrine of Humean supervenience concerning objective probabilities, or as they are sometimes called, *chances*. (Again, details will be found below.) In a recent triplet of articles (Lewis 1994, Thau 1994 and Hall 1994), Lewis and two of his associates have concluded that this 'one reservation', when examined more closely, is enough to undermine entirely the Principal Principle as we know it. Hall and Lewis claim that that principle—now deemed the 'Old Principle'—

ought to be replaced with a new rule they call the 'New Principle'. Thau seems to concur.

What we are left with is "a revolutionary view", according to Thau. The old, intuitive Principal Principle is dead; a new, unfamiliar, yet supposedly more rational principle has taken its place. It is notable, however, that none of the three authors is able to supply compelling reasons to think that the New Principle is legitimate. At best, the evidence is suggestive: the New Principle does not suffer from the problems of the old, and it supports approximately the same (intuitive) inferences as the old, in most situations.

It is the aim of this paper to supply the necessary justification for the New Principle, though not quite in accordance with the wishes of Lewis, Hall and Thau. By closely scrutinising the history of the 'revolution'—the process by which the three authors replace the old principle with the new—I will show that we do not have to abandon the old principle at all. On the contrary, it turns out that two principles are consistent with one another, and that they have their roots in the same intuitions and time-honoured inferential habits.. Thus we are able to keep what is intuitive (the Principal Principle), and motivate what was apparently unmotivated (the New Principle).

The Principal Principle

The Principal Principle tells us how to set our subjective probability (credence, degree of belief) that an event A will occur, given a hypothesis about the objective probability that A will occur. To be precise, it tells us to set the subjective probability *equal to* the objective probability. For example, given the hypothesis that the chance of a tossed coin's landing heads is one half, the Principal Principle tells us that we ought to set our subjective probability that the coin will land heads to one half.

More formally, suppose we are given a type of chance experiment X (e.g., a coin toss), a possible outcome E (e.g., heads), and a probabilistic theory T that tells us that the chance of a type X experiment's producing a type E outcome is p. Let A be the proposition that some particular type X experiment produces an E outcome (e.g., that some particular coin toss lands heads). Then we ought to set our subjective probability function C(·) according to the following rule:

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

where K is any other body of information.[1] (When the Principal Principle is used as a rule of direct inference, as in the cases under consideration in this paper, K will be taken to represent all background knowledge.)

The above rule holds only provided that K includes no *inadmissible* evidence. Inadmissibility is discussed at length below. For now I will simply mention what I call Lewis's *working characterisation* of admissibility: a piece of information is admissible just in case either (1) it is information about some state of affairs that obtains prior to the experiment in question, or (2) it is information about chance itself, that is, the information provided by T (Lewis 1980, 92–4). In the case above, for example, information about the actual outcome of the coin toss is inadmissible, as it concerns a state of affairs that does not obtain before the experiment begins. (Intuitively, if we know that the coin is going to land heads, then we ought not to set our subjective probability that it will land heads to one half.)

I will write the objective probability p as a function of A, so that p = $P_T(A)$. (Thus $P_T(\cdot)$ is just the probability function encapsulated in T.) The Principal Principle is then written:

$$C(A|TK) = P_T(A).$$ [2]

---

1 Lewis's version of the principle is more complex, as it allows the probability of the event described by A to change with time. For expository convenience, I avoid this complication by assuming that the event is specified to be an event of a kind to which the theory T explicitly assigns objective probabilities.

2 A frequentist (e.g., Reichenbach 1949) might prefer to avoid postulating the existence of single case probabilities by formulating the principle as

In what follows, it will be convenient to consider the Principal Principle as being made up of two parts: a standard of admissibility, and a rule for setting subjective probabilities that must be followed if all evidence meets this standard. (Thus the rule together with the standard of admissibility make up the complete, qualified rule that Lewis calls the Principal Principle.)

The Structure Of The Argument

As I have said above, I will make my case in the course of a careful examination of the process by which Lewis, Hall and Thau reject the Principal Principle in favour of the New Principle. Like the principle itself, this process has two parts. In what I will call round one, Lewis's old working characterisation of admissibility is discarded in favour of a new characterisation, described by Thau and endorsed by Lewis. In round two, the old rule for setting subjective probabilities is abandoned, and replaced by a new rule which is the heart of the New Principle, namely,

$C(A|TK) = P_T(A|T)$.[3]

I will dispute neither of these two amendments. On the contrary, I will argue that Thau and Lewis's new characterisation of admissibility is the correct definition, and that the new rule for setting subjective probabilities is a valid rule. However, in understanding more clearly why Lewis, Hall and Thau are right about these things, we will in

---

$C(A|TK) = P_T(E)$, where, as stated above, A is a type E event. Interestingly, the admissibility constraint then enforces the rule that Reichenbach and other frequentists take to solve the reference class problem, that E should be the narrowest event type such that (i) A is of type E, and (ii) we have reliable knowledge of P_T (E).

