An Epistemic Theory of Objective Chance

Richard Johns
Department of Philosophy
University of British Columbia
johns@interchange.ubc.ca
January 26, 2007

ABSTRACT

A theory of objective, single-case chances is presented and defended. The theory states that the chance of an event E is its epistemic probability, given maximal knowledge of the possible causes of E. This theory is uniquely successful in entailing all the known properties of chance, but involves heavy metaphysical commitment. It requires an objective rationality that determines proper degrees of belief in some contexts.

I. INTRODUCTION

In this paper I will develop and defend a propensity theory of objective chance, which one might be tempted to call a subjective theory of objective chance. (I will resist this temptation, however, and call it the *epistemic* view of chance.) Physical chances, on this view, have all of their generally-accepted properties, so that I am not offering some mere epistemic surrogate for chance, but the real thing. Chances are, in fact, rational degrees of belief (in the right epistemic state).

Immediately one might think that this third view is a contradiction in terms, that when we talk of *objective* chance, or propensity, we mean something that is *not* a degree of belief. I think this would be a mistake, however. As I argue in the next section, the objectivity of chance consists of certain ideas that are not only consistent with the epistemic view, but *entailed* by it. Moreover, these platitudes about chance are not entailed by any other theory of chance, so that the epistemic view is presently the only explanation of them.

I realise that the epistemic view of chance will initially seem implausible to many, but this may be mitigated by the following two clarifications.

First, one must be careful to distinguish between the physical tendencies themselves and the numerical measurement of them. (The term 'propensity' is unfortunately used for both.) Consider, for example, the differences in stability between isotopes. If you buy some copper piping at a hardware store, it will be mostly composed of copper 63 and 65. It will contain very little copper 62, as this isotope decays so quickly. Copper 63, on the other hand, lasts indefinitely. We say that nuclei of copper 62 have a much stronger tendency to decay than those of copper 63. This difference in tendencies is certainly not epistemic, since it results somehow from a difference in the number of neutrons per nucleus: copper 63 has 34 neutrons, one more than copper 62 has.

Chances are not physical tendencies, however, but numerical *measures* of tendencies. By 'measure' here I mean a real-valued function that represents the tendencies, which moreover satisfies the axiomatic requirements for a normalised measure, or probability function.[1] This measure is indeed epistemic, in the sense that the numbers are degrees of rational belief.

A comparison with secondary qualities may be helpful here. The redness of a ripe tomato is in some sense an objective feature of it, depending on chemical bond resonances and so on, but is also dependent on the observer. (A red object is roughly one that appears red to a standard observer under the proper conditions.) In a similar way, the tendency of copper 62 to decay rapidly is an objective feature of it, but the chance of decay is roughly the degree of belief in decay that a perfectly rational person would have, given maximal information about the causal context.

A second clarification is to distinguish between a quantity being observer-dependent and its *varying* between observers. For chance, on the epistemic view, is observer-dependent but does not vary between observers. How is this possible? It results from the "observer" in this case being a rather ideal one, who is not only (perfectly) rational but who also has maximal information about all the factors that might help to

[1] A normalized measure function is roughly one that is closed under Boolean operations, is additive, maps the empty set to zero and the sample space to unity. It is a purely structural property of some functions.

cause the event in question. In this way, the inevitable differences between human observers in knowledge and in quality of reasoning are beside the point.

This position does require, of course, that rationality be a feature of the world that goes beyond human culture and human physiology. The half life of Copper 62 is about 9.74 minutes, I'm told, and I doubt that this number somehow depends on Western civilization or the layout of the cerebral cortex. (The value 9.74 minutes as a measure of the tendency of Copper 62 to decay is surely independent of humanity altogether, beyond trivial matters such as our definition of one minute as one 1440^{th} part of one earth day.) Such a high view of rationality is not uncommon among analytic philosophers, however. Logicians are still reluctant to see logical truths such as *modus ponens* or the laws of probability as socially constructed or dependent on human evolution. Thus I do not see the requirement for such an independent rationality to be much of a drawback for the epistemic view. Instead, I see it as a useful argument against those who would make rationality a mere human construct. I am happy to say that such a claim is hard to square with the existence of objective chances.

