4

Opinions and Chances

I

Ramsey was one of the few philosophers who have fully appreciated the fundamental picture of metaphysics that was originally sketched by Hume. In this picture the world—that which makes proper judgement true or false—impinges on the human mind. This, in turn, has various reactions: we form habits of judgement and attitudes, and modify our theories, and perhaps do other things. But then—and this is the crucial mechanism—the mind can express such a reaction by 'spreading itself on the world'. That is, we regard the world as richer or fuller through possessing properties and things that are in fact mere projections of the mind's own reactions: there is no reason for the world to contain a fact corresponding to any given projection. So the world, on such a metaphysic, might be much thinner than common sense supposes it. Evidently the picture invites us to frame a debate: how are we going to tell where Hume's mechanism operates? Perhaps everywhere: drawing us to idealism, leaving the world entirely noumenal; or perhaps just somewhere; or nowhere. Hume's most famous applications of his mechanism, to values and causes, are extended by Ramsey to general propositions, which to him represented not judgements but projections of our habits of singular belief, and also to judgements of probability, which are projections of our degrees of confidence in singular beliefs.

If we are to assess his views we must be sure of what counts as an argument for or against this projectivist picture. The main burden of my essay is that most ways of framing the debate underestimate the resources available to the projectivist. I think it is also clear, particularly from his 1929 paper 'General Propositions and Causality', that Ramsey himself was optimistic about those resources, in a way that has not been widely recognised.

The usual way of attacking the projectivist is this: He is saddled with a particular view of the meaning of remarks made in the area in question. This view is then shown not to correspond with some feature of the meaning that we actually give to those remarks. It is triumphantly concluded that projectivism is inadequate, and that we must adopt a realistic theory, seeing the remarks as straightforward descriptions of a part of the world that we are (somehow) able to cognize. This kind of attack is clearly worthless unless it is

75

76   *Metaphysics*

clear that the projectivist is indeed committed to the theory of meaning attributed to him. Yet, I shall argue, the theory of meaning to be linked with Hume's picture is variable, subtle, and obscure: if, as I suspect, it is as yet unclear what resources the projectivist has in this matter, it follows that all such attacks so far made are unsuccessful.

A couple of examples may help. It used to be thought that a subjective theory of value entailed identifying the assertion that X is good with the assertion that the speaker himself liked X. This is properly refuted by pointing out that the two have entirely different truth-conditions (or assent-conditions), and the subjectivist is discomfited. But it is now widely recognized that only a very naive subjective theory of value commits this error: a theory of value as a projection of our attitudes can adopt a much better account of what then is said by attributions of value—primarily in terms of expression of such attitudes. Again, it used to be thought (perhaps it sometimes still is) that someone who, like Hume, thinks of the world as a succession of distinct events, and who accords no real distinct existence to necessary connections between those events, must think that we mean no more than regular succession when we talk of cause. But there is no reason for saying that this is what we *must* mean when we project a certain habit of reliance on a regularity, or some other attitude toward it, onto the world. Perhaps, for example, we express some special attitude to the regularity or dignify it in a certain way, and then many stock objections to regularity theories (factory whistles blowing at the same time and so on) are entirely irrelevant. They simply draw attention to regularities that, for some reason yet to be explored, we do not dignify.[1]

But expressive theories of meaning are themselves attacked. It is probably necessary at present to distinguish two kinds of rejection. One, which I shall be coming to, joins issue over some particular aspect of meaning, such as the occurrence of the disputed remarks in subordinate clauses, to which the theory is supposed to be inadequate. The other is hostile to the whole idea of there being a debate. I have in mind the conservative, pessimistic, and perhaps Wittgensteinian view that we cannot do much with our language except speak it, or at best put down rules for buildng up meanings by establishing rules that govern the components of sentences with those meanings. But if that enterprise leaves us with such things as the placid truths that 'good' is satisfied by good things, 'chance' refers to chance, and that A, B satisfies 'x causes y' when A causes B, then we should rest and be thankful. Sometimes, indeed, it is felt that the very endeavour to find semantic structure in a rule-governed way rules out Hume's kind of theory by committing us to a *correspondence* theory of truth, as if coherence theorists, or pragmatists, or the projectivists I am interested in, must half-wittedly deny that 'London' refers to London, and so on. In fact, in constructing such a theory of semantic contribution we use our language simply to describe itself, and leave perfectly untouched the question of which metaphysics is appropriate to that use. It is puzzling to think why people

1. This is explored further in essay 5.

Opinions and Chances   77

still associate the creation of a formal truth-theory for a language with particular views of truth.[2] To exorcise this temptation, imagine a formal theory of Arabic numerals enabling us to deduce which number a sequence of digits refers to, given axioms saying to what individual digits refer; the one essential rule merely captures the way the position of a digit indicates the power of ten by which it is multiplied. Does such a theory, or the interest of creating it, tell us what it is to refer to a number, or commit us to a correspondence theory of such a thing? Of course not. It is entirely silent on the issue, and merely uses the notion, while telling us nothing about it.[3]

I think the argument to the contrary that confuses people goes: a correspondence theory of truth needs to identify some fundamental word-to-world relations; Tarski's style of theory can be taken to offer the relations of reference and satisfaction for this job; hence Tarski's style of theory is, or at least helps, a correspondence theory. This would be fine if there were independent argument that *what it is* for words to refer to things, or things to satisfy predicates, is well thought of in terms of correspondence. But it might just as well be thought of in terms of the predicates taking part in a coherent system; or being used in promoting certain ends; or in terms of the things having the reactions of the mind projected onto them. In other words, we might just as well argue that a coherence, pragmatist, or subjective theory needs to identify some fundamental word-to-world relations; hence Tarski's style of theory helps them too. This shows its irrelevance to this issue. Yet, although it gains nothing from truth-theories, we cannot dismiss the view that Hume's theory is not debatable. The difficulty that perplexes me is that if, as I shall suggest, the projectivist can make perfect sense of apparently realistic practice, it is not clear what intellectual quirks mark him off, nor what is left to fight over except harmless images and metaphors. The interest, at any rate, comes in seeing what he can do by way of incorporating apparently realistic practice: this is a programme that can be called 'quasi-realism', and I see Ramsey as one if its patrons.