3 In the versions of this rule presented by Lewis and Hall, K is restricted to historical information. Provided that an appropriate admissibility constraint is added to the rule, this restriction is not important. I will return to this point at the end of the paper, where I discuss the necessity of such a constraint.

addition see that they are wrong about a third thing: they are mistaken in their claim that the old Principal Principle must be discarded.

Round One

Problems of Consistency

Between them, Lewis, Hall and Thau raise two problems with the Principal Principle. The first is a conflict with the metaphysical doctrine of the Humean supervenience of objective probabilities. This problem is the common concern of all three authors. The second problem, mentioned only by Hall, is the alleged inability of the Principal Principle to deal with what Hall calls ‘crystal balls’. Both may be considered to be problems of consistency: the Principal Principle is apparently inconsistent with certain accounts of probability, and with the existence of crystal balls.

Let me begin with the problem that has received the more attention, that of Humean supervenience. Objective probabilities are said to be *Humeanly supervenient* just in case the facts about probabilities are (roughly speaking) entirely determined by particular non-probabilistic matters of fact. Note that on such an account, it is typically the case that, for any given event, the matters of fact that determine the probability of that event inhabit the past, present *and* future.

A good example of a Humean supervenience account of objective probability (i.e., an account on which the probabilities are Humeanly supervenient) is the *actual frequency* account. According to this view, if A is the proposition that some particular experiment of type X will produce a type E outcome, then the probability of A is the actual frequency with which type X experiments—including those yet to be conducted—produce E outcomes.[4]

The Principal Principle, Lewis argues, is inconsistent with any Humean supervenience account of objective probability that allows what Lewis, Thau and Hall call *undermining futures*. An account of objective probability allows for an undermining future in the following circumstances.

First, the theory must assign probabilities to the particular matters of fact that *determine* those same probabilities. We can see that this is, for example, true of the actual frequency account: the account assigns probabilities to the outcomes of type X experiments, yet those outcomes themselves determine the probability that is assigned. As a consequence of this circular relation between probabilities and outcomes, there will be a certain probability (perhaps zero) assigned to a future in which the probability-determining facts turn out differently than they in fact do. That is, though it is a fact (about the future) that the probability-determining facts turn out one way, there is a probability that they turn out another way. But this means that there is a probability assigned to the event of the *probabilities* being other than they are. (For example, in the case of the actual frequency account, there is a certain probability that the frequency will be other than it is.)

Second, if it is possible for this probability to be *non-zero*, then the underlying philosophical account of probability allows undermining futures. It allows for a set of probability ascriptions T that, among other things, ascribes a positive probability that T—the set of ascriptions itself—will turn out *not to be the case*. More formally, it is allowed that there are Ts such that $P_T(\sim T) > 0$. (Again, the actual frequency account exhibits this feature because, assuming the independence of the probabilities of individual outcomes, the probability of a given finite

4 For expository convenience, I am assuming that the frequentist desires to derive so-called single case objective probabilities from the frequencies. This is true of Lewis's frequentism, but is not necessary; see footnote 2 for a version of the Principal Principle that allows one to reason about single cases without dealing with single case probabilities.

frequency is always less than one.) This is certainly a curious fact, but Lewis is prepared to tolerate it. (Others are not—see Bigelow, Collins and Pargetter 1993.) Indeed, as Lewis thinks that his own frequency-based account of probability commits him to the existence of undermining futures, he does not have much choice.

Now, says Lewis, consider a Humean supervenience theory for which there exist undermining futures. What is the proper value for our subjective probability C(~T|T)? According to the Principal Principle, C(~T|T) ought to be set to P_T(~T), which is by assumption greater than zero. But the principle must be wrong in requiring this, since the axioms of probability require that C(~T|T) be zero. So the Principal Principle cannot be true for any theory that allows undermining futures. Since Lewis's Humean supervenience account of probability *does* allow undermining futures, he must find a substitute for the Principal Principle that allows him to ignore the fact that P_T(T) is greater than zero.