II. WHAT DO WE KNOW ABOUT OBJECTIVE CHANCE?

There are some beliefs about physical chance that are very widely (and reasonably) held, by physicists as well as propensity theorists, so that any theory of chance might be expected to account for them. I will call them the Basic Assumptions. They are as follows:

(1) Chance is a normalised measure.

(2) Chance measures the strength of tendency of a type of event to occur in a specific context. A chance of zero represents an (almost) physically impossible event, and a chance of one means that the event is (almost) physically necessary.

(3) Knowledge of the chance of an event E (in some epistemic states where one knows suitably little else about the experiment) rationally obliges one to have an equal

degree of belief that E will occur. (This was named the 'Principal Principle' by David Lewis.)

(4) The chance of an event E is determined by its causal context, or "generating conditions", so that in identical experiments, where all the possible causes of E are exactly replicated, it always has the same chance of occurring.

(5) Causally-independent events have chances that are probabilistically independent, at least for 'classical' systems. For quantum systems, there seem to be exceptions to this.

(6) The stability of long-run relative frequencies in repeated identical experiments (such as in quantum mechanics) is correctly explained by positing independent, identical, single-case (iid) chances, from which the (probable) existence of stable long-run frequencies can be mathematically derived, using the probability calculus.

(7) Chance can be estimated (approximately and fallibly) by measuring those stable long-run relative frequencies in repeated identical experiments.

These Assumptions are not of course logically independent. David Lewis has shown, for example, that (7) follows from (1), (3), (4) and (5) together.

Looking at 1-7 above, we can see why chance is thought to be an objective quantity, not a subjective one. (4) for example says that the chance of E depends only on the causal context, i.e. the physical events and laws that might cause E, and not on anyone's beliefs about those causes. Also, the fact that chances can be empirically measured (7) surely entails that it is an objective quantity. On the other hand, it is easy to derive all of 1-7 from the epistemic view of chance, as is shown below.

Assumptions (4) and (5) are not often stated explicitly, but I think they are accepted by virtually all physicists. The main difficulty in expressing them is that they involve the notion of causation, which is understood in so many different ways. In particular the common idea, that a cause is something that raises the chance of its effect, cannot be the variety of causation intended here. For if that were the meaning of 'cause' in this context, then (5) becomes nonsensical. The two senses of 'independent', causal and probabilistic, become almost synonymous, so that the first part of (5) becomes a tautology and the second part a contradiction. Assumption (4) also makes little sense from this point of view, as the causal context should surely include any event whose occurrence *lowers* the chance E, as well as those that raise it.

If 'cause' is not being used in its probabilistic sense here, then what is meant by it? I mean something like the notion of 'derivativeness' that Elizabeth Anscombe draws attention to in her 1971 inaugural lecture *Causality and Determination*. For convenience, I will refer to this kind of causation as *physical causation*, following Dowe (2000), although I do not share Dowe's view of this relation. For contrast, I'll call an event whose occurrence raises the chance of E a *positive causal factor* for E.

To see the important difference between physical causes and positive causal factors one can consider examples where the chance of an event E is raised by some event that ends up not being physically connected to E, and which hence does nothing to bring E about in the physical sense. For a simple example of this, consider a projectile that is buffeted by random forces, so that the point where it will hit the ground is not causally determined at the time it is launched. Let the event E be that the projectile fails to strike its target. A second stochastic projectile is aimed at the first, so that the two might collide. (Whether they will collide isn't predetermined, due to the random forces acting on them.) If they collide, then the first projectile has practically no chance of reaching its target so that E is practically certain. It is clear, in this case, that the mere *firing* of the second projectile raises the chance of E. Now suppose that the projectiles fail to collide, yet (by luck) E occurs anyway. (Assume that they causally interact only by collision.) The outcome E in this case was not caused, in any way at all, by the firing of the second projectile, even though the firing increased the chance.[2]

Many authors have also noted that one can alter a causal process, introducing new elements, or replacing old elements with new ones, in such a way that (i) the chance of some event E is thereby *lowered*[3], or left unchanged[4], and (ii) those new elements might

---

[2] In another context, Harry Frankfurt has presented other well-known cases where the chance of an event depends upon factors that are only possible causes of it. These usually involve a spectator who did nothing in fact, but who was ready to intervene just in case the process had appeared to be headed in an undesirable direction.