Specific charges against projectivist theories will concentrate upon ways in which our thinking about the area in question appears to accord an objective or independent standing to the things allegedly projected. Primarily, chances, laws, and causes (not to mention values and goods) are all things about which, we say, we can be ignorant. Our opinions about them can be wrong, defective, in various ways. We allow the possibility that we think of them as existing when they really don't and that we are unaware that they exist when they really do. We acknowledge experts, so that some person's views of, say, probabilities, become authoritative enough to count as knowledge and to enter into books as physical constants like values of masses and densities (and other people's opinions are often not worth a straw). Yet even the experts might be wrong: it is not their opinion that defines laws and chances; the laws

---

2. M. Platts, *Ways of Meaning* (London: Routledge & Kegan Paul, 1979).

3. This issue was also joined in *Spreading the Word*, chapters 7 and 8. The idea that truth theories dominate metaphysics is not nearly so popular now as when this was written.

78   *Metaphysics*

and chances would have been what they are regardless of whether people had known about them.

Ramsey is usually thought of as one of the fathers of a 'subjective' theory of probability that denies or at best struggles with such facts. On that theory, a distribution of confidence across any totality of propositions is coherent if it satisfies some very weak constraints. But those constraints allow for the most bizarre confidences and agnosticisms. Yet coherence is all that there is. As Kyburg and Smokler put it in the introduction to their 1964 collection, for subjectivism any degree of belief in any statement is permissible, but there are restrictions placed on the distribution of degrees of belief among sets of related statements.[4] Since there is nothing to be wrong about, the view has been summarized as claiming that 'sincerity is enough'. On the more modern version, which I discuss in part III, it is also mandatory to stick by opinions through time and as various kinds of observation are made. However, even this gives us no title to say that a man who announces a quite outrageous set of confidences is 'wrong'; the only vice he could display would be a kind of fickleness as time goes by. It is easy to see why this implausible theory is fostered onto a projectivist. In probability, as in the theory of value, if projection is all there is, there is surely nothing to be wrong *about*. But all this flies in the face of the objectivity of our usage and renders the theory an easy prey to criticism.

But Ramsey was well aware of the shortcomings of a purely subjective theory of laws and chances. He explicitly denies, for instance, that chances correspond to anyone's actual degrees of belief;[5] he knows that we believe in unknown laws (and he would have said the same about chances);[6] he knows that some opinions about chances are much better than others.[7] His effort is to show that these phenomena do not refute an antirealist, projectivist, theory of chances and laws, but actually are explicable given such a theory. It is the fact that he made this quasi-realist attempt that seems to me to show that Ramsey was much better aware of the resources of projectivism than many of his apparent followers (Carnap[8] being an honourable exception). How far can his programme succeed?

II

Hume forged the essential tool for the projectivist to use as he attempts to reconcile his theory with the objectivity of usage. In his great essay 'On the

4. H. E. Kyburg and H. E. Smokler, eds., *Studies in Subjective Probability* (New York: Routledge & Kegan Paul, 1964), p. 7.

5. F. P. Ramsey, *Foundations of Mathematics*, ed. R. B. Braithwaite (London: Routledge & Kegan Paul, 1931), p. 206.

6. F. P. Ramsey, *Foundations*, ed. D. H. Mellor (London: Routledge & Kegan Paul, 1978), pp. 139, 150.

7. Ibid., p. 95.

8. R. Carnap, *Logical Foundations of Probability*, 2nd. ed. (Chicago: Chicago University Press, 1962), p. 16 ff.

Standard of Taste', he points out, in effect, that it is no part of a projectivist metaphysic to claim that one projection is as good as another. Some may be inferior, some superior, and even the best may, in principle, be capable of improvement. Thus, let us take the difficult case of moral evaluations. If values are projections of a habit of forming some kind of attitude to some kinds of thing, how can I be aware that my *own* attitudes might be defective, and capable of improvement? (If we prefer societies to individuals, the question becomes: 'how can *we* be aware that *ours* are'?) The answer is that I know that people are capable of habits of projection which from my own standpoint are deplorable: they judge things of which they are ignorant, and their views are the function of fears and fantasies, blind traditions, prejudice, and so on. But then who am I to be sure that I am free of these defects? This thought is quite sufficient to enable me to understand the possibility of my attitudes improving. They ought to be formed from qualities I admire—the proper use of knowledge, real capacity for sympathy, and so on. If they are not, and if the use of those capacities and the avoidance of the inferior determinants of opinion would lead me to change, then the resulting attitudes would be not only different, but better. It is true that in saying this I am presupposing one kind of evaluation in giving sense to the possible deficiencies of the other. An attitude to the processes of attitude formation is used to give sense to the possibility not merely of change but of improvements in moral judgement. But this gives nothing an axiomatic status: at the end of a process of reevaluation, everything may have changed. The right analogy is with the rebuilding of Neurath's boat, and we know that in principle the result of that might be an improved boat. Equally we can understand and fear the possibility of deterioration. It follows that a projectivist picture of values need have little to do with the frivolities of traditional moral subjectivism ('one opinion is as good as another' and so forth). By pursuing the point we might begin to see how a projectivist can incorporate notions of truth and knowledge.