I now turn to the second objection to the Principal Principle, the problem of the mechanisms that Ned Hall calls *crystal balls*. A crystal ball is a device that allows us to see the outcome of future events, such as coin tosses or measurements of fundamental particles. Clearly, if we have a crystal ball, and we see that a tossed coin is going to land tails, then we ought to set our subjective probability in that coin's landing heads to zero. More formally, let A be the proposition that the coin will land heads, let K include the fact that the crystal ball shows the coin landing tails, and let T (the probabilistic laws) state that the crystal ball successfully predicts outcomes with probability one. Then KT entails that A is false. Thus C(A|TK) must, as a matter of logic, be set to zero.

The Principal Principle, however, tells us to set C(A|TK) to one half—unless the information from the crystal ball is inadmissible. If the principle is to be consistent with the possible existence of crystal balls, then, the evidence gleaned from a crystal ball had better not be admissible. Unfortunately, this is not the case, at least according to Lewis's working characterisation of admissibility. On that characterisation, information about any state of affairs prior to the experiment in question is admissible. But the image of the coin landing tails hovers in the crystal ball before the toss takes place, so information about that image cannot be declared inadmissible.

Admissibility Reconsidered

Lewis and Thau first attempt to solve the inconsistency problems just described by amending the characterisation of admissibility. The solution eventually turns out to have its own problems, leading to round two and the introduction of the New Principle. Let us begin, however, by considering admissibility in isolation. (Note that the views in this section are those of Lewis and Thau only. Hall has nothing to say about admissibility, as he considers it unnecessary once the New Principle is adopted. Later in this paper we will see that he is mistaken.)

When Lewis first introduced the Principal Principle, he wrote that he had “no definition of admissibility to offer”, but would “suggest sufficient (or almost sufficient) conditions for admissibility” (1980, 92). He then went on to present what I have called his working characterisation, described above. It is clear that he did not consider the working characterisation to be canonical, but rather regarded it as an approximation to the correct definition, which is, we may assume, lodged somewhere deep in the recesses of the intuition.

In their 1994 papers, Thau and Lewis appeal to these intuitions to defend a notion of admissibility that is incompatible with the working characterisation in certain key respects—respects which, as we will see, allow a rather different treatment of the apparent inconsistencies described above. The new characterisation, we are led to suppose, is closer to—or perhaps even identical with—the correct definition of admissibility. Thus it turns out that the Principal Principle’s ‘inconsistencies’ are artifacts of Lewis’s original provisional characterisation of admissibility, and disappear when admissibility is properly understood.

The new characterisation of admissibility, as stated by Thau (Lewis follows Thau), is contained in the following assertion:

> A proposition is inadmissible if it provides direct information about what the outcome of some chance event is. (Thau 1994, 500)

It is clear from the context that ‘if’ here has the force of ‘if and only if’, and that ‘direct’ information is information over and above that contained in the relevant assertions of probability. We can thus rephrase Thau’s characterisation as follows: an admissible piece of information, relative to an application of the Principal Principle to some objective probability $P_T(A)$, is one that contains information about A only insofar as it contains information about the *chance* of A, $P_T(A)$. We ought to add that, if a proposition B is to be admissible, it must also be the case that B *conjoined with T* tells us no more about A than does $P_T(A)$. (It will normally be simpler to neglect this qualification, but it becomes important in the section on crystal balls immediately below.)

In short, then, given an application of the Principal Principle to some objective probability $P_T(A)$, a piece of information B is inadmissible if B (or BT) tells us something about A that the value of $P_T(A)$ does *not* tell us. (Rather than continually saying “relative to an application of the Principal Principle to the objective probability $P_T(A)$” I will from now on say simply “relative to $P_T(A)$”.)

To see how this characterisation works, we may take two simple examples. First, if some piece of evidence B contains no information about A, then B will be admissible (relative to $P_T(A)$). Thus information about the future courses of the planets is (except in extremely unusual circumstances) admissible relative to probabilities concerning the outcomes of coin tosses. Second, if A (or ~A) is a logical consequence of B or BT, then B is inadmissible relative to $P_T(A)$, unless $P_T(A)$ is one (zero). For example, information to the effect that a tossed coin will land heads is inadmissible relative to the probability of the outcome of the toss. (It is immediately apparent that this new characterisation of admissibility is different from the working characterisation. On Thau’s characterisation, information about the future may be admissible, and information about the past inadmissible. For further comments about the relation between the two characterisations, see Thau, p. 500.)