[3] Consider, for example, that replacing a striker with a defender, during a soccer match, will reduce the chance of that team scoring a goal before the end of the game. Yet if the team does produce a goal then that defender may play a role in it, and is even a possible scorer of such a goal. See Dowe (2000) for more such examples.

help to bring E about. There is therefore no prospect of understanding physical causation in terms of altering chances.

With this physical understanding of causation, assumption (4) seems accurately to capture a belief among physicists that is universal, or virtually so. To suggest to physicists that one might alter the chance of an event E without changing the physical context in which E might arise is to invite incredulous stares. After all, physical events don't appear all by themselves – they are produced by earlier physical events (deterministically or otherwise). Thus a change in the tendency for an event to arise equates to a change in its tendency to be caused. It seems absurd that a fixed complete set of all the possible physical causes of E might vary in its tendency to produce E.

Assumption (4) also sheds light on the design of experiments to measure chances. Why does one run the "same" experiment over and over again? And what constitutes the "same" experiment? The idea is to determine the tendency of a given set of possible causes to produce E, and it is taken for granted that *if the possible causes of E remain the same, then so will the chance of E.* The world inevitably changes as the experiments are performed: the earth rotates, a distant country gets a new government, and so on. But these will not affect the chance of E so long as they are not possibly involved in the production of E.

III. THE EPISTEMIC VIEW

David Lewis has rightly drawn attention to the Principal Principle as central to our understanding of chance.[5] It has an interesting consequence, when added to the other Basic Assumptions, as follows:

(8) The chance of an event E is numerically equal to the rational degree of belief in E, given maximal knowledge of the causal background of E (and nothing else).

---

[4] Suppose one replaces one component of a chancy process with another that is exactly similar in all of its intrinsic physical properties. This cannot affect the chance of E, yet the new component may help to cause E.

[5] Lewis (1994) says, "Don't call any alleged feature of reality 'chance' unless you've already shown that you have something knowledge of which could constrain rational credence."

How does this follow from the Assumptions? Well, since the chance of an event is determined by the causal context, a rational person with maximal knowledge of the causal context would know what the chance is. Then, by the Principal Principle, their degree of belief would equal that chance.

Lewis himself derived something like this statement (8), regarding it as an alternative formulation of the Principal Principle. He then noted that it could be formulated as an *analysis* of chance, namely:

(8') The chance of an event E $=_{df}$ the rational degree of belief in E, given maximal knowledge of the causal background of E (and nothing else).

Thus is the epistemic view of chance. As you may know, Lewis dismissed rather quickly his version of (8') as an analysis of chance. It seems to me, however, that his dismissal was an unfortunate consequence of his overly strict metaphysical scruples. Moreover, in view of the extreme difficulty of finding a satisfactory theory of chance, a strong case can be made for relaxing metaphysical constraints, if this allows such a theory to emerge. I suggest that we first find a theory of chance that *works* (in the sense of satisfying our basic assumptions) and worry about the metaphysical commitment later.

The first thing to note is that this definition of chance (8') entails the Basic Assumptions 1-7. I will now show this. For convenience, during these arguments, I will say 'the *causal credence* of E' to refer to the rational degree of belief in E, given maximal knowledge of the causal background of E (and nothing else). The epistemic view is then just that the chance of an event is its causal credence, so I have to show that the seven Basic Assumptions hold of causal credence.

1. Since causal credences are epistemic probabilities, one can use the usual arguments (e.g. Dutch book arguments) that they must conform to the probability calculus.

2. A causal credence of 1 means that a rational person with maximal knowledge of the causal context is certain[6] that the event will occur. In other words, the event's occurrence is logically entailed[7] by a maximal description of the causes, so that the event is physically determined. In a similar way, an event with zero causal credence is (almost) determined not to occur. Events whose causal credences are between zero and one therefore have some degree of determination by their causes. The *degree of determination* by a causal context is, I think, the only reasonable interpretation of the 'strength of causal tendency' within that context.