In the case of empirical judgements of chance the matter is much easier in two respects. The first is not my main concern in this paper, but it is worth noticing. Projectivism in moral philosophy is open to attack on the grounds that the reaction of the mind that is supposedly projected is itself only identifiable as a reaction to a cognized *moral* feature of the world. The specific attitudes and emotions (approval, indignation, guilt, and so on) can, it is argued, be understood only in terms of perception of right and wrong, obligations, rights, *etc.*, which therefore cannot be reflections of them. Myself, I do not think that this is true, nor do I think, if it were true, that it would refute projectivism. For it is not surprising that our best vocabulary for identifying the reaction should be the familiar one using the predicates we apply to the world we have spread. Thus, to take a parallel, many people would favour a projectivist view of the comic, and they may well be right even if our best way of describing the reaction which we are projecting onto a situation we describe as comic is 'that reaction we have when we find something funny'. I don't think a behaviourist analysis is either required or helpful, for obviously the behaviour, to someone with no sense of humour, would be incomprehensible.

80   *Metaphysics*

In any case, a projectivist theory of probability meets no such objection. For it is easy to identify the main component projected when we attribute a good chance to an event or a high probability to a judgement: it is, of course, simply a degree of confidence.

Degrees of confidence in propositions are 'intervening variables' in psychological theory. We can know about them through interpreting the behaviour that, supposedly, they explain. The measurement of degrees of confidence is not necessarily straightforward, but this is no problem. It may even be indeterminate, very often, what a person's degree of confidence in a proposition actually is, or whether indeed he has one. But in the same way it is often not straightforward to know what a person's belief is; it may be indeterminate what his belief about some matter is, and even indeterminate whether he has one. Yet the notion of a belief is a proper theoretical concept in psychological explanation. Here we should notice that a projectivist needs no more degrees of confidence than a person has beliefs about chances; it is no part of his view that, for instance, a real number should be in principle assigned to every proposition an agent has ever thought of, representing his degree of belief in it. We need no such extravagance: often we express ourselves by saying that we have no idea what the chance is—and this attitude need not co-exist with a particular confidence in a proposition. We shall see how to interpret it on a projectivist picture later. If the whole notion of a degree of confidence were suspect, as some authors claim, then we would need at least to indicate another projected psychological state. I am inclined to suggest that it would not be fatal if, as in the moral or even the causal example, our success in doing this without using the vocabulary of chance were only partial. But, in fact, I doubt whether it would be too difficult.

The second respect in which things are easier for probability is the backing we can give to our standards for evaluating projections. We need standards for assessing projections of degrees of confidence, enabling us to say that some are better than others, that even the best may be capable of improvement, that some are worthless. But there are obvious sources for such standards: an opinion might be formed in the light of experience of observed frequencies or fit into an otherwise successful scheme of projections, and most fundamentally it might give its possessor the habit of belief in what happens and disbelief in what does not. And it is this which is the lynchpin of Ramsey's theory. We could say that it imports a pragmatic standard for evaluating projections, but this might be misleading. For it is not as though the standard is in any way optional or avoidable if we adopt different goals or purposes. It is necessary that truth count as success in judgement, and that the proportion of successes achieved by a habit of making judgements be a measure of the confidence that ought to be felt in the beliefs to which that habit leads one. The standard is mandatory.

For suppose we have a thin, Humean view of the world. What is our purpose in projecting onto it chances and probabilities? Ramsey writes that we 'judge mental habits by whether they work, *i.e.* whether the opinions they lead to are for the most part true, or more often true than those which

alternative habits would lead to'. The opinions he is talking about are of course not opinions about probabilities, for that would get nowhere, but the opinions of particular matters of fact that judgements of probability will lead us to form. Fortunately the world displays patterns allowing us to have successful habits of particular belief: faced with partial or complete regularities we can form partial or complete confidence in new cases, and the world grants us success if we are careful. As Ramsey writes, the best habit of belief formation will have us forming confidence of a strength proportionate to the ratio of particular truths to falsities that the habit leads us to believe in. (For why, see section IV.) So it seems that Ramsey is going beyond mere coherence of sets of belief, in a thoroughly sensible and necessary way. That standard is too permissive, since, on the face of it, a set of beliefs may possess the virtue of coherence while having the disadvantage of enjoining confidence that things happen which never do, or that things don't happen which often do. Pragmatism must supplement coherence.

But this charge ignores the work that has recently been done on the relation between subjectivism and 'learning from experience'. This work makes it plausible to believe that a subjective theory of probability, relying only on the constraint of coherence, can show that the process called conditionalization is obligatory. It may then seem as though conditionalizing is itself a process that forces opinions to converge, and that what they converge upon is a value for probabilities in accord with observed frequencies. If all this is true, then the constraint of coherence would be *sufficient* to give the standard of evaluation that Ramsey wants. In the next section I assess this argument, and in the last I go on to develop a projective theory in more detail.

III

In the succeeding sections I shall use the standard terminology, in which a *chance setup* exists, and there is an actual or hypothetical series of *trials* that yield a stable frequency of various outcomes. The set of trials we can call A, and the outcome in which we are interested, B. There is a slight strain in adapting this terminology to, say, the chance of a person being a gin drinker, but for the moment we are not interested in any problems caused by extending the notions.

We now suppose that we have conducted a reasonably extended investigation, and in a large number of trials the proportion of Bs is tending to stabilize around some figure, p. Our problem is to give a projectivist account of the natural judgement that such evidence (in the absence of other evidence) would lead to 'the chance of an A being B is p'. But this way of putting it blurs a vital distinction. There are two sorts of chance judgement that could be made. One is local, and concerns only the chance of one of the trials being B in the set that makes our evidence. The other is not: it is a judgement that concerns other trials and has implications for our confidence in future cases. It

Metaphysics

is one thing to say 'the chance was p of one of the examples of A we have considered having been B'. It is quite another to say 'the chance is p of any A, including ones yet to be realized, being B'. The first judgement is local or restricted to the class of trials in which we already know the frequency. The second involves a prediction, or more accurately, a commitment to a particular kind of confidence in situations not yet brought about. It involves an apparently inductive step. Yet, although it will evidently be more complex to identify the thoughts that license it, I shall consider it first.