Of course, B may contain information about A without B or BT entailing A. However, the fact that entailment means inadmissibility is sufficient to remove any inconsistency with undermining futures and crystal balls, for the following reasons.

Crystal balls: Let T state that a certain crystal ball is infallible, and let K include the information that this ball has predicted that a certain event will not occur, and thus that the proposition A that states that the event will occur is false. Hall's problem (which occurs only when $P_T(A)$ is greater than zero) is caused by the fact that the working characterisation rules that K is admissible. On Thau's characterisation of admissibility, however, K is inadmissible, since KT entails ~A. Thus the Principal Principle cannot be applied to $P_T(A)$, and we may set C(A|KT) to zero, as the crystal ball suggests we should.

Conflict with Humean Supervenience: (This result is sketched by both Lewis and Thau in their 1994 papers.) The conflict is a result of the fact that the working characterisation rules that T is admissible relative to $P_T(\sim T)$ even when $P_T(\sim T)$ is greater than zero. It will be seen that on Thau's characterisation of admissibility, this is not the case (because T entails that ~T is false). Thus the Principal Principle cannot be applied to $P_T(\sim T)$, and we are free to set C(~T|KT) to zero, as logic demands.

How close is Thau's characterisation to the correct definition of admissibility? For the purposes of my argument, the question could remain unanswered, because at the time of writing there is no other characterisation of admissibility in play. Lewis has renounced his old working characterisation in favour of Thau's (1994, 485) and Hall tries to do without admissibility altogether. Since we are provisionally agreed on admissibility, we can proceed, provisionally, on the basis of this agreement.

There is, however, one strong consideration that may be advanced in favour of Thau's characterisation. Admissibility (or rather, inadmissibility) seems to be a phenomenon that arises in every epistemic context, not just those involving probabilities. Consider: given any propositions A, B and K, when should C(A|KB) not be set equal to C(A|B)? The answer: just in case K tells us something about A over and above what B tells us.[5] To put it another way, C(A|KB) can be set equal to C(A|B) just in case K is admissible, where K is admissible if it contains no new information about A. This observation is just as true when reasoning about chances as when reasoning about anything else. Now we can see that, of the two parts of the Principal Principle—the rule for setting subjective probabilities equal to objective probabilities, and the admissibility constraint—the first is proper to the epistemology of probability, but the second is just the local application of a universal rule concerning conditional subjective probability.

In light of this observation, it becomes quite plausible that Thau's characterisation captures the outlines of the correct definition of admissibility. (I say 'outlines' because the characterisation must be fleshed out with an account of which propositions 'contain information over and above' that contained in other propositions.[6])

---

[5] B may of course contain new information that just confirms the subjective probability suggested by K. Thus 'over and above' must be read so as to imply some tension between B and K.

[6] There are more or less objectivistic ways of fleshing out the notion of admissibility. The more objectivism, the stronger the admissibility constraint will be. On the most subjectivistic notion of admissibility, the admissibility constraint would amount to this: you may set C(A|KB) equal to C(A|B) just in case you believe that K contains no information about A over and above that contained in B, regardless of the grounds of this belief. The constraint is strengthened as we impose requirements on the situations in which one proposition may be believed to contain, or not to contain, information about another. An obvious and uncontroversial requirement is that needed to save the Principal Principle from inconsistency, that if KB entails A (and B does not) then K is inadmissible. At the objective end of the spectrum, it may be required that facts about 'information containment' be determined by some kind of inductive logic. Fortunately, for my purposes it is not necessary to make any difficult decisions about these things.

Round Two

At the end of round one, we saw that the Principal Principle could be made consistent with the existence of undermining futures and crystal balls by appealing to an independently plausible and intuitive characterisation of admissibility, one which is in any case agreed upon by all parties. In this section we see why this result is at best only "the beginning of a solution" (Lewis 1994, 485).

A Problem of Utility

Adopting a better characterisation of admissibility saves the Principal Principle from the inconsistency problems, but it creates a new and far less academic problem, described by both Lewis and Thau. It seems that on most Humean supervenience accounts of chance, information about chance itself is inadmissible. As a consequence, the Principal Principle, though consistent with Humean supervenience, cannot be applied to Humeanly supervenient probabilities. In short, for the Humean, the principle is completely useless.