3. The Principal Principle is a fairly straightforward consequence of the epistemic view, since PP has the structure of an authority (or expert) principle. It is, in certain circumstances, rational to defer to known epistemic authorities. If I know that some person X knows everything I do about a certain subject, and more besides, and is rational, and I know that X's degree of belief in some proposition is q, then my degree of belief should be q as well.

Now suppose that I know the causal credence q of an event E, and nothing else. In that case, I know that a rational person with maximal knowledge of the causal context of E would believe to degree q that E occurs. Now that person is an epistemic authority for me, so I should defer to her, in which case my degree of belief in the occurrence of E will be q well.

4. That the causal credence of an event is determined by its causal context follows rather trivially from the epistemic view, provided one has a sufficiently high view of rationality. All that is required is that there be precisely one maximal (true) proposition describing the causal context of the event in question. (This is easily shown. There cannot be more than one maximum, and the maximal proposition can be constructed as the conjunction of all the true propositions describing the causal context.)

---

[6] Or almost certain. I claim that the rational number (quotient) 1 represents certainty, whereas the real number (Dedekind cut or Cauchy sequence) 1 represents a degree of belief that is very close to certainty.

[7] At least on a conceptualist account of logical consequence!

5. The fifth Assumption, that outcomes of separate experiments are probabilistically independent, with respect to causal credence, but in the classical case only, is rather tricky. Its derivation is too lengthy to be given here. The 'classical' assumption I make is that the maximal description of a pair of systems <X, Y> is always equivalent to a logical conjunction of descriptions, one concerning X only and the other Y only. The derivation of independence in this classical case also makes crucial use of the fact that the causal credence for a trial is based on a *maximal* specification of the causal context for that trial, so that learning the outcomes of other trials cannot increase one's knowledge about the causes on that trial. I also need a Humean assumption that separate objects have no *a priori* epistemic relevance to each other.

6. It is clear from the above that causal credences are objective, single-case probabilities. The probability calculus then entails that in repeated identical experiments, which are probabilistically independent from (5), the long-run relative frequency for an outcome E has a high causal credence for being close to the causal credence for E. Then, since causal credences are expert probabilities, we can reasonably infer that the relative frequency for E will be close to the causal credence. Since we have inferred the relative frequency of E from a hypothesis about the causes of E, we have explained this relative frequency.[8]

7. I will not argue in detail that causal credences are empirically measurable, according to the epistemic view, since Lewis has already shown this. I will merely summarise the argument. Suppose one repeats an experiment 10,000 times, in order to estimate the causal credence of some event-type E. First, one has some initial epistemic probability distribution over the possible values of the causal credence of E. (These are the priors.) Second, one uses the independence property (5), the constant-probability property (4), and the probability calculus (1) to calculate the causal credence of the observed value of the relative frequency of E in the 10,000 trials, for each possible value of the causal credence of E. From the Principal Principle (3) these values are equal to the likelihoods,

[8] Explanation is yet another tricky topic, but I take the view that to explain an event is to infer it (to some degree, not necessarily with certainty) from a hypothesis about the causes of the event.

i.e. the epistemic probabilities of the observed relative frequency of E, given each hypothesis about the causal credence. With the priors and the likelihoods in place, one then uses Bayes's theorem to calculate the epistemic probability for each possible value of the causal credence of E, given the observed relative frequency of E.

The results of the calculation are straightforward. As long as the prior distribution over the causal credences is fairly flat, then the posterior is sharply peaked very close to the observed frequency. In other words, one is warranted in using the observed relative frequency of E as an estimate for the causal credence.

IV. OBJECTIONS TO THE EPISTEMIC VIEW

Some criticisms have been levelled at the epistemic view of chance, which I will now discuss.

1. The Problem of Stochastic Laws.

The "possible causes" referred to in the epistemic view must obviously include the dynamical properties of the system in question. But how could these properties be presented, if not in the form of a chancy law such as "Every F has chance *q* to cause G"? If the dynamical properties of an indeterministic system can only be specified using chances, then such chances are irreducible, and not degrees of belief.

My response is surely the only possible one here, namely that there must be some way to specify the dynamical properties of an indeterministic system without referring to chances. In order to be plausible, this response will clearly require the development of an account of physical laws in general, not just stochastic laws. There isn't space here to present such an account, but the following will give the rough idea.