It is not surprising that a man observing and recording results from a process that generates a certain frequency of outcomes B among events A, with no discernible pattern, should come to have a degree of confidence, proportionate to the frequency of Bs among As, that an arbitrary A, such as the next one, will be B. But why *should* he? The simplest answer has two stages. First, the man has the inductive habit: he expects the process to go on generating roughly the frequency it has done so far. Second, given that this is so, he will be right to have the degree of confidence in a particular outcome identical with its relative frequency, because that is the standard for rationality. There is nothing mysterious about this second point. If the inductive expectation is right, then the relative frequency remains stable. If that is so, then a habit of adopting and acting upon any *other* degree of confidence in particular expectations would lay you open to certain loss if you are required to act out your confidence by buying or selling bets at the corresponding rate, given that your partner is someone more straightforward. It would be like having confidence other than 0.25 that a card from a shuffled pack is a heart: a hopeless position if you are required to buy or sell bets at a corresponding rate. Nor should we worry that there is anything unrealistic (undemocratic, as it were) in criticizing someone for having a set of confidences that would lead to loss if he were *required* to post odds on which he could be *required* to buy or sell bets: it is not a satisfactory defence to reply that we are not often required to gamble. The defective degrees of confidence are like bad dispositions that may nevertheless remain unrealized. They can still be criticized by pointing out what would happen if they were to be acted upon.

The two-stage answer relies on induction, and the rationality of that is left dangling. It is tempting, therefore, to hope that work on conditionalizing achieves an answer without relying on specifically inductive habits. It is not, indeed, likely that this could be so, since induction appears to be a necessary component of any answer, in that if there were no reason for expecting the process to generate the same frequency as hitherto there would indeed be no reason for expecting the next A to be a B with any particular confidence. So unless work on conditionalization provided some justification of induction, it could not provide the requisite standard.

The central argument in this area is credited variously to David Lewis and Patrick Suppes. It is, in effect, an extension of the standard Dutch book argument for coherence to an agent's probability distributions through time. The standard argument makes coherence at any one time a necessary condition of a rational distribution of confidence. The new argument, which we can

call the dynamic Dutch book argument, or DDB, extends this to prevent a
rational agent from wiping his slate clean at any time and forming whatever
new confidences he fancies (although until Ian Hacking first made it clear, it
was not widely recognized that one might do this, so that subjectivists had
happily helped themselves to conditionalization anyhow).[9] The new argument
is designed to prove the connection between rationality and conditionaliza-
tion. Following Paul Teller we can see the DDB like this.[10] We imagine an
agent with a set of beliefs at time 0, described by a function P_o giving the
confidence with which every proposition in the domain is believed. A change
in belief in A is described as conditionalized upon evidence E if, at time n,
after E becomes known (so that $P_n(E) = 1$), $P_n(A)$ is equal to $P_o(A \& E)/
P_o(E)$. That is, at time n the new confidence in A is equal to the old condi-
tional probability of A upon E. As Teller shows, we can generalize everything
to the case where E merely changes probability, but this does not matter.

Now it is quite clear that, sometimes, changes of belief that are not condi-
tionalizations are legitimate. One may rethink a problem afresh, and come to
regret one's old confidences. One may think up new alternatives. But we can
avoid objections based on this by restricting ourselves to cases where nothing
of this sort occurs, but where someone has *in advance,* at time 0, a settled
policy or habit of not conditionalizing. In other words, one has a policy or
habit that should E come about, will lead him to some confidence in A
greater or less than his present confidence of A upon E. The DDB shows that
a man known to have such a plan, and required to buy and sell bets according
to his confidences, can be made to buy and sell bets on which he has a net
loss, whatever happens, by an opponent who knows no more than him (ex-
cept, perhaps, that he has the habit or policy). The general proof is complex,
but its principle is quite easy to grasp. Suppose I am following a plan, or have
a habit, which means that I now have a large confidence that A will occur if E
does, but which enjoins that if E does occur I will only have small confidence
in A. Suppose I think there is a 60 percent chance that John will play a spade
and a 90 percent chance that if he does so he will be left with a court card.
But, flouting conditionalization, I am settled that I will have only 30 percent
confidence that John will have a court card after he has played a spade. The
nub of your strategy for profit is this. You sell *to* me a bet to yield (say) 1 if he
plays a spade and has a court card, and 0 otherwise. I will pay a relatively
large amount for that (0.54, in fact). You plan that *if* he plays a spade you will
cover that bet by buying *from* me a bet to yield you 1 if he has a court card, 0
otherwise, and you know you will be able to do that cheaply, for since I will
then have little confidence in the court card, I will want little for such a bet
(0.3, in fact). You then only need to arrange side bets to give you a modest
profit if he does *not* play a spade, and you will profit in any event. If you can
sell me bets when I am confident, and buy when I am less so, you profit.
Conversely, if I had announced that although now I regard it as only 30

9. I. Hacking, 'Slightly More Realistic Personal Probability', *Philosophy of Science* 1967,
pp. 311–25.

10. P. Teller, 'Conditionalization and Observation', *Synthèse* 1973, pp. 218–58.

84   *Metaphysics*

percent probable that if he plays a spade he will have a court card, yet, I agree, if he does play a spade I will be very confident that he has a court card (I know I get excited). You buy *from* me a bet on both things now, and sell *to* me a bet on the court card later, if he plays a spade. Again, arranging side bets to cover him not playing a spade, you profit whatever happens. I ask little now for the first bet, and am prepared to pay a lot later for the bet which covers it.

It is not quite right to say that this gives an effective method of profiting from a nonconditionalizing agent. The direction of his departure from the present value of conditional probabilities (the probability of *A* upon *E*, now) must be known. If we know that he will inflate his confidence in *A* we can profit, and if we know that he will deflate it we can profit, but it does not follow, and it is not true, that if we know that he will do one or the other we can profit. We have to know the direction of his aberration before we know whether to buy or sell bets, but of course a general tendency in an agent one way or another could also be exploited over time.