How could it be that a chance theory contains information over and above that provided by the chances? Let us work with a frequency account of chance. Observe that if our proposition A concerns a type E event, information about future frequencies of type E events is inadmissible. For example, if we know that nine out of the next ten tosses of a fair coin will land heads, we will not set our subjective probability that the very next toss lands tails to one half. In general, this kind of frequency information is relevant to the outcome described by A but is not conveyed by the probability $P_T(A)$, and is thus not admissible, defeating our use of the Principal Principle when we have such evidence. But on a frequency account of chance, *all* information about probabilities is information, in part, about future frequencies. So information about probabilities is itself inadmissible. In short, a frequency theory T necessarily contains two sorts of information: (a) the value of $P_T(A)$, and (b) further information about A, namely information about the future frequency of type E events—that is, inadmissible information.

Could it be that this frequency information is not information 'over and above' the chances, in the sense (see footnote 5) that the frequency information in T serves only to confirm what the chances say? A certain observation suggests otherwise: given a frequency theory T, $P_T(A)$ is not equal to $P_T(A|T)$. (See Lewis 1994, n. 9, for a worked example.) That is, T contains extra information which alters the probability of A, even though the original probability of A was supplied by T itself. As we will see immediately below, in any such situation, T is inadmissible.

The New Principle

To make use of the Principal Principle, one must have information about chances, but no inadmissible information. When the chances are Humeanly supervenient, this seems impossible, for the information about chances is itself inadmissible.

Lewis, Hall and Thau conclude that the Principal Principle must be replaced by a new rule for setting subjective probabilities that the proponent of Humean supervenience can use. Their suggestion is motivated by the observation of the last section, that the effect of the extra information contained in a Humeanly supervenient T seems to be captured in the conditional chance $P_T(A|T)$. Thus Lewis and Hall replace the old rule for setting subjective probabilities

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

with a new rule they call the New Principle:

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

Lewis says nothing explicit about the role of admissibility with respect to this new rule. According to Hall, however, once we adopt the New Principle we no longer require *any* admissibility constraint on K. As mentioned above, I will later argue that Hall is wrong.

The status of the New Principle is uncertain. There is something quite disturbing about the way that T appears twice on the right hand side, when once ought to be quite enough. Hall and Lewis each put forward a plausibility argument for the New Principle, somewhat similar to that at the end of the last section, but they offer no compelling motivation. What I will argue in the rest of this section is that the New Principle and the Principal Principle are just two aspects of a single rule relating subjective and objective probabilities. This rule, furthermore, is soundly based in our intuitive inferential customs and habits. It is, in effect, the *real* Principal Principle, whereas the rule that Lewis presents in his 1980 paper is just a consequence of that principle that does not capture its full content.

To avoid confusion, let LP denote the rule presented by Lewis in "Subjectivist's Guide" (and described at the beginning of this paper). Henceforth, I will use the name 'Principal Principle' to name the rule from which LP and NP are derived, a rule that states, very roughly, that subjective probabilities may always be set equal to corresponding objective probabilities, subject to the now familiar admissibility constraint.

I will explain the nature of this rule, and make clear the way in which it differs from LP, by looking at our everyday dealings with conditional probabilities. Consider the following case. The probability that the outcome of a die toss is a '2' is one sixth. Intuitively, we are warranted in adopting a subjective probability of one sixth that any die toss results in a '2'. Now suppose that we know that the outcome of some tossed die is even. The probability of a die toss producing a '2', given that the outcome is even, is one third. This conditional probability, like the unconditional probability before it, provides an intuitive warrant for an equal subjective probability. That is, if a die has been tossed, and we know that the outcome is even (but nothing more), then we should set our subjective probability to one third that the outcome was a '2'.

I take it, then, that both intuition and observation of our everyday inferential habits suggest that subjective probabilities be set equal to relevant objective probabilities, whether these objective probabilities are conditional or not. This practice can be broken into two parts, described by two rules:

(LP) $C(A|TK) = P_T(A)$, and

(CP) $C(A|BTK) = P_T(A|B)$,

both subject to an admissibility constraint. Together, then, these two rules capture our intuitive grasp of the relationship between subjective and objective probabilities. (Note that CP entails LP, and that LP entails many applications of CP.[7] I take it that, although these logical relations hold, they play no part in our practice. We do not derive LP from CP, or vice versa. Rather, intuitively, these two propositions have equal status.)