Laws are plausibly taken to be "pushy explainers", i.e. they are *causes* of events and thus help to explain what happens. For example, an undisturbed body travels uniformly in a straight line *because* Newton's first law says it must. Laws are thus said to "govern" behaviour. Yet, of course, sticks and stones are blissfully unaware of

Newton's laws, so why do they obey them? Does God compel them to? A more plausible idea sounds rather Aristotelian, that it is simply in the *nature* of a stone to move in straight lines, unless acted upon. It is the *nature* of the stone, its being that kind of thing, that causes the motion.

What then is Newton's first law? First, it is a description of the *motion* of the stone, i.e. of its behaviour. It is a proposition. What makes this proposition a law? It is the fact that it is a logical consequence of the stone's nature. More precisely, a law is a description of the behaviour of a system that is a logical consequence of (a maximal description of) the system's nature. Physical laws thus have a certain kind of necessity, namely logical necessity relative to the nature of the system. Presumably the nature of a system is itself an essential property of it.

On this account of laws, a stochastic system will be one whose nature and initial state do not logically entail any unique history. There may, however, be a well-defined epistemic probability for each possible history, given maximal information about the system's nature and the initial state.

If this general account of laws is accepted, then this objection to the epistemic theory of chance disappears.

2. The problem of logical probabilities

A fairly common view about epistemic probability is that it is constrained only by the axioms of probability, together with the Principle of Conditionalization, $P(A|B) = P(A\&B)/P(B)$. It follows that rationality doesn't provide any probabilities by itself, but only generates probabilities from other probabilities. There are, it is said, no "logical probabilities". If this is true, then the epistemic view of chance is false.

The blanket denial of logical probabilities is certainly wrong, however. If one is careful, then one can construct rather artificial epistemic states in which probabilities are logically determined by symmetry constraints. So the difficulty here is really that there may not be *strong enough* constraints on epistemic probability, for causal credences to take precise values. I worry about this. But it isn't a very solid objection, since the dynamical natures of real stochastic systems are unknown, so we really have no idea to what extent knowledge of them will constrain epistemic probability. It is also worth noting that some objective chances could, for all we know, be intervals rather than precise numbers. Such interval chances, if they exist, can readily be explained by the epistemic theory as cases where rationality (in knowledge of the dynamical nature and initial state) fails to mandate a precise degree of belief.

V. ALTERNATIVE PROPENSITY THEORIES

The epistemic view, it might be said, is just one of many propensity accounts of probability. Why then should it be accepted, given that it faces such difficulties? The short answer to this is that the alternative propensity accounts are either empty or inadequate.

Popper's definition of propensity as a hypothetical limiting relative frequency (in an infinite set of trials) is in the latter category. It fails to entail the Principal Principle, and hence cannot explain the empirical measurability of chance. Further, the idea of defining chances as frequencies of some sort contradicts the widely held assumption (6), that stable frequencies are explained by iid chances. More recent propensity theories (Giere, Fetzer) are more modest, in declining to say what propensity is, leaving it basic and irreducible. The seven Basic Assumptions, which we derived from the epistemic view, are simply stipulated on these accounts. These accounts also have great difficulty in understanding inverse chances, which are straightforward for the epistemic view.

For this reason, I do not view these accounts as rivals to the epistemic view. It is presently the only propensity theory that is even minimally adequate.

VI. CONCLUSION

The epistemic theory provides a clear and straightforward account of chance that exactly matches all of our assumptions about it. Moreover, it is the only such account. In view of this, its metaphysical commitments are really quite light. Dynamical natures may offend Humean sensibilities, but their usefulness makes them well worth the modest cost. It is widely accepted that the axioms of probability and the conditioning principle are objective rational constraints on partial belief, so the existence of additional such constraints is hardly scandalous.

It is fairly clear that any account of objective chance that agrees with our intuitions will involve some sizable metaphysical commitment. Thus the choice before us is either to allow some such commitment, or to repudiate physical chances altogether. I urge the former, and also to recognise the strength of the epistemic theory as (presently) the only satisfactory account of chance.