What does this argument show? It shows that an agent known to plan a definite confidence in some proposition receives certain evidence that is either higher or lower than the value he now attaches to the conditional probability of the proposition on the evidence, can be made to lose, whatever happens, if he is required to act out those confidences. Let us agree that such a plan is irrational. Does it follow that we should expect rational confidence to converge upon frequencies, thereby bypassing the apparently inductive step? The feeling that it may have something to do with it comes like this. Suppose we antecedently hand people a number of hypotheses about the chance of an *A* being a *B*, and invite them to form a distribution of confidence among them avoiding the pathological values 0 or 1. We then amass frequencies, and since the agents must conditionalize, the posterior probabilities gradually increase for hypotheses giving the chance a value near the observed frequency, and fall away for the others. Eventually opinion converges upon a high probability for the chance being as near as possible to the observed frequency. And the value given the chance dictates our confidence in the next *A* being a *B*.

As an attempt to either bypass or cast light upon the inductive step, this argument clearly fails. It falls to a dilemma: either the original hypotheses are consistent with changes in chances over time, or they are not but relate solely to the trials already conducted. They are local in my sense of the term. If they are, then whatever our confidence that the chance of an *A* being a *B* took a certain value in generating the frequencies we have so far observed, we need inductive confidence to transpose that confidence to the future. It is no easier to argue that since the chances of an *A* being a *B* have always been good, they will continue to be, than it is to argue that since nature has always been ordered, it will continue to be so. On the other hand, if the original hypotheses describe eternal chances, so that in accepting a hypothesis concerning chance I would indeed be committing myself to a uniformity—to the probability that a stable observed frequency of *A*s among *B*s can be extended indefi-

nitely with a similar value—then the induction is presupposed in setting up the request; it is only if we have an inductive faith in such uniformities that we should be inclined to distribute confidence over the initial finite selection of hypotheses. Otherwise we should simply point out the many other things that nature might do instead of giving a uniform chance for an *A* being a *B*.

This is easily seen if we imagine a concourse of souls in an antecedent heaven, each of whom is handed a ticket describing a different course that nature might take in the world they are about to enter. Some give eternal constancy to the relation between *A*s and *B*s, some see it as altering over time, so that although up until a certain time a certain degree of confidence in an *A* being a *B* accords with frequencies, after that time a different one does. These souls may dream up conditional probabilities for themselves, giving the world an *x* percent chance of continuing to conform to their ticket if it does so until a certain time. But unless they have an *a priori* reason for expecting the world they are to enter to favour uniformities, there is no reason for *x* to be different for those with straight tickets and for those with bent ones. And if all possible hypotheses are ticketed, then *x* must be vanishingly small, corresponding to the fact that at any time there will be an indefinite number of tickets conforming to the world up until that time, but subsequently divergent. Of course, I am not here denying that we may be able to think of a reason why *x* should be placed higher for those with straight tickets: that is solving the problem of induction. But it is quite clear that if this can be done, it is not by simply proving the virtue of conditionalizing. For *that* is something which the bent can do with the straight. In a nutshell, conditionalizing appears interesting only if we pose the problem in a way that presupposes inductive good sense. (Philosophers of a sociological and Popperean bent are liable to point out that only hypotheses about stable chances would be of interest to scientists, who would regard the others as crazy. This is true. It is true because scientists, like the rest of us, possess inductive good sense.)

It appears, then, that hypotheses about chances that carry implications for future distributions of confidence are not automatically the outcome of conditionalizing changes of opinion. Induction is needed, as indeed we might have expected. Furthermore, we can escape the unrealistic element of seeing learning from experience in terms of conditionalizing: namely, the nebulous nature of the prior distributions of confidence needed. It is usually much more natural to see our experience as putting us in mind of some hypothesis about chance, rather than merely modifying the degree of confidence with which we used to hold one. The rational man does not have to spring from the womb fully armed with an infantile probability distribution across all the hypotheses that experience teaches him to believe. But if we now turn to purely local assessments of chance, things may appear easier. Here we merely want to say that the chance of an *A* being a *B* on the trials we have conducted was *p*. Whether it remains so can be, so far as that judgement goes, entirely up to the gods, and depending on our opinion of them we can expect what we like about the next *A*. What is the rationale for such local judgements?

The connection between judgements of chance and confidence that we

86   *Metaphysics*

have relied upon has been quite simple. We have imagined judgements of chance dictating a corresponding degree of confidence: the reason why that must conform to relative frequency is evident if we demand that the confidence could be acted upon in willingness to buy and sell bets. If we imagine a closed set, like a pack of cards, any confidence in an outcome (e.g. a card being a heart) with a known frequency must conform to that frequency, for otherwise, if the confidence were acted out, loss would be certain. So it might seem as though everything ought to be very simple for local judgements of chance. We know the frequency of As in our set being Bs; the right confidence to have that an arbitrary member of the set is a B must conform to that frequency, so the right chance judgement must be one that expresses that degree of confidence, that is, the judgement that the chance is identical with the frequency. But the trouble is that although this may be correct in the sense that, if we knew nothing else, it would be the proper estimate of the chance, it does not follow that we regard it as true.

For there is actually no compulsion on us to identify local chances with local frequencies: we all know that there is a chance that in any finite set of trials the frequency with which an outcome actually occurred differs markedly from the *chance* of that outcome occurring. The chance of getting a B might have been q, even if the obtained relative frequency of Bs was p. Here realism seems to triumph: what kind of account can a projectivist give of this modal claim? Equally we may persist in actually believing that the As in our set had a chance q of being B, although we know that the relative frequency was p, so if we were told to bet on whether an arbitrary A had been a B, p would be the right figure to act upon. Again, what account can a projectivist give of the belief about chance, when, as in this case, it appears to diverge from the right degree of confidence to have about arbitrary members of the set?