Shortly, I will show that NP follows from CP. First, however, I will discuss the role played by admissibility in the application of CP. Both B and K must be admissible relative to $P_T(A|B)$, and they will satisfy this requirement just in case they tell us nothing about A above and beyond what we are told by $P_T(A|B)$. But B will always be admissible. It cannot contain information about A that is not conveyed by $P_T(A|B)$, because $P_T(A|B)$ must reflect the presumed truth of B, and thus of all the information that comes with B. In effect, then, the only real constraint on the application of CP is the admissibility of K.[8]

Now consider the significance of admissibility when $P_T(A|B)$ is not equal to $P_T(A)$. In such a case CP is applicable, provided that K is admissible. LP, however, is not applicable, because B is inadmissible relative to $P_T(A)$. Why? The fact that conditionalizing on B changes the

---

7 The applications that cannot be handled by LP alone are (a) the derivation of NP, discussed below, and (b) cases in which the conditional probability $P_T(A|B)$ is not derived from $P_T(AB)$ and $P_T(B)$ by way of the usual definition of conditional probability, for example, cases in which $P_T(A|B)$ is given directly by some probabilistic law of nature relating A and B.

8 This is not quite true, since B and T might separately tell us nothing special about A, yet their conjunction BT might (crystal balls are an example). For my purpose, however (deriving the New Principle), this qualification will disappear, for B will be none other than T.

probability of A implies that B contains information not reflected in $P_T(A)$. This is just as well, because the Principal Principle (the conjunction of LP and CP) would otherwise engender a contradiction, as follows.

Suppose that $P_T(A|B)$ is not equal to $P_T(A)$. Then by CP, C(A|BTK) ought to be set to $P_T(A|B)$. But LP tells us to set C(A|BTK) to $P_T(A)$. For example, let A be the proposition that a certain die throw produces a '2', and B be the proposition that the same die throw produces an even outcome. As we saw above, by CP, C(A|BTK) should be set to one third. But by LP, C(A|BTK) should be set to $P_T(A)$, which is one sixth.

Fortunately, as we have just seen, in such cases B is inadmissible relative to $P_T(A)$. It follows that the second of the two inconsistent applications of the Principal Principle (i.e., the application of LP) is invalid. In the example, knowing that the outcome of the die throw is even is inadmissible relative to the probability of one sixth, because it gives us information about the outcome of the toss over and above the information contained in the one sixth probability.

To summarise this section's conclusions about admissibility:

1. We now have a new and much broader test for inadmissibility. If $P_T(A|B)$ is not equal to $P_T(A)$, then B is inadmissible relative to $P_T(A)$.[9]
2. B is always admissible relative to $P_T(A|B)$.
3. CP and LP can never come into conflict. When $P_T(A|B)$ is not equal to $P_T(A)$, B is inadmissible relative to $P_T(A)$, and LP is inapplicable. We apply CP instead (provided, as always, that K is admissible).

Now we come to the denouement of this section. Recall that LP was considered useless to the Humean because Humean theories of chance T are inadmissible relative to the very probabilities $P_T(A)$ that they concern. We can now use another aspect of the Principal Principle,

---

[9] It may seem that we can now define admissibility as follows: B is admissible relative to $P_T(A)$ just in case $P_T(A|B)$ is equal to $P_T(A)$. This is inadequate, however, because some background information may not fall into the domain of $P_T(\cdot)$.

CP, to set a subjective probability for A without assuming that T is admissible. Supposing that K is admissible, and putting T for B in CP:

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

which is of course the rule that Lewis and Hall call the 'New Principle'. So to subscribe to the Principal Principle is also to subscribe to the New Principle. (I should stress the fact that although T may not be admissible relative to $P_T(A)$, it is admissible relative to $P_T(A|T)$, because T cannot contain information about A that is not taken into consideration in $P_T(A|T)$.)

As we have seen, $P_T(A)$ may differ from $P_T(A|T)$. If we ignore admissibility, it appears that there are two rules, Lewis's LP and his more recent New Principle, which are at odds, for LP implies that $C(A|T)$ should be set to $P_T(A)$, the New Principle that it should be set to $P_T(A|T)$. We can now see that there is no disagreement between the two rules. When $P_T(A)$ is not equal to $P_T(A|T)$, T is inadmissible relative to $P_T(A)$, and LP does not apply. In such cases, it is the New Principle that supplies the correct subjective probabilities. LP does not disagree; rather, it exits gracefully by way of its admissibility clause. But LP and the New Principle are not merely consistent—they are part of a greater whole, namely, the Principal Principle, which has guided our inferences time out of mind.