The answer can come only from seeing the parts that induction and science play. If we think the chance of an A being a B was actually q, we think that q would have been the right betting rate antecedently to the set of trials. We can believe this because we can believe that there was or could have been a longer set of trials in which the proportion of Bs tended to stabilize on the correct figure, q, and we believe that none of the things that affect such frequencies was different on our actual trials. This last is a scientific belief, in the sense that our causal theories of the world are what tell us whether particular factors that do influence the frequency of Bs were present. Of course, if the figure p arose from a sufficiently long series, this in itself will be evidence that such a factor was present, whatever it may have been. But the point is that we may not be forced to think that, and it may be easier to believe the reverse. The judgement of local chance, when it diverges from actual frequency, is then an expression of the confidence that should be felt in a hypothetical situation: a situation where none of the things that, we believe, affect frequencies of Bs among As would be different from those which obtained on our trials. Similarly the modal claim expresses our fear that the hypothetical series might exhibit the different frequency q, and our actual set of trials may be a very poor indication of it. However, there is no reason to be depressed by these

Opinions and Chances   87

possibilities, for induction works, and what produces a frequency of outcomes one day is very likely to do substantially the same the next.

It seems to me that this is an account of our apparently realistic talk (in this case, modal claims and claims apparently divorcing local judgements from expressions of confidence) that yet concedes nothing to realism. Chances have not entered as real facts, capable of explaining or causing events. They remain projections, even if we are interested in spreading them not only over the actual events that have confronted us or that will confront us, but over events that could have happened as well.

In my last section I want to take up another problem that may tempt us to realism: the problem of our willingness to talk about knowledge of chance, and our subtlety in so doing.

IV

Let us suppose that the best evidence is that the frequency of Bs among As indeed approximates to p, and will keep doing so. Clearly, then, a man with the degree of confidence p has the habit of singular belief that meets Ramsey's standard: he has a degree of confidence proportionate to the number of times he is right. But there is another standard needed. For it does not follow that we should endorse this confidence, nor this judgement of chance. Consider that we may be able effectively to divide the As into two classes and rightly predict a high ratio of Bs among the CAs and a low ratio among the others. I shall call this 'effecting a partition of the As'. If this were done we would be in a better position than him in this very straightforward sense. Using our knowledge we can gamble with him and consistently win, by buying bets at a rate corresponding to p on CAs being B, and selling on $\sim CA$s. And in a variety of less mercenary ways we can see that our habit is more useful than his. It is more accurate, more efficient. It is not necessarily incorrect to say that the chance of an A being a B is p, even when effective partitions exist. We can, after all, talk of the chance of an animal being a carnivore, or of a human baby being Chinese. But if an acceptable practical application is to be made of such a remark, then the context must be one in which the participants themselves cannot easily put the subject matter into one or the other partition. If they could, then there would be something deficient about the distribution of confidence: it could effectively be improved. We might express this by saying that there is no case of an A, neither those which are C nor those which are not, on which we should accept such an estimate of the chance of it being a B. But if an acceptable single-case judgement is to follow from such a remark, then the conversational context will be one in which the participants cannot easily put the subject into one or other of the partitions. It remains true that if this can be done, the judgement of chance is deficient. (Of course, there is no implication here about negligence or otherwise in being ignorant about C.)

We here have the beginnings of a reasonably clear view of two topics that sometimes perplex analyses of probability: the rationality of seeking the nar-

88   *Metaphysics*

rowest reference class when we want single-case judgements, and the propriety of restricting the terminology of chance to phenomena which satisfy von Mises's second condition for an empirical collective, namely that there should be no effective selection procedure for singling out a subset of the members with a different overall frequency of the relevant property. (This is felicitously called 'the requirement of excluded gambling systems'.) To take the second issue first, if we believe that a selection procedure (corresponding to the property C in the last paragraphs) can *easily* be found, then it follows that we believe that our distribution of confidence can *easily* be improved, by the standards we have seen. While we believe such a thing we are obviously in a deficient position, and, depending on the consequences of a judgement and the degree of ease of improving our knowledge, we may wish to attach no weight to our judgement—to suspend it, in fact. The projectivist thus has excellent pragmatic reasons for confining our judgements of chance of an A being a B to reference classes for which we do not expect an easy method of partition. There is no metaphysics of randomness required. It is just that if we cannot effect a partition, a judgement of chance leads to our best possible habit. If we can, it does not.

The single-case problem is in effect the same. If you want to have a degree of confidence in the judgement of a particular philosopher that he drinks gin, the statistic of the proportion of men who drink gin helps. But that statistic can easily be refined: middle-class people, academics, academics in arts subjects, of a certain age . . . The pragmatic motive for seeking the narrowest reference class is just the same as before: a degree of confidence based solely on the wider statistic leaves its possessor in an inferior position vis-à-vis someone who can partition the class of men and attach different degrees of confidence to singular judgements depending upon the subclass in which the subject is found. Again, of course, the extent to which it is worth seeking statistics for narrower reference classes depends upon the expected benefit of the more discriminating judgement and the expected homogeneity of the reference class. If nothing much hangs on it and if our prior judgement is that gin drinking is not likely to be much different in any subsets we can think of, it may not be worth a research programme to find out. Nevertheless the point remains that unless we make our evidence as weighty, in Keynes' sense, as we can, by considering possible partitions of our class, we cannot be sure that we are properly serving the purposes of judgement, and the right think to do may be to form no opinion. A good example of such reticence comes from legal suspicion of a 'mere' statistic putting a defendant into a class with a high frequency of guilt. Unless the class is as weighty as it can be made, it would be impermissible to be confident that the defendant is guilty.[11]

It seems to me that this pragmatic perspective on these issues has a clear advantage over a realist metaphysics. Notoriously, in trying to give sense to the single-case judgement, a realist metaphysic of chance becomes tangled in

11. L. J. Cohen, *The Probable and the Provable* (Oxford: Clarendon Press, 1977), esp. chapter 7.

the issue of determinism. If it is determined now that I will, or will not, catch a cold next winter, it is hard to see what sense to make of talk of my propensity to do so or not, just as, if I did not die before I was thirty, it is very unclear what could be meant by saying that I had a propensity to do so. Yet I had a chance of doing so. Equally, only if a chance setup is indeterministic will it be, for a realist, true that trials on it form a collective. Otherwise, it is in principle possible to select trials with a different long-run frequency of outcomes. For the projectivist this becomes simply irrelevant. Our purpose in making judgements of chance, and the standards to use in following out those purposes, are perfectly indifferent to whatever secret springs and principles lie behind the empirical collectives that form the subjects of those judgements. When our best judgement is that we should treat the trials on some setup as forming a collective, we are entitled to project a chance and form confidence accordingly; when we suspect that we can find a partition, we should not. But there need have been nothing wrong with us if we have treated something as a collective, but at some improved state of knowledge a partition is found. Not all ignorance is culpable.

The realist is apt to be impatient with such *laissez-aller* attitudes. True, he will say, you can tell us when we talk of empirical chance, and perhaps you can give some pragmatic understanding of why we do. But for all that, we may be *wrong* to do so in some cases even when we are not culpable. Warranted assertibility is not truth. A later discovery that a trial on a chance setup A fell into a subclass with a different ratio of Bs from that shown overall would show that we had spoken falsely; but what made our remark false was not the discovery but the fact about the trial. Unless indeterminism is true, there are always such facts, in principle, forcing falsity on all attributions of chance other than 0 or 1. This argument is all the more persuasive because it cannot be avoided by mentioning warranted assertibility in the long run, the usual pragmatist substitute for truth. Once we admit that there are facts determining whether a particular outcome will occur or not, we cannot very well claim that a long run of improving investigation would not find them; and in that case determinism will entail that all chances of particular events are 0 or 1, even on this definition.

But it is the definition that is at fault. A projectivist need have no use for truth, about the chance of a single case, as the limit of the degree of confidence to which progressive omniscience would tend. He *does* need a proper account of fallibility, enabling him to admit that a particular estimate of chance might in principle be improved. This we have provided him, without any involvement in the metaphysics of determinism. This means that along with a judgement he has the concept of a standpoint yielding a possibly improved judgement. It does *not* mean that he needs the concept of a standpoint from which *all* other judgement is seen to be wrong, and this is what the unnecessary notion of the limit is attempting to import.[12]

In practice this means that although we have plenty of use for warranted,

12. See also essays 1 and 2.

90   *Metaphysics*

defensible, careful, estimates of chances of particular events, we have less use for claims that we know such chances. The claim to knowledge entails, I think, the claim that no improved standpoint, yielding a revised estimate, is possible. To know something is to know that no judgement contradicting one's own could be really preferable. To know that the chance of an outcome on a single trial is p we would need theoretical knowledge that no partition exists. thus we are entitled to say that we know the chance of an individual outcome on an individual trial only if we are entitled to say that we know that the chance setup admits of no partition. Now we can adopt different standards for saying this. The clearest case (where, as it were, even God cannot partition the trials) is one where we know that the system is indeterministic. Yet our standards need not be so absolute. We may know that whatever God could do, there could be no practicable project of partitioning the trials. The systems on which people gamble are designed so that this is so. In such a system we have effectively ruled out the possibility of an improved judgement about any single trial, and we can properly express belief that we have succeeded in doing this by claiming knowledge of the chance. In other cases, responsible judgement is all that we want or need.

How far can our quasi-realism, our attempt to found apparently realist practice on a subjective basis, succeed? A possible stumbling block would be talk of chances as explaining events, or as themselves being things that need explanation. This would seem to invest them with some ontological standing, with a real influence on the world, which fits ill with the projectivist picture. But perhaps the appearance is deceptive, for everything will depend on the interpretation we make of such explanatory claims. Clearly, if we have a generalized antirealist attitude toward science, then chances could be as honest inhabitants of theory as any other posit can be. But even without this, once we have incorporated the notion of truth, or right opinion, there will be natural things to ask and say about why chances are what they are. We can ask what it is about the world that makes it the case that one particular distribution of confidence over propositions about some subject matter is right. And we can, in asking why something is the case, cite that some such distribution is right as part of our answer, even if this looks dangerously like giving some chances a real, causal, place. It need not be doing so, for instance, because if we endorse one distribution (say, give a 0.5 chance to heads in a coin toss) it will standardly follow that we ought to endorse others (a 0.25 chance of two heads on two trials) and it is no surprise that if asked why this is the chance on two trials, we reply by citing the first 'fact'. This is not the place to enter into all the moves a projectivist might make in tackling explanatory contexts, but their mere existence is unlikely to be much of a problem for him. The empirical part of science connects frequencies with whatever factors influence them, and our reaction to this knowledge is our talk of chance. If, in expanding such reactions, we find ourselves talking of chances explaining things and needing explanation, the proper response is to ask what we are projecting onto the world by making these remarks.

A more serious threat is that with its very success quasi-realism takes much

of the impetus out of subjectivism. Responsible subjectivism is less fun. If, one by one, a quasi-realist programme takes over the things that realists used to consider their special private property, then the view that we have real metaphysical options becomes more doubtful. We all thought we knew what we meant by subjectivism, as opposed to propensity theories, frequency theories, and so on. But, if I am right, the intellectual practices supposedly definitive of these different positions can be available to all, so that the old definitions and divisions appear quite artificial.

In any case, if there is a metaphysical issue, then the subjectivists, so far as we have yet discovered, may have been right about it. But the subjectivism is not the irresponsible and therefore inefficient brand with which Ramsey is wrongly associated. It maintains standards for proper projection, and those standards go beyond coherence and beyond the dynamic coherence that I discussed. They involve a proper respect for frequencies, arising from a proper respect for induction. The main consequence of this responsibility is that subjectivism, as a metaphysic, becomes immune to a large variety of abuses. It also finds itself able to give sense to most of the thoughts that tempt us to realism. However, it maintains what can be seen as clear gains: it avoids the metaphysical problems of indeterminism, since proper single-case projections can be made when we are perfectly indifferent to that issue, and by trying to purchase realist practice from a more austere metaphysic we may come to feel more secure in that practice. There would, of course, be much more to be said if we were to expand the quasi-realist programme, particularly to cover the hypotheticals that are involved in the modal claims I discussed. But my instinct is that if there are obstacles on this route, then they afford opportunities for delightful scrambles, rather than excuses for retreat. And I think that was Ramsey's view as well.