Hall on Admissibility

There is one substantive issue that must be resolved before I evaluate the various claims that have been made about the New Principle. Can we do without admissibility? Hall claims that if we adopt the New Principle, and limit ourselves to historical information, we can. I claim that the New Principle is a consequence of the Principal Principle, and that as such, it is subject to the same qualification concerning inadmissible evidence (as is every appropriately phrased epistemic principle). More exactly, if CP, that is

$$C(A|BTK) = P_T(A|B),$$

is invoked, K must be admissible relative to $P_T(A|B)$.

In this section I resolve this difference between Hall and myself concerning the need for an admissibility constraint. I can make my case by pointing out a flaw in Hall's treatment of crystal balls. Recall that a crystal ball is a device that is able to predict the outcomes of yet-to-be-conducted probabilistic experiments. Information about the image seen floating in a crystal ball is information about a past state of affairs, but also tells us—indirectly—about a future event. On the working characterisation of admissibility, such information was not ruled inadmissible but, as demonstrated above, on the correct characterisation (Thau's) it *is* inadmissible.

Hall asserts that the New Principle is able to deal with the case of crystal balls without resorting to some admissibility constraint. In other words, he holds that information obtained from crystal balls can be considered admissible, and will lead, by way of the New Principle, to appropriate subjective probabilities. Let us formalise this claim. Suppose that we are concerned to set our subjective probability for a proposition A concerning some future event, and that the probability of A, $P_T(A)$, is one half. Suppose also that our (purely historical) background knowledge K includes the fact that a crystal ball has predicted that the event in question will not occur, and that our body of probabilistic laws T tells us that the crystal ball is completely reliable. Now the subjective probability of A, C(A|KT), ought to be zero. Hall wishes to show that the New Principle assigns a zero subjective probability, as logic demands.

I will simplify Hall's treatment in what follows, since the important point does not depend on any of its finer details. Hall notes that KT entails ~A, and argues as follows. Like the Principal Principle, the New Principle has its conditional version:

$$C(A|BKT) = P_T(A|BT)$$

Substituting K for B, we obtain $C(A|KT) = P_T(A|KT)$. Since KT entails ~A, $P_T(A|KT)$ must be zero. So the New Principle sets the subjective probability for A to zero, as desired.

So far, so good. Consider, however, what happens if we set C(A|KT) not by the conditional version, but the ordinary version of the New Principle:

$$\begin{aligned} C(A|KT) &= PT(A|T) \\ &\approx PT(A) \\ &\approx \tfrac{1}{2} \end{aligned}$$

We have a direct contradiction with the result obtained in the last paragraph.

To save the New Principle from inconsistency, Hall must do one of two things. One possibility is to claim that $P_T(A|T)$ is equal to zero. Hall in fact does make this claim, although he does not explain why the claim is important. (Nor does he explain why, though he thinks that $P_T(A|T)$ is zero, he resorts to the *conditional* version of the New Principle to deal with crystal balls.) In any case, the claim cannot be correct. T may tell us that our crystal ball's prediction is absolutely reliable, but it does not tell us what that prediction is. The mere existence of the crystal ball entails nothing about A. Thus $P_T(A|T)$ will be equal to $P_T(A)$. The only way to rescue NP, then, is to rule that K is inadmissible relative to $P_T(A|T)$, so that C(A|KT) cannot be set equal to $P_T(A|T)$. The claim that the New Principle requires no admissibility constraint cannot be maintained.[10]

---

10 One could do away with an admissibility constraint by adopting the following consequence of the old New Principle as a new New Principle:

$$C(A|KT) = P_T(A|KT).$$

No one has made such a suggestion, but in any case, it faces a serious difficulty: some information in K may not fall into the domain of $P_T(\cdot)$.

This is the same difficulty I raised in the previous footnote, with respect to a certain definition of admissibility. If everything did fall into the domain of $P_T(\cdot)$, one could adopt either the new definition of admissibility (keeping the original New Principle) or the new New Principle. But in fact, the one is equivalent to the other. Either explicitly or implicitly, then, admissibility makes its appearance.

The Principal Principle Endures

I will conclude by examining the various positions currently held concerning the status of LP, in the light of the observations and results described above.

My Position: The venerable inferential rule that I call the Principal Principle endures, with Lewis's working characterisation of admissibility replaced by Thau's. The New Principle may then be derived from the old, and will apply to any account of chance for which $P_T(A)$ is in general not equal to $P_T(A|T)$.

Lewis: Lewis agrees that the working characterisation of admissibility should be replaced. However, he would discard the old Principal Principle—or rather, what he takes to be the old principle, namely LP—and replace it with the New Principle. (Thau also wants to replace LP, but does not say what that replacement would be. For this reason, I will not include Thau in the following discussion.) Lewis notes that if chance is non-Humean, and in particular, if $P_T(A)$ is equal to $P_T(A|T)$ for all A, then applications of the New Principle will be indistinguishable from applications of LP (and thus, of what I take to be the real Principal Principle).

Hall: Hall, like Lewis, would replace the old principle with the new, discarding the admissibility constraint in the process.

Of these three, Hall's can be dismissed immediately, for the reason described in the last section: without admissibility, the New Principle is inconsistent. That leaves Lewis's view and my own. On either view, something like LP is used for non-Humean chances, and something like the New Principle for Humean chances. The difference between Lewis's and my views, then, is this: on my view, the New Principle is to be derived from an old and compelling inferential rule, the Principal Principle. On Lewis's view, the New Principle is taken to be foundational.

My view is, I think, to be preferred, simply because it takes the more intuitive principle as the foundational principle. This better reflects our psychology, and better reflects our logic too, if we wish our logical axioms to be what are intuitively the most fundamental principles of our reasoning.

My position does have one strange consequence—that in a Humean world, the Principal Principle, in its simple, intuitive form, is useless, and is replaced by the New Principle, which still looks decidedly odd. (Lewis's view, of course, has the same consequence, and with less motivation.) This peculiarity, however, is not a result of my views about the Principal Principle, but of Humean accounts of chance themselves, in particular, of the fact that on a Humean account, statements about chance necessarily contain information above and beyond chances. For this reason, strange though it seems, $P_T(A)$ is not in general equal to $P_T(A|T)$. It is this inequality, and not the perfectly intuitive CP, that is responsible for the peculiar aspect of the New Principle. (And for this reason, there is certainly no reason to force non-Humeans, for whom $P_T(A)$ is always equal to $P_T(A|T)$, to go by way of that ungraceful principle.)

In conclusion, provided that the New Principle can be seen as a special case of the Principal Principle, there is no reason not to take the old principle as more basic. My view could be undermined only by calling into question the derivation of the new principle from the old. This would involve casting doubt on the notion of admissibility I have been using, the notion, due to Thau, that allows me to argue that T is always admissible relative to $P_T(A|T)$. There is, however, no alternative characterisation of admissibility in sight. Furthermore, if I am right in thinking that admissibility is an instance of a more general constraint on setting subjective probabilities, then it would appear that Thau's characterisation is indeed the correct one.

Suppose, then, that we accept my view over that of Lewis. What has become of the 'revolutionary' overthrow of our original intuitions and habits regarding the connection between subjective and objective probability? In fact, the old regime—the Principal Principle—is as much in control as ever before. What have been presented as innovations are rather recoveries of aspects our inferential customs that were imperfectly described in Lewis's 1980 paper. First, the change in the characterisation of admissibility represents a return to the correct, intuitive definition of admissibility, to which Lewis's working characterisation was only ever an approximation. Second, the augmentation of LP restores a part of our practice—concerning conditional probabilities—that had received inadequate attention in "A Subjectivist's Guide".

What next? There are two major qualities that may be claimed for a rule such as the Principal Principle: that the rule is *intuitive*, and that the rule is *correct*. It seems that the Principal Principle, as I have described it, is indeed the intuitive rule for setting subjective probabilities. It has also survived a modest assault on its correctness; namely, that recently conducted by Lewis, Thau and Hall. However the most important question has yet to be answered: why should we accept that the Principal Principle dictates the uniquely correct way that objective probability constrains our beliefs?[11]

Works Cited

Bigelow, J., J. Collins and R. Pargetter (1993). "The Big Bad Bug: What are the Humean's Chances?" *British Journal for the Philosophy of Science* 44:443–462.

Hall, Ned (1994). "Correcting the Guide to Objective Chance". *Mind* 103:504–517.

Lewis, David (1980). "A Subjectivist's Guide to Objective Chance", reprinted in *Philosophical Papers Volume Two*, 1986, Oxford: Oxford University Press.

Lewis, David (1994). "Humean Supervenience Debugged". *Mind* 103:473–490.

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

---

11 Thanks to Ned Hall, David Lewis, David Papineau and especially Barry Loewer for much helpful discussion concerning this paper.

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