Addendum

It may be instructive to add some notes comparing the resolutely antirealist approach sketched in this paper (and of course much more technically elaborated in works such as Jeffrey's *The Logic of Decision* (Chicago: University of Chicago Press, 1983) and Skyrms's *Choice and Chance* (Encino, Calif.: Dickenson, 1975) with the eclectic theory of David Lewis's 'A Subjectivist's Guide to Objective Chance'. Lewis believes that the practice of science requires, in addition to measures of reasonable partial belief, a notion of 'objective chance'. However, our best grip on this notion is given via our conception of what reasonable credence concerning it amounts to: the Principal Principle. This identifies the chance of an event at a world and a time as the unique degree of belief in the event that would arise from any reasonable initial degree of belief, by conditionalizing on the complete history up to the time and 'theory of chance' at the world.

This is the upbeat side (I owe the description to James Logue). It seems, initially, grist to the antirealist mill, since direct evaluations of objective

92   *Metaphysics*

chance now turn out to be the kinds of thing a subjectivist can make—provided, of course, he can access the history and the 'theory of chance'. Objective chance might stand to confidence like values stand to attitudes. In the philosophy of value we might say that the objective value of a state is or reflects the attitude you would have to it given the complete facts ('history') and theory of value. Provided that having a theory of value is something like having a sensibility, that is, a fact-in/attitude-out grinder (or having such a grinder plus proper concern for coherence and attitudes of fallibility, describable by the antirealist), all is well. It would not be well if having such a theory were having a description of a particular part of the way of the world, because the antirealist will resist the metaphysics that this seems to be bringing with it.

Lewis does not see it like this. He wants to leave objective chance as a real aspect of things, and it is instructive to see the difficulties this causes (the downbeat side). These centre around finding a coherent status for the theory of chance, or the 'history-to-chance' conditionals. Lewis himself appears to vacillate over the crucial feature of this theory already mentioned, namely whether we are to think of it as *true* or not. When it first comes in 'it may or may not be a complete theory, a consistent theory, a systematic theory, or a credible theory. It might be a miscellany of unrelated propositions about what the chances would be after various fully specified particular courses of events. Or it might be systematic . . .' (p. 96). But it rapidly transforms itself into the theory that *holds* at *w*, or in other words, the truth about how history determines chance at *w* (p. 97). The transformation is critical for the subjectivist versus objectivist opposition, of course. In the former guise the theory is naturally thought of as the subject's own set of procedures or norms for distributing confidence in the face of empirical knowledge; in the latter guise it seems more like an objective, potentially unknown and even unknowable specification of how real chances relate to real events in the particular world. Only in the latter guise does it sustain Lewis's realism about chance. It is also worth remarking that only in the latter guise does it seem the kind of thing chances might supervene upon. We do not normally think of values, for instance, as supervening on our own dispositions to value this or that, and it can falsify practical reasoning to think of them doing so (see chapter 9 for more on this. The theme is usefully taken up in Philip Pettit and Michael Smith, 'Backgrounding Desire', *Philosophical Review*, 1989).

The difficulty that I interpret Lewis's paper to be exploring is not this, but rather the problematic modal status for history-to-chance conditionals. Either they are necessary, or they are contingent. If necessary, then any two worlds sharing the same Humean mosaic (or even the same historical Humean mosaic, or set of admissible events) share the same chances. If contingent, what are they contingent upon? Lewis dislikes the first option because he cannot see belief in such conditionals as rationally compulsory. But we are fairly familiar with the idea of necessary truths that are not rationally compulsory. My own reason for disliking it would be that it appears to fly in the face of the methodology implicit in judgements of chance. We make such judgements in the light of finite evidence in accordance with various methodologies from the

Opinions and Chances   93

family of laws of large numbers. When a finite pattern is interpreted as a Bernoulli sequence in which each event has a chance p, we understand that the same pattern *could* have been the outcome of a different chance set up: a non-Bernoullian sequence in which chances change, or a sequence in which every event had a somewhat different chance than p. The history-to-chance conditionals are not treated as necessarily true. This is not the 'assertibility condition' for chances.

But if contingent, contingent upon what? The question will not trouble an antireductionist realist, content to see chance as another self-standing domain of fact, which may or may not vary with other things. It may not trouble a realist roughly of Armstrong's type, content to think in terms of universals and relations between them forming self-standing 'gridlocks' or instructions for the ways things fall out. But it rightly troubles Lewis (who doubts whether these theories can give any account of the Principal Principle, which sums up all we take ourselves to know about chance), and it troubles anyone who thinks that the Humean mosaic is all that there is. Perhaps one might think this but also think that a chance exists (how?) of there being other things that would in turn have altered the chances at the Humean world. But let us call a resolutely Humean world one where the mosaic makes up all that there is and *all that there is any chance of being*. What can history-to-chance conditionals be contingent upon in a resolutely Humean world? How could they be true in one such mosaic and not true in another identical mosaic? And if they could not so vary, we seem to be back with the necessity of history-to-chance conditionals.

The dilemma is sharp. But it should trouble Lewis only because of his realist leanings. If we jettison those we can give a perfectly good description of the way in which norms governing belief in chances should be governed by patterns in the Humean mosaic, the way in which those norms are neither rationally compulsory, nor, in the light of the large-numbers methodology, thought of as necessarily yielding the truth (the last word about chance at a world). There is no final puzzle to dispel over how chances ('the things themselves') can be contingently related to the mosaic: finding that intractable is a symptom of the truth-conditional hankerings I tried to diagnose for modality in essay 3. (My thanks to David Lewis for extensive reaction to an earlier commentary on his paper, and especially to James Logue